Algebraic Topology I

Allen Hatcher

Copyright ©1999 by Allen Hatcher

Paper or electronic copies for noncommercial use may be made freely without explicit permission from the author.

All other rights reserved.

Table of Contents

Chapter 0. Some Underlying Geometric Notions 1

1. Homotopy and Homotopy Type 1

Deformation Retractions 2. Homotopy of Maps 3.
Homotopy Equivalent Spaces 4. Contractible Spaces 5.

2. Cell Complexes 8

Definitions and Examples 8. Subcomplexes 11.
Some Basic Constructions 11. Two Criteria for Homotopy Equivalence 14.
The Homotopy Extension Property 17.

Chapter 1. Fundamental Group and Covering Spaces . . 22

1. The Fundamental Group 27

Paths and Homotopy 27. The Fundamental Group of the Circle 32.
Induced Homomorphisms 36.

2. Van Kampen's Theorem 40

Free Products of Groups 40. The van Kampen Theorem 42.
Applications to Cell Complexes 47.

3. Covering Spaces 53

Lifting Properties 57. Classification of Covering Spaces 61.
Deck Transformations and Group Actions 65.

4. Additional Topics 77

a. Graphs and Free Groups 77. b. K(G,1) Spaces 82.
c. Graphs of Groups 87.

Chapter 2. Homology 92

1. Simplicial and Singular Homology 97

Δ-Complexes 97. Simplicial Homology 99. Singular Homology 103.
Homotopy Invariance 105. Exact Sequences and Excision 108.
The Equivalence of Simplicial and Singular Homology 122.

2. Computations and Applications 128

Degree 128. Cellular Homology 131. Euler Characteristic 140.
Split Exact Sequences 141. Mayer-Vietoris Sequences 143.
Homology with Coefficients 145.

3. The Formal Viewpoint 151

Axioms for Homology 151. Categories and Functors 154.

4. Additional Topics 159

a. H_1 and π_1 159. b. Theorems of Brouwer and Hopf 161.
c. Simplicial Approximation and the Lefschetz Fixed Point Theorem 166.

Chapter 3. Cohomology 173

1. Cohomology Groups 174

The Universal Coefficient Theorem 175. Cohomology of Spaces 182.

2. Cup Product 189

The Cohomology Ring 193. External Cup Product 196.
Spaces with Polynomial Cohomology 201.

3. Poincaré Duality 213

Orientations 216. Cap Product 218. Proof of Poincaré Duality 221.
Cup Product and Poincaré Duality 228. Other Forms of Duality 232.

4. Additional Topics 241

a. The Universal Coefficient Theorem for Homology 241.
b. The Künneth Formula 250. c. H-Spaces and Hopf Algebras 261.
d. The Cohomology of SO(n) 273. e. Bockstein Homomorphisms 283.
f. Limits 292. g. More About Ext 300. h. Transfer Homomorphisms 305.
i. Local Coefficients 313.

Chapter 4. Homotopy Theory 324

1. Homotopy Groups 325

The Long Exact Sequence 329. Whitehead's Theorem 333.
The Hurewicz Theorem 335. Eilenberg-MacLane Spaces 344.

2. Homotopy Properties of CW Complexes 357

Cellular Approximation 357. Cellular Models 359.
Excision for Homotopy Groups 366. Stable Homotopy Groups 373.

3. Fibrations 382

The Homotopy Lifting Property 382. Fiber Bundles 384.
Path Fibrations and Loopspaces 393. Postnikov Towers 397.
Obstruction Theory 405.

4. Additional Topics 412

a. Basepoints and Homotopy 412. b. The Hopf Invariant 415.
c. Minimal Cell Structures 417. d. Cohomology of Fiber Bundles 420.
e. Cohomology Theories and Ω-Spectra 437.
f. Spectra and Homology Theories 448. g. Eckmann-Hilton Duality 452.
h. Stable Splittings of Spaces 458. i. The Loopspace of a Suspension 461.
j. Symmetric Products and the Dold-Thom Theorem 467.
k. Steenrod Squares and Powers 479.

Appendix 513

Topology of Cell Complexes 513.
The Compact-Open Topology 523.

Bibliography 525

Index 530

Preface

This book was written to be a readable introduction to Algebraic Topology with rather broad coverage of the subject. Our viewpoint is quite classical in spirit, and stays largely within the confines of pure Algebraic Topology. In a sense, the book could have been written thirty years ago since virtually all its content is at least that old. However, the passage of the intervening years has helped clarify what the most important results and techniques are. For example, CW complexes have proved over time to be the most natural class of spaces for Algebraic Topology, so they are emphasized here much more than in the books of an earlier generation. This emphasis also illustrates the book's general slant towards geometric, rather than algebraic, aspects of the subject. The geometry of Algebraic Topology is so pretty, it would seem a pity to slight it and to miss all the intuition which it provides. At deeper levels, algebra becomes increasingly important, so for the sake of balance it seems only fair to emphasize geometry at the beginning.

Let us say something about the organization of the book. At the elementary level, Algebraic Topology divides naturally into two channels, with the broad topic of Homotopy on the one side and Homology on the other. We have divided this material into four chapters, roughly according to increasing sophistication, with Homotopy split between Chapters 1 and 4, and Homology and its mirror variant Cohomology in Chapters 2 and 3. These four chapters do not have to be read in this order, however. One could begin with Homology and perhaps continue on with Cohomology before turning to Homotopy. In the other direction, one could postpone Homology and Cohomology until after parts of Chapter 4. However, we have not pushed this latter approach to its natural limit, in which Homology and Cohomology arise just as branches of Homotopy Theory. Appealing as this approach is from a strictly logical point of view, it places more demands on the reader, and since readability is one of our first priorities, we have delayed introducing this unifying viewpoint until near the end of the book.

There is also a preliminary Chapter 0 which introduces some of the basic geometric concepts and constructions which play a central role in both the homological and homotopical sides of the subject.

Each of the four main chapters concludes with a section called Additional Topics which the reader can sample at will, independent of the basic core of the book contained in the other sections. Many of these extra topics are in fact rather important in the overall scheme of Algebraic Topology, though they might not fit into the time constraints of a first course. Altogether, these Additional Topics amount to nearly half the book, and we have included them both to make the book more comprehensive and to give the reader

who takes the time to delve into them a more substantial sample of the true richness and beauty of the subject.

A sequel Volume II is being written as well. This will begin with the important but somewhat more sophisticated topic of spectral sequences. It was very tempting to include something about this marvelous tool in the present volume, but spectral sequences are such a big topic that it seemed best to start with them afresh in a new volume. There is also a third book in progress, on vector bundles, characteristic classes, and K-theory. This will be largely independent of Volume II and also of much of Volume I. In the present book this future volume is referred to as [VBKT], its provisional title being 'Vector Bundles and K-Theory.'

In terms of prerequisites, this first volume assumes the reader has some familiarity with the content of the standard undergraduate courses in algebra and point-set topology. One topic which is not always a part of a first point-set topology course but which is quite important for Algebraic Topology is quotient spaces. A good place to read about this is Chapter 4 of the textbook of Armstrong, where quotient spaces are called 'identification spaces.'

It is the author's intention to keep this book available online permanently, as well as publish it in the traditional manner for those who want the convenience of a bound copy. With the electronic version it will be possible to continue making revisions and additions, so comments and suggestions from readers will always be welcome. The web address is:

http://math.cornell.edu/~hatcher

One can also find here the parts of the other two books which are currently available.

Chapter 0
Basic Geometric Notions

The aim of this short preliminary chapter is to introduce a few fundamental geometric ideas which recur frequently throughout algebraic topology. The exposition is somewhat informal, with no theorems or proofs until the last couple pages, and it should be read in this informal spirit, skipping bits here and there. In fact, this whole chapter could be skipped now, to be referred back to later for basic definitions.

Standard Notations

$\mathbb{R}^n$, n-dimensional Euclidean space, with real coordinates

$\mathbb{C}^n$, complex n-space

$I = [0, 1]$, the unit interval

S^n, the unit sphere in $\mathbb{R}^{n+1}$, all vectors of length 1

D^n, the unit disk or ball in $\mathbb{R}^n$, all vectors of length ≤ 1

$\partial D^n = S^{n-1}$, the boundary of the n-disk

$\mathbb{1}$, the identity function from a set to itself

$\amalg$, disjoint union

$\approx$, isomorphism

$\mathbb{Z}_n$, the integers mod n

To avoid overusing the word 'continuous' we adopt the convention that maps between spaces are always assumed to be continuous unless otherwise stated.

1. Homotopy and Homotopy Type.

One of the main ideas of algebraic topology is to consider two spaces to be equivalent if they have 'the same shape' in a sense much weaker than simply being homeomorphic. To take an everyday example, the letters of the alphabet can be written either as unions of finitely many straight and curved line segments, or in thickened forms which are compact subsurfaces of the plane.

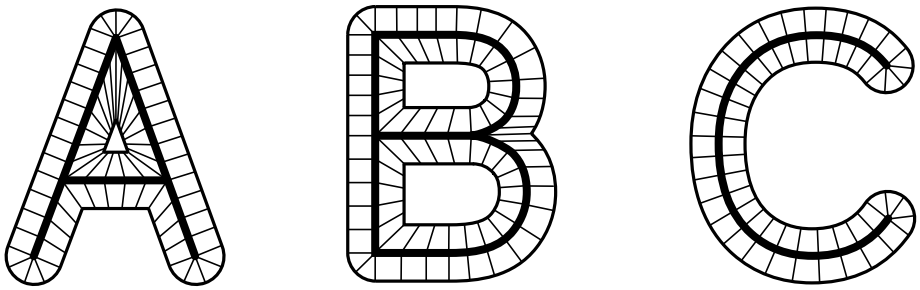

In each case the thin letter is a subspace of the thick letter, and we can continuously shrink the thick letter to the thin one. A nice way to do this is to decompose a thick letter, call it $\mathbf{X}$, into line segments connecting each point on the outer boundary of $\mathbf{X}$ to a unique point of the thin subletter X, as indicated in the figure. Then we can shrink $\mathbf{X}$ to X by sliding each point of $\mathbf{X} - \mathsf{X}$ into X along the line segment which contains it. Points which are already in X do not move.

This shrinking process we can think of as taking place during a time interval $0 \le t \le 1$, and then it defines a family of functions $f_t : \mathbf{X} \to \mathbf{X}$ parametrized by $t \in I = [0,1]$, where $f_t(x)$ is the point to which a given point $x \in \mathbf{X}$ has moved at time t. Naturally we would like $f_t(x)$ to depend continuously on both t and x, and this will be true if we have each $x \in \mathbf{X} - \mathsf{X}$ move along its line segment at constant speed so as to reach its image point in X at time $t = 1$, while points $x \in \mathsf{X}$ are stationary as remarked earlier.

Deformation Retractions.

These examples lead to the following general definition. A **deformation retraction** of a space X onto a subspace A is a family of maps $f_t : X \to X$, $t \in I$, such that $f_0 = \mathbb{1}$, $f_1(X) \subset A$, and $f_t \,|\, A = \mathbb{1}$ for all t. The family f_t should be continuous in the sense that the associated map $X \times I \to X$, $(x,t) \mapsto f_t(x)$, is continuous.

Here are a few more examples, in which the deformation retraction f_t can again be constructed by sliding along line segments.

Example 1. A disk with two smaller open subdisks removed deformation retracts onto several different subspaces.

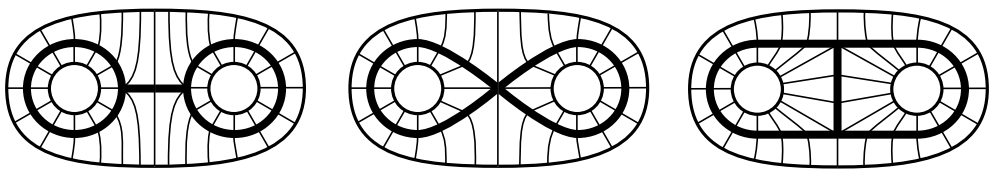

Example 2. The Möbius band deformation retracts onto its core circle.

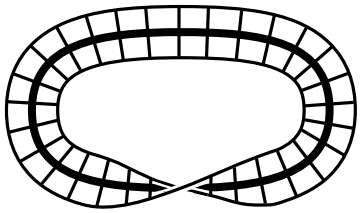

In all the preceding examples the structure which gives rise to the deformation retraction can be described by means of the following definition. For a map $f: X \to Y$, the **mapping cylinder** M_f is the quotient space of the disjoint union $(X \times I) \amalg Y$ obtained by identifying each $(x, 1) \in X \times I$ with $f(x) \in Y$.

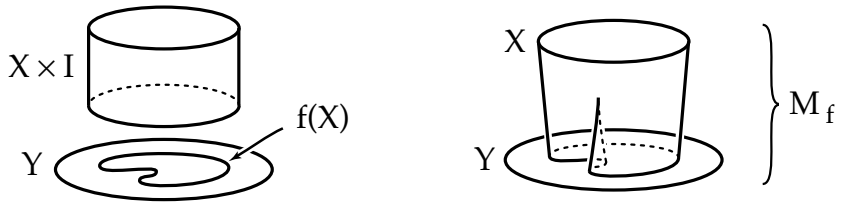

In the letter examples, the space X is the outer boundary of the thick letter, Y is the thin letter, and the map $f: X \to Y$ sends the outer endpoint of each line segment to its inner endpoint. A similar description applies to Examples 1 and 2. Then it is a general fact that a mapping cylinder M_f deformation retracts to the subspace Y by sliding each point (x, t) along the segment $\{x\} \times I \subset M_f$ to the endpoint $f(x) \in Y$.

Not all deformation retractions arise in this way from mapping cylinders, however. To give a fairly trivial example, one could modify one of the letter examples by deleting the boundary of the thick letter, or any subset of the boundary, and the remainder of the thick letter would still deformation retract to the thin letter. But more subtle deformation retractions can easily be constructed. For example, the thick **X** deformation retracts to the thin X which in turn deformation retracts to the point of intersection of its two crossbars. The net result is a deformation retraction of **X** onto a point, during which certain pairs of points follow paths that merge before reaching their final destination.

Later in this section we will describe a considerably more complicated example of a deformation retraction, involving the so-called 'house with two rooms.'

Homotopy of Maps.

A deformation retraction $f_t: X \to X$ is a special case of the general notion of a **homotopy**, which is simply any family of maps $f_t: X \to Y$, $0 \le t \le 1$, such that the associated map $F: X \times I \to Y$ given by $F(x, t) = f_t(x)$ is continuous. One says that two

maps $f_0, f_1 : X \to Y$ are **homotopic** if there exists a homotopy f_t connecting them, and one writes $f_0 \simeq f_1$.

In these terms, a deformation retraction is a homotopy from $f_0 = \mathbb{1}$ to a **retraction**, a map $f_1 : X \to X$ such that $f_1(X) \subset A$ and $f_1 \big| A = \mathbb{1}$. One could equally well regard a retraction as a map $X \to A$ restricting to the identity on the subspace $A \subset X$. Not all retractions arise from deformation retractions. For example, every space X retracts onto any point $x_0 \in X$ via the map sending all of X to x_0. But a space which deformation retracts onto a point must certainly be path-connected, since a deformation retraction of X to a point x_0 gives a path joining each $x \in X$ to x_0. It is less trivial to show there are path-connected spaces which do not deformation retract onto a point. One would expect this to be the case for the letters 'with holes,' A, B, D, O, P, Q, R. In Chapter 1 we will develop techniques to prove this.

A homotopy $f_t : X \to X$ which gives a deformation retraction of X onto a subspace A has the additional property that $f_t \big| A = \mathbb{1}$ for all t. In general, a homotopy $f_t : X \to Y$ whose restriction to a subspace $A \subset X$ is independent of t is called a homotopy **rel** A, a homotopy *relative to* A. Thus, a deformation retraction of X onto A is equivalent to a retraction of X onto A which is homotopic rel A to the identity map of X.

Homotopy Equivalent Spaces.

If a space X deformation retracts onto a subspace A via $f_t : X \to X$, then if $r : X \to A$ denotes the resulting retraction and $i : A \to X$ the inclusion, we have $ri = \mathbb{1}$ and $ir \simeq \mathbb{1}$, the latter homotopy being given by f_t. Generalizing this situation, a map $f : X \to Y$ is called a **homotopy equivalence** if there is a map $g : Y \to X$ such that $fg \simeq \mathbb{1}$ and $gf \simeq \mathbb{1}$. One says the spaces X and Y are **homotopy equivalent**, or that they have the same **homotopy type**, and one writes $X \simeq Y$. It is easy to check this is an equivalence relation, in contrast with the nonsymmetric notion of deformation retraction. For example, the three graphs in Example 1 above are all homotopy equivalent since they are deformation retracts of the same space, but none of the three is a deformation retract of any other.

It is true in general that two spaces X and Y are homotopy equivalent if and only if there exists a third space Z containing both X and Y as deformation retracts. For the less trivial implication one can in fact take Z to be the mapping cylinder M_f of any homotopy equivalence $f : X \to Y$. We observed above that M_f deformation retracts to Y, so what needs to be proved is that M_f also deformation retracts to its other end X if f is a homotopy equivalence. This is shown at the end of this chapter.

One could say that the main task of algebraic topology is to provide techniques to decide whether two given spaces are homotopy equivalent.

Contractible Spaces.

A space having the homotopy type of a point is called **contractible**. This amounts to requiring that the identity map of the space be **nullhomotopic**, i.e., homotopic to a constant map. In general, this is slightly weaker than saying the space deformation retracts to a point; see the exercises for an example distinguishing these two notions.

Example: The house with two rooms. This two-dimensional subspace of $\mathbb{R}^3$ shown in the figure below is contractible but not in any obvious way. To build this space, start with a box divided into two chambers by a horizontal rectangle R, where by a 'rectangle' we mean not just the four edges of a rectangle but its interior too. Access to the two chambers from outside the box is provided by two vertical tunnels. The upper tunnel is made by punching out a square from the top of the box and another square directly below it from R, then inserting four vertical rectangles, the walls of the tunnel. This tunnel allows entry to the lower chamber from outside the box. The lower tunnel is formed in similar fashion, providing entry to the upper chamber. Finally, two vertical rectangles are inserted to form 'support walls' for the two tunnels.

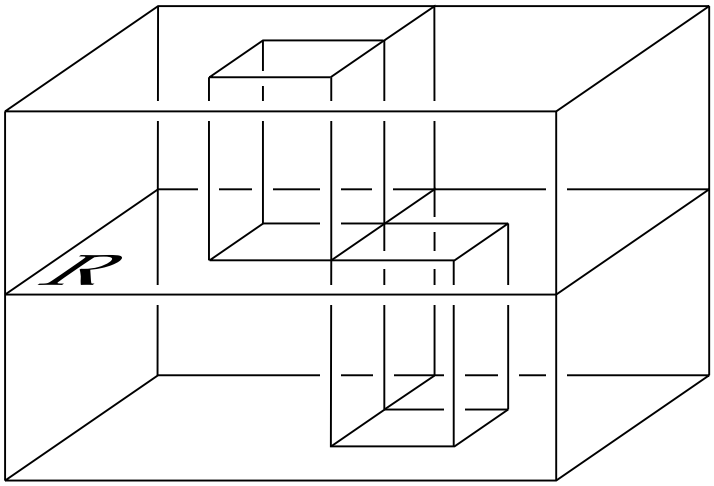

The resulting space X thus consists of three horizontal pieces each homeomorphic to an annulus $S^1 \times I$, plus all the vertical rectangles which form the walls of the two chambers: the exterior walls, the walls of the tunnels, and the two support walls.

To see that X is contractible, consider a closed ε-neighborhood $N(X)$ of X. This clearly deformation retracts onto X if ε is sufficiently small. In fact, $N(X)$ is the mapping

cylinder of a map from the boundary surface of $N(X)$ to X. Less obvious is the fact that $N(X)$ is homeomorphic to D^3, the unit ball in $\mathbb{R}^3$. To see this, imagine forming $N(X)$ from a ball of clay by pushing a finger into the ball to create the upper tunnel, then gradually hollowing out the lower chamber, and similarly pushing a finger in to create the lower tunnel and hollowing out the upper chamber. Mathematically, this process gives a family of embeddings $h_t : D^3 \to \mathbb{R}^3$ starting with the usual inclusion $D^3 \hookrightarrow \mathbb{R}^3$ and ending with a homeomorphism onto $N(X)$.

Thus we have $X \simeq N(X) = D^3 \simeq \textit{point}$, so X is contractible since homotopy equivalence is an equivalence relation.

In fact, X deformation retracts to a point. For if f_t is a deformation retraction of the ball $N(X)$ to a point $x_0 \in X$ and if $r : N(X) \to X$ is a retraction, for example the end result of a deformation retraction of $N(X)$ to X, then the restriction of the composition rf_t to X is a deformation retraction of X to x_0. However, it is not easy to see exactly what this deformation retraction looks like! A slightly easier test of geometric visualization is to find a nullhomotopy in X of the loop formed by a horizontal cross-section of one of the tunnels. We leave this as a puzzle for the reader.

Exercises

1. Show that the torus with one point deleted deformation retracts onto a graph consisting of two circles intersecting in a point, namely, longitude and meridian circles of the torus.

2. Show that S^{n-1} is a deformation retract of $\mathbb{R}^n - \{0\}$.

3. Show that the composition of homotopy equivalences $X \to Y$ and $Y \to Z$ is a homotopy equivalence $X \to Z$. Deduce that homotopy equivalence is an equivalence relation.

4. A **deformation retraction in the weak sense** of a space X onto a subspace A is a homotopy $f_t : X \to X$ such that $f_0 = 1\!\!1$, $f_1(X) \subset A$, and $f_t(A) \subset A$ for all t. Given this data, show that the inclusion $A \hookrightarrow X$ is a homotopy equivalence.

5. (a) Let X be the subspace of $\mathbb{R}^2$ consisting of the horizontal segment $[0,1] \times \{0\}$ together with all the vertical segments $\{r\} \times [0, 1-r]$ for r a rational number in $[0,1]$. Show that X deformation retracts to any point in the segment $[0,1] \times \{0\}$, but not to any other point.

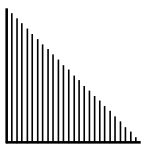

(b) Let Y be the subspace of $\mathbb{R}^2$ which is the union of an infinite number of copies of X arranged as in the figure below, and let Z be the zigzag subspace homeomorphic to $\mathbb{R}$ indicated by the heavier line. Show that there is a deformation retraction in the weak sense (see the preceding problem) of Y onto Z, but no true deformation retraction. Deduce from this that Y is contractible but does not deformation retract onto any point.

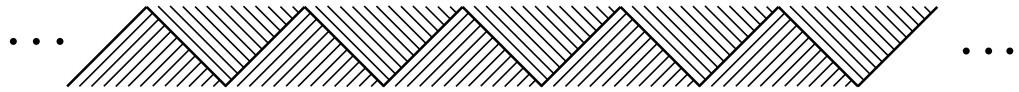

6. For $n > 2$, construct an n-room analog of the house with two rooms.

7. Show that a retract of a contractible space is contractible.

8. Show that a space X is contractible iff every map $f : X \to Y$, for arbitrary Y, is nullhomotopic. Similarly, show X is contractible iff every map $f : Y \to X$ is nullhomotopic.

9. Show that a map $f : X \to Y$ is a homotopy equivalence if there exist maps $g, h : Y \to X$ such that $fg \simeq \mathbb{1}$ and $hf \simeq \mathbb{1}$. More generally, show that f is a homotopy equivalence if fg and hf are homotopy equivalences.

10. Show that a homotopy equivalence $f : X \to Y$ induces a bijection between the set of path-components of X and the set of path-components of Y, and that f restricts to a homotopy equivalence from each path-component of X to the corresponding path-component of Y.

11. Show that any two deformation retractions r_t^0 and r_t^1 of a space X onto a subspace A can be joined by a continuous family of deformation retractions r_t^s, $0 \le s \le 1$, of X onto A, where continuity means that the map $X \times I \times I \to X$, $(x, s, t) \mapsto r_t^s(x)$, is continuous.

2. Cell Complexes.

In this section we introduce the class of spaces which is by far the most convenient for purposes of algebraic topology.

Definitions and Examples

A familiar way of constructing the torus $S^1 \times S^1$ is by identifying opposite sides of a square. More generally, an orientable surface M_g of genus g can be constructed from a $4g$-sided polygon by identifying pairs of edges, as shown in the figure for the cases $g = 1, 2, 3$.

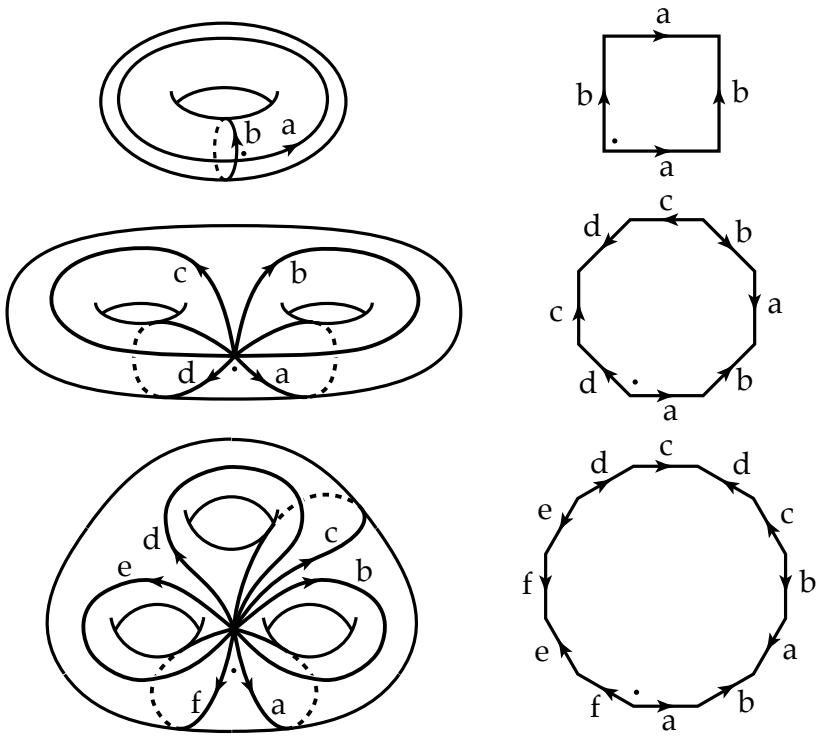

The $4g$ edges of the polygon become a union of $2g$ circles in the surface, all intersecting in a single point. One can think of the interior of the polygon as an open disk, or **2-cell**, attached to the union of these circles. One can also regard the union of the circles as being obtained from a point, their common point of intersection, by attaching $2g$ open arcs, or **1-cells**. Thus the surface can be built up in stages: Start with a point, attach 1-cells to this, then attach a 2-cell.

A natural generalization of this is to construct a space by the following procedure:

(1) Start with a discrete set X^0, whose points are regarded as 0-cells.

(2) Inductively, form the $\boldsymbol{n}$**-skeleton** X^n from X^{n-1} by attaching n-cells e^n_α via maps $\varphi_\alpha : S^{n-1} \to X^{n-1}$. That is, X^n is the quotient space of the disjoint union of X^{n-1}

with a collection of n-disks D^n_α under the identifications $x \sim \varphi_\alpha(x)$ for $x \in \partial D^n_\alpha$. Thus as a set, $X^n = X^{n-1} \coprod_\alpha e^n_\alpha$.

(3) One can either stop this inductive process at a finite stage, setting $X = X^n$ for some $n < \infty$, or one can continue indefinitely, setting $X = \bigcup_n X^n$. In the latter case X is given the weak topology: A set $A \subset X$ is open (or closed) iff $A \cap X^n$ is open (or closed) in X^n for each n.

A space X constructed in this way is called a **cell complex**, or more classically, a **CW complex**. The explanation of the letters 'CW' is given in the Appendix, where a number of basic topological properties of cell complexes are proved. The reader who wonders about various point-set topological questions which lurk in the background of the following discussion should consult the Appendix for details.

Example 1. A one-dimensional cell complex $X = X^1$ is what is called a **graph** in algebraic topology. It consists of vertices (the 0-cells) to which edges (the 1-cells) are attached. The two ends of an edge can be attached to the same vertex.

Example 2. The house with two rooms, pictured in the previous section, has a visually obvious two-dimensional cell complex structure. The 0-cells are the vertices where three or more of the depicted edges meet, the 1-cells are the edges connecting these vertices, and the 2-cells are rectangles except for the floors and ceilings of the rooms, which contribute 2-cells of a more complicated shape. If one counts up, one finds there are 29 0-cells, 51 1-cells, and 23 2-cells, with the alternating sum $29 - 51 + 23$ equal to 1. This is the **Euler characteristic**, which for a cell complex with finitely many cells is defined to be the number of even-dimensional cells minus the number of odd-dimensional cells. As we shall show in Chapter 2, the Euler characteristic of a cell complex depends only on its homotopy type, so the fact that the house with two rooms has the homotopy type of a point implies that its Euler characteristic must be 1, no matter how it is divided up into cells.

Example 3. The sphere S^n has the structure of a cell complex with just two cells, e^0 and e^n, the n-cell being attached by the constant map $S^{n-1} \to e^0$. This is equivalent to regarding S^n as the quotient $D^n/\partial D^n$.

Example 4. $\mathbb{R}\mathrm{P}^n$, **real projective n-space**, is defined to be the space of lines through the origin in $\mathbb{R}^{n+1}$. Each such line is determined by a non-zero vector in $\mathbb{R}^{n+1}$, unique up to scalar multiplication, and $\mathbb{R}\mathrm{P}^n$ is topologized as the quotient space of $\mathbb{R}^{n+1} - \{0\}$ under the equivalence relation $v \sim \lambda v$ for scalars $\lambda \neq 0$. We can restrict to vectors of length 1, so $\mathbb{R}\mathrm{P}^n$ is also the quotient space $S^n/(v \sim -v)$, the sphere with antipodal points

identified. This is equivalent to saying that $\mathbb{R}\mathrm{P}^n$ is the quotient space of a hemisphere D^n with antipodal points of ∂D^n identified. Since ∂D^n with antipodal points identified is just $\mathbb{R}\mathrm{P}^{n-1}$, we see that $\mathbb{R}\mathrm{P}^n$ is obtained from $\mathbb{R}\mathrm{P}^{n-1}$ by attaching an n-cell, with attaching map the quotient projection $S^{n-1} \to \mathbb{R}\mathrm{P}^{n-1}$. Hence, by induction on n, $\mathbb{R}\mathrm{P}^n$ has a cell complex structure $e^0 \cup e^1 \cup \cdots \cup e^n$ with one cell e^i in each dimension $i \leq n$.

Example 5. Since $\mathbb{R}\mathrm{P}^n$ is obtained from $\mathbb{R}\mathrm{P}^{n-1}$ by attaching an n-cell, the infinite union $\mathbb{R}\mathrm{P}^\infty = \bigcup_n \mathbb{R}\mathrm{P}^n$ becomes a cell complex with one cell in each dimension. We can view $\mathbb{R}\mathrm{P}^\infty$ as the space of lines through the origin in $\mathbb{R}^\infty = \bigcup_n \mathbb{R}^n$.

Example 6. $\mathbb{C}\mathrm{P}^n$, **complex projective n-space**, is the space of complex lines through the origin in $\mathbb{C}^{n+1}$, i.e., one-dimensional vector subspaces of $\mathbb{C}^{n+1}$. As in the case of $\mathbb{R}\mathrm{P}^n$, each line is determined by a non-zero vector in $\mathbb{C}^{n+1}$, unique up to scalar multiplication, and $\mathbb{C}\mathrm{P}^n$ is topologized as the quotient space of $\mathbb{C}^{n+1} - \{0\}$ under the equivalence relation $v \sim \lambda v$ for $\lambda \neq 0$. Equivalently, this is the quotient of the unit sphere $S^{2n+1} \subset \mathbb{C}^{n+1}$ with $v \sim \lambda v$ for $|\lambda| = 1$. It is also possible to obtain $\mathbb{C}\mathrm{P}^n$ as a quotient space of the disk D^{2n} under the identifications $v \sim \lambda v$ for $v \in \partial D^{2n}$, in the following way. The vectors in $S^{2n+1} \subset \mathbb{C}^{n+1}$ with last coordinate real and non-negative are precisely the vectors of the form $(w, \sqrt{1 - |w|^2}\,) \in \mathbb{C}^n \times \mathbb{C}$ with $|w| \leq 1$, so such vectors form a disk D^{2n}_+ bounded by the sphere $S^{2n-1} \subset S^{2n+1}$, the vectors $(w, 0)$. Each vector in S^{2n+1} is equivalent under the identifications $v \sim \lambda v$ to a vector in D^{2n}_+, and the latter vector is unique if its last coordinate is nonzero. If the last coordinate is zero, we have just the identifications $v \sim \lambda v$ for $v \in S^{2n-1}$.

From this description of $\mathbb{C}\mathrm{P}^n$ as the quotient of D^{2n}_+ under the identifications $v \sim \lambda v$ for $v \in S^{2n-1}$ it follows that $\mathbb{C}\mathrm{P}^n$ is obtained from $\mathbb{C}\mathrm{P}^{n-1}$ by attaching a cell e^{2n} via the quotient map $S^{2n-1} \to \mathbb{C}\mathrm{P}^{n-1}$. So by induction on n we obtain a cell structure $\mathbb{C}\mathrm{P}^n = e^0 \cup e^2 \cup \cdots \cup e^{2n}$ with cells only in even dimensions. Similarly, $\mathbb{C}\mathrm{P}^\infty$ has a cell structure with one cell in each even dimension.

Each cell e^n_α in a cell complex X has a **characteristic map** $\Phi_\alpha : D^n_\alpha \to X$ which extends the attaching map φ_α and is a homeomorphism from the interior of D^n_α onto e^n_α. Namely, we can take Φ_α to be the composition $D^n_\alpha \hookrightarrow X^{n-1} \coprod_\alpha D^n_\alpha \to X^n \hookrightarrow X$ where the middle map is the quotient map defining X^n. For example, in the canonical cell structure on S^n described in Example 3, a characteristic map for the n-cell is the quotient map $D^n \to S^n$ collapsing ∂D^n to a point. For $\mathbb{R}\mathrm{P}^n$ a characteristic for the k-cell, $k \leq n$, is the quotient map $D^k \to \mathbb{R}\mathrm{P}^k \hookrightarrow \mathbb{R}\mathrm{P}^n$ identifying antipodal points of ∂D^k, and similarly for $\mathbb{C}\mathrm{P}^n$.

Subcomplexes.

A **subcomplex** of a cell complex X is a closed subspace $A \subset X$ which is a union of cells of X. Since A is closed, the characteristic map of each cell in A has image contained in A, and in particular the image of the attaching map of each cell in A is contained in A, so A is a cell complex in its own right.

As a simple example, each skeleton X^n of a cell complex X is a subcomplex. Particular cases of this are the subcomplexes $\mathbb{R}\mathrm{P}^k \subset \mathbb{R}\mathrm{P}^n$ and $\mathbb{C}\mathrm{P}^k \subset \mathbb{C}\mathrm{P}^n$ for $k \leq n$. These are in fact the only subcomplexes of $\mathbb{R}\mathrm{P}^n$ and $\mathbb{C}\mathrm{P}^n$.

There are natural inclusions $S^0 \subset S^1 \subset \cdots \subset S^n$, but these subspheres are not subcomplexes of S^n in its usual cell structure, the one given in Example 3 above. This can be achieved however by giving S^n a different cell structure having more cells, by regarding each S^k as being obtained inductively from the equatorial S^{k-1} by attaching two k-cells, the components of $S^k - S^{k-1}$. The infinite-dimensional sphere $S^\infty = \bigcup_n S^n$ then becomes a cell complex too. Note that the two-to-one quotient map $S^\infty \to \mathbb{R}\mathrm{P}^\infty$ which identifies antipodal points of S^∞ also identifies the two n-cells of S^∞ to the single n-cell of $\mathbb{R}\mathrm{P}^\infty$, for each n.

In the examples of cell complexes given so far, the closure of each cell is a subcomplex, and more generally the closure of any collection of cells is a subcomplex. This is in fact true for most naturally-arising cell structures, but it is not true in general. For example, if we start with S^1 with its canonical cell structure consisting of a 0-cell and a 1-cell, and attach to this a 2-cell by a map $S^1 \to S^1$ whose image is some nontrivial subarc of S^1, then the closure of the 2-cell is not a subcomplex, since it contains only a part of the 1-cell.

Some Basic Constructions.

Cell complexes have a very nice mixture of rigidity and flexibility, with enough rigidity to allow many arguments to proceed in a combinatorial cell-by-cell fashion, and enough flexibility to allow many natural constructions to be performed on them. Here are some of those constructions.

Products. If X and Y are cell complexes, then $X \times Y$ has the structure of a cell complex with cells the products $e^m_\alpha \times e^n_\beta$ where e^m_α ranges over the cells of X and e^n_β ranges over the cells of Y. For example, the cell structure on the torus $S^1 \times S^1$ described at the beginning of this section is obtained in this way from the standard cell structure on S^1. In the general case there is one small complication, however: The topology on $X \times Y$ as a cell complex is sometimes slightly weaker than the product topology, with more open sets

than the product topology has, though the two topologies coincide if either X or Y has only finitely many cells, or if both X and Y have countably many cells. This is explained in the Appendix. In practice this subtle point of point-set topology rarely causes problems.

Quotients. If (X, A) is a **CW pair** consisting of a cell complex X and a subcomplex A, then the quotient space X/A inherits a natural cell complex structure from X. The cells of X/A are the cells of $X - A$ plus one new 0-cell, the image of A in X/A. For a cell e^n_α of $X - A$ attached by $\varphi_\alpha : S^{n-1} \to X^{n-1}$, the attaching map for the corresponding cell in X/A is the composition $S^{n-1} \to X^{n-1} \to X^{n-1}/A^{n-1}$.

For example, if we give S^{n-1} any cell structure and build D^n from S^{n-1} by attaching an n-cell, then the quotient D^n/S^{n-1} is S^n with its usual cell structure. As another example, take X to be a closed orientable surface with the cell structure described at the beginning of this section, with a single 2-cell, and let A be the complement of this 2-cell, the 1-skeleton of X. Then X/A has a cell structure consisting of a 0-cell with a 2-cell attached, and there is only one way to attach a cell to a 0-cell, by the constant map, so X/A is S^2.

Suspension. For a space X, the **suspension** SX is the quotient of $X \times I$ obtained by collapsing $X \times \{0\}$ to one point and $X \times \{1\}$ to another point. The motivating example is $X = S^n$, when $SX = S^{n+1}$ with the two 'suspension points' at the north and south poles of S^{n+1}, the points $(0, \cdots, 0, \pm 1)$. One can regard SX as a double cone on X, the union of two copies of the **cone** $CX = (X \times I)/(X \times \{0\})$. If X is a CW complex, so are SX and CX as quotients of $X \times I$ with its product cell structure, I being given the standard cell structure of two 0-cells joined by a 1-cell.

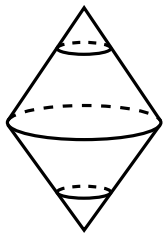

Suspension becomes increasingly important the farther one goes into algebraic topology, though why this should be so is certainly not evident in advance. One especially useful property of suspension is that not only spaces but also maps can be suspended. Namely, a map $f : X \to Y$ suspends to $Sf : SX \to SY$, the quotient map of $f \times \mathbb{1} : X \times I \to Y \times I$.

Join. The cone CX is the union of all line segments joining points of X to an external vertex, and similarly the suspension SX is the union of all line segments joining points of X to two external vertices. More generally, given X and a second space Y, the **join** $X * Y$ is defined to be the quotient space of $X \times Y \times I$ under the identifications $(x, y_1, 0) \sim (x, y_2, 0)$ and $(x_1, y, 1) \sim (x_2, y, 1)$. Thus we are collapsing the subspace $X \times Y \times \{0\}$ to X and $X \times Y \times \{1\}$ to Y. For example, if X and Y are both closed intervals, then we are collapsing two opposite faces of a cube onto line segments so that it becomes a tetrahedron:

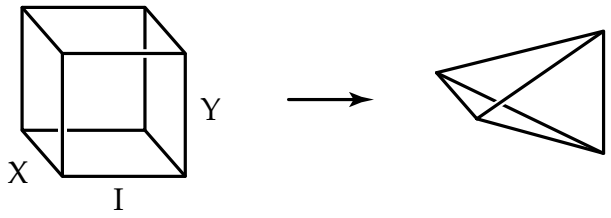

In the general case, $X * Y$ contains copies of X and Y at its two 'ends,' and every other point (x, y, t) in $X * Y$ is on a unique line segment joining the point $x \in X \subset X * Y$ to the point $y \in Y \subset X * Y$, the segment obtained by fixing x and y and letting the coordinate t in (x, y, t) vary.

A nice way to write points of $X * Y$ is as formal linear combinations $t_1 x + t_2 y$ with $0 \leq t_i \leq 1$ and $t_1 + t_2 = 1$, subject to the rules $0x + 1y = y$ and $1x + 0y = x$ which correspond exactly to the identifications that define $X * Y$. In much the same way, an iterated join $X_1 * \cdots * X_n$ can be regarded as the space of formal combinations $t_1 x_1 + \cdots + t_n x_n$ with $0 \leq t_i \leq 1$ and $t_1 + \cdots + t_n = 1$, with the convention that terms $0t_i$ can be omitted. This viewpoint makes it easy to see that the join operation is associative. A very special case which plays a central role in algebraic topology is when each X_i is just a point. For example, the join of two points is a line segment, the join of three points is a triangle, and the join of four points is a tetrahedron. The join of n points is a convex polyhedron of dimension $n - 1$ called a **simplex**. Concretely, if the n points are the n standard basis vectors for $\mathbb{R}^n$, then their join is the subspace $\Delta^{n-1} = \{ (t_1, \cdots, t_n) \in \mathbb{R}^n \mid t_1 + \cdots + t_n = 1 \text{ and } t_i \geq 0 \}$.

Another interesting example is when each X_i is S^0, two points. If we take the two points of X_i to be the two unit vectors along the i^{th} coordinate axis in $\mathbb{R}^n$, then the join $X_1 * \cdots * X_n$ is the union of 2^n copies of the simplex Δ^{n-1}, and radial projection from the origin gives a homeomorphism between $X_1 * \cdots * X_n$ and S^{n-1}.

If X and Y are CW complexes, then there is a natural CW structure on $X * Y$ having the subspaces X and Y as subcomplexes, with the remaining cells being the product cells of $X \times Y \times (0, 1)$. As usual with products, the CW topology on $X * Y$ may be weaker than the quotient of the product topology on $X \times Y \times I$.

Wedge Sum. This is a fairly trivial but still useful operation. Given spaces X and Y with chosen points $x_0 \in X$ and $y_0 \in Y$, then the **wedge sum** $X \vee Y$ is the quotient of the disjoint union $X \amalg Y$ obtained by identifying x_0 and y_0 to a single point. For example, $S^1 \vee S^1$ is a figure 8, two circles touching at a point. More generally one could form the wedge sum $\bigvee_\alpha X_\alpha$ of an arbitrary collection of spaces X_α by starting with the disjoint union $\coprod_\alpha X_\alpha$ and identifying points $x_\alpha \in X_\alpha$ to a single point. In case the spaces X_α

are cell complexes and the points x_α are 0-cells, then $\bigvee_\alpha X_\alpha$ is a cell complex since it is obtained from the cell complex $\coprod_\alpha X_\alpha$ by collapsing a subcomplex to a point.

For any cell complex X, the quotient X^n/X^{n-1} is a wedge sum of n-spheres $\bigvee_\alpha S^n_\alpha$, with one sphere for each n-cell of X.

Smash Product. Like suspension, this is another construction whose importance becomes evident only later. Inside a product space $X \times Y$ there are copies of X and Y, namely $X \times \{y_0\}$ and $\{x_0\} \times Y$ for selected points $x_0 \in X$ and $y_0 \in Y$. These two copies of X and Y in $X \times Y$ intersect only at the point (x_0, y_0), so their union can be identified with the wedge sum $X \vee Y$. The **smash product** $X \wedge Y$ is then defined to be the quotient $X \times Y/X \vee Y$. One can think of $X \wedge Y$ as a reduced version of $X \times Y$ obtained by collapsing away the parts which are not genuinely a product, the separate factors X and Y.

The smash product $X \wedge Y$ is a cell complex if X and Y are cell complexes with x_0 and y_0 0-cells, assuming that we give $X \times Y$ the cell-complex topology rather than the product topology, when these two topologies differ. For example, $S^m \wedge S^n$ has a cell structure with just two cells, of dimensions 0 and $m + n$, hence $S^m \wedge S^n = S^{m+n}$. In particular, when $m = n = 1$ we see that collapsing longitude and meridian circles of a torus to a point produces a 2-sphere.

Two Criteria for Homotopy Equivalence.

Earlier in this chapter the main tool we used for constructing homotopy equivalences was the fact that a mapping cylinder deformation retracts onto its 'target' end. By repeated application of this fact one can often produce homotopy equivalences between rather different-looking spaces. However, this process can be a bit cumbersome in practice, so it is useful to have other techniques available as well. Here is one which can be quite helpful:

(1) *If (X, A) is a CW pair consisting of a CW complex X and a contractible subcomplex A, then the quotient map $X \to X/A$ is a homotopy equivalence.*

A proof of this will be given in the next subsection, but let us look at some examples now.

Example 1: Graphs. Consider the three graphs in the second figure of this chapter. These three graphs are homotopy equivalent since each is a deformation retract of a disk with two holes, but we can also deduce this from statement (1) above since collapsing the middle edge of the first and third graphs produces the second graph.

More generally, suppose X is any graph with finitely many vertices and edges. If the two endpoints of any edge of X are distinct, we can collapse this edge to a point,

producing a homotopy equivalent graph with one fewer edge. This simplification can be iterated until all edges of X are loops, hence each component of X is either an isolated vertex or a wedge sum of circles.

This raises the question of whether two such graphs, having only one vertex in each component, can be homotopy equivalent if they are not in fact just isomorphic graphs. An exercise at the end of the previous section reduces the question to the case of connected graphs. Then the task is to prove that a wedge sum of m circles is not homotopy equivalent to a wedge sum of n circles if $m \neq n$. This sort of thing is hard to do directly. What one would like is some sort of algebraic object associated to spaces, depending only on their homotopy type, and taking different values for $\bigvee_m S^1$ and $\bigvee_n S^1$ if $m \neq n$. In fact the Euler characteristic does this since $\bigvee_m S^1$ and $\bigvee_n S^1$ have Euler characteristics $1-m$ and $1-n$. But it is a rather nontrivial theorem that the Euler characteristic of a space depends only on its homotopy type. A different algebraic invariant that works equally well for graphs, and whose rigorous development requires less effort than the Euler characteristic, is the fundamental group of a space, the subject of Chapter 1.

Example 2. Consider the space X obtained from S^2 by attaching the two ends of an arc A to two distinct points on the sphere, say the north and south poles. Let B be an arc in S^2 joining the two points where A attaches. Then X can be given a CW complex structure with the two endpoints of A and B as 0-cells, the interiors of A and B as 1-cells, and the rest of S^2 as a 2-cell. Since A and B are contractible, X/A and X/B are homotopy equivalent to X. The space X/A is the quotient S^2/S^0, the sphere with two points identified, and X/B is $S^1 \vee S^2$. Hence S^2/S^0 and $S^1 \vee S^2$ are homotopy equivalent, which might not have been entirely obvious *a priori*.

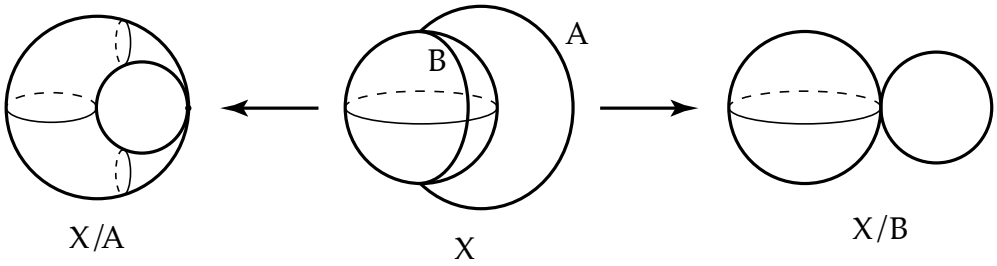

Example 3. Let X be the union of a torus with n meridional disks. To obtain a CW structure on X, choose a longitudinal circle C in X, intersecting each of the meridional disks in one point. These intersection points are then the 0-cells, the 1-cells are the rest of C and the boundaries of the meridional disks, and the 2-cells are the remaining regions of the torus and the interiors of the meridional disks.

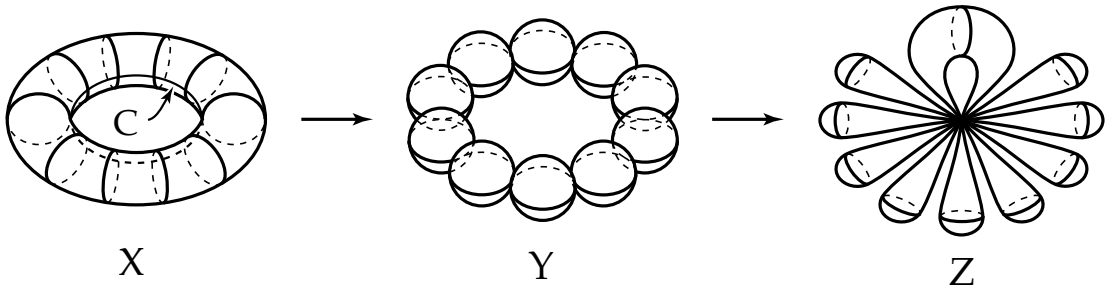

Collapsing each meridional disk to a point yields a homotopy equivalent space Y, a 'necklace with n beads.' We can further collapse all but one of the edges in C to get a homotopy equivalent space Z which is the wedge sum of $n-1$ spheres with the space S^2/S^0 in the preceding example. Using the argument which showed that $S^2/S^0 \simeq S^1 \vee S^2$, we conclude that Z, hence also X, is homotopy equivalent to the wedge sum of S^1 with n spheres.

Another common way to change a space without changing its homotopy type involves the idea of continuously varying how its parts are attached together. A general definition of 'attaching one space to another' which includes the case of attaching cells is the following. Suppose we are given two space X and Y, a subspace $A \subset Y$, and a map $f: A \to X$. Then we can form a quotient space of $X \amalg Y$ by identifying each point $a \in A$ with its image $f(a) \in X$. We denote this quotient space by $X \cup_f Y$, thinking of it as X with Y attached along A via f. When $(Y, A) = (D^n, S^{n-1})$ we have the case of attaching an n-cell to X via a map $f: S^{n-1} \to X$.

Mapping cylinders are examples of this construction, since the mapping cylinder M_f of a map $f: X \to Y$ is the space obtained from Y by attaching $X \times I$ along $X \times \{1\}$ via f. Closely related to the mapping cylinder M_f is the **mapping cone** $C_f = Y \cup_f CX$ where CX is the cone $(X \times I)/(X \times \{0\})$ and we attach this to Y along $X \times \{1\}$ via the identifications $(x, 1) \sim f(x)$. If we identify X with the subspace $X \times \{0\} \subset M_f$, then C_f is just the quotient M_f/X. For example, the space obtained by attaching an n-cell to a space X by a map $f: S^{n-1} \to X$ is the same as the mapping cone C_f.

Here is our second criterion for homotopy equivalence:

(2) *If (Y, A) is a CW pair and we have two attaching maps $f, g: A \to X$ which are homotopic, then $X \cup_f Y \simeq X \cup_g Y$.*

Again let us defer the proof and look at some examples.

Example 4. Let us rederive the result in Example 2 above that a sphere with two points identified is homotopy equivalent to $S^1 \vee S^2$. The sphere with two points identified can be obtained by attaching S^2 to S^1 by a map which wraps a closed arc A in S^2 around

S^1, as shown in the figure. Since A is contractible, this attaching map is homotopic to a constant map, and attaching S^2 to S^1 via a constant map of A yields $S^1 \vee S^2$. The result then follows from (2) since (S^2, A) is a CW pair, S^2 being obtained from A by attaching a 2-cell.

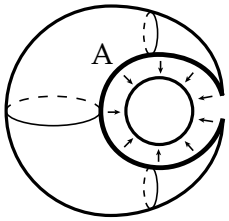

Example 5. In similar fashion we can see that the necklace in Example 3 is homotopy equivalent to the wedge sum of a circle with n 2-spheres. The necklace can be obtained from a circle by attaching n 2-spheres along arcs, so the necklace is homotopy equivalent to the space obtained by attaching n 2-spheres to a circle at points. Then we can slide these attaching points around the circle until they all coincide, producing the wedge sum.

Example 6. Here is an application of the earlier fact that collapsing a contractible subcomplex is a homotopy equivalence: If (X, A) is a CW pair, consisting of a cell complex X and a subcomplex A, then $X/A \simeq X \cup CA$, the mapping cone of the inclusion $A \hookrightarrow X$. For we have $X/A = (X \cup CA)/CA \simeq X \cup CA$ since CA is a contractible subcomplex of $X \cup CA$.

Example 7. If (X, A) is a CW pair and A is contractible in X, i.e., the inclusion $A \hookrightarrow X$ is homotopic to a constant map, then $X/A \simeq X \vee SA$. For, by the previous example we have $X/A \simeq X \cup CA$. And then, since A is contractible in X, the mapping cone $X \cup CA$ of the inclusion $A \hookrightarrow X$ is homotopy equivalent to the mapping cone of a constant map, which is $X \vee SA$. For example, $S^n/S^i \simeq S^n \vee S^{i+1}$ for $i < n$, since S^i is contractible in S^n if $i < n$. In particular this gives $S^2/S^0 \simeq S^2 \vee S^1$, which is Example 2 again.

The Homotopy Extension Property.

In this final subsection of the chapter we shall actually prove a few things, in particular the two criteria for homotopy equivalence described above and the fact that any two homotopy equivalent spaces can be embedded as deformation retracts of the same space.

The proofs depend upon a technical property which arises in many other contexts as well. Consider the following problem. Suppose one is given a map $f_0 : X \to Y$, and on a subspace $A \subset X$ one is also given a homotopy $f_t : A \to Y$ of $f_0 \big| A$ which one would like to extend to a homotopy $f_t : X \to Y$ of the given f_0. If the pair (X, A) is such that this extension problem can always be solved, one says that (X, A) has the **homotopy extension property**. Thus (X, A) has the homotopy extension property if every map $X \times \{0\} \cup A \times I \to Y$ can be extended to a map $X \times I \to Y$.

In particular, the homotopy extension property for (X, A) implies that the identity map $X \times \{0\} \cup A \times I \to X \times \{0\} \cup A \times I$ extends to a map $X \times I \to X \times \{0\} \cup A \times I$, so $X \times$

$\{0\}\cup A\times I$ is a retract of $X\times I$. Conversely if there is a retraction $X\times I\to X\times\{0\}\cup A\times I$, then by composing with this retraction we can extend every map $X\times\{0\}\cup A\times I\to Y$ to a map $X\times I\to Y$. Thus the homotopy extension property for (X,A) is equivalent to $X\times\{0\}\cup A\times I$ being a retract of $X\times I$.

This retraction criterion allows one to give simple examples of pairs (X,A) which do not have the homotopy extension property, such as $\big(I,(0,1)\big)$, since a retraction $I\times I\to I\times\{0\}\cup(0,1)\times I$ would have compact image. A slightly less trivial example is the pair (I,S) where S is the sequence $\{1/n\}$ together with its limit point 0. It is not hard to show that there is no retraction $I\times I\to I\times\{0\}\cup S\times I$, so the homotopy extension property fails for (I,S).

A quite useful result in the positive direction is:

Proposition 0.1. *If (X,A) is a CW pair, then $X\times\{0\}\cup A\times I$ is a deformation retract of $X\times I$, hence (X,A) has the homotopy extension property.*

Proof: There is a retraction $r:D^n\times I\to D^n\times\{0\}\cup\partial D^n\times I$, for example the radial projection from the point $(0,2)\in D^n\times\mathbb{R}$. Then $r_t=tr+(1-t)\mathbb{1}$ is a deformation retraction of $D^n\times I$ onto $D^n\times\{0\}\cup\partial D^n\times I$. This deformation retraction gives rise to a deformation retraction of $X^n\times I$ onto $X^n\times\{0\}\cup(X^{n-1}\cup A^n)\times I$ since $X^n\times I$ is obtained from $X^n\times\{0\}\cup(X^{n-1}\cup A^n)\times I$ by attaching copies of $D^n\times I$ along $D^n\times\{0\}\cup\partial D^n\times I$. If we perform the deformation retraction of $X^n\times I$ onto $X^n\times\{0\}\cup(X^{n-1}\cup A^n)\times I$ during the t-interval $[1/2^{n+1},1/2^n]$, this infinite concatenation of homotopies is a deformation retraction of $X\times I$ onto $X\times\{0\}\cup A\times I$. □

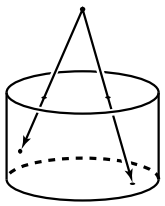

It is interesting to see what the extended homotopy given by this proof actually looks like in the simplest case $X=A\cup e^n$. Points near enough the center of e^n just move radially outward for a while and stop, but points nearer the boundary of e^n move radially outward until they meet the boundary, and then they follow an initial segment of the existing homotopy on A.

Now we can prove the following generalization of statement (1) in the previous subsection:

Proposition 0.2. *If the pair (X,A) satisfies the homotopy extension property and A is contractible, then the quotient map $q:X\to X/A$ is a homotopy equivalence.*

Proof: Let $f_t:X\to X$ be a homotopy extending a contraction of A, with $f_0=\mathbb{1}$. We

shall construct a diagram:

$$\begin{array}{ccc} X & \xrightarrow{f_t} & X \\ {\scriptstyle q}\downarrow & \overset{g}{\nearrow} & \downarrow{\scriptstyle q} \\ X/A & \xrightarrow[\overline{f}_t]{} & X/A \end{array}$$

The lower homotopy $\overline{f}_t$ is induced by f_t and exists since $f_t(A) \subset A$. The map g gives a factorization $f_1 = gq$, which is possible since $f_1(A)$ is a point, the point to which A contracts. The maps g and q are inverse homotopy equivalences since $gq \simeq \mathbb{1}$ via f_t, and $qg \simeq \mathbb{1}$ via $\overline{f}_t$. □

Another application of the homotopy extension property, giving a slightly more refined version of the statement (2) in the previous subsection, is the following:

Proposition 0.3. *If (Y, A) is a CW pair and we have two attaching maps $f, g : A \to X$ which are homotopic, then $X \cup_f Y \simeq X \cup_g Y$ rel X.*

Here the definition of $W \simeq Z$ rel X for pairs (W, X) and (Z, X) is that there are maps $\varphi : W \to Z$ and $\psi : Z \to W$ restricting to the identity on X, such that $\psi\varphi \simeq \mathbb{1}$ and $\varphi\psi \simeq \mathbb{1}$ via homotopies which restrict to the identity on X at all times.

Proof: If $F : A \times I \to X$ is a homotopy from f to g, consider the space $X \cup_F (Y \times I)$. This contains both $X \cup_f Y$ and $X \cup_g Y$ as subspaces. A deformation retraction of $Y \times I$ onto $Y \times \{0\} \cup A \times I$ as in Proposition 0.1 induces a deformation retraction of $X \cup_F (Y \times I)$ onto $X \cup_f Y$, and similarly $X \cup_F (Y \times I)$ deformation retracts onto $X \cup_g Y$. Both these deformation retractions restrict to the identity on X, so together they give a homotopy equivalence $X \cup_f Y \simeq X \cup_g Y$ rel X. □

Here is a technical result whose proof involves a more elaborate application of the homotopy extension property:

Proposition 0.4. *Suppose (X, A) and (Y, A) satisfy the homotopy extension property, and $f : X \to Y$ is a homotopy equivalence with $f \,|\, A = \mathbb{1}$. Then f is a homotopy equivalence* rel A.

Proof: Let $g : Y \to X$ be a homotopy inverse for f. The first step of the proof will be to improve g so that $gf \simeq \mathbb{1}$ rel A. Let h_t be a homotopy from $gf = h_0$ to $\mathbb{1} = h_1$. Since $f \,|\, A = \mathbb{1}$, $h_t \,|\, A$ provides a homotopy from $g \,|\, A$ to $\mathbb{1}$. Extend this to a homotopy $g_t : Y \to Y$ from $g = g_0$ to a map g_1 with $g_1 \,|\, A = \mathbb{1}$. Let $k_t : X \to X$ be the homotopy

h_{1-t} followed by the homotopy $g_t f$, that is, $k_t = h_{1-2t}$ for $t \in [0, \frac{1}{2}]$ and $k_t = g_{2t-1} f$ for $t \in [\frac{1}{2}, 1]$. The homotopy $k_t \,|\, A$ starts and ends with the identity, and its second half simply retraces its first half. Therefore there is a 'homotopy of homotopies' $k_{st} : A \to A$ from $k_t = k_{0t}$ to the constant homotopy $\mathbb{1} = k_{1t}$ with $k_{s0} = \mathbb{1} = k_{s1}$ for all s. Using the homotopy extension property for $(X \times I, A \times I)$, extend k_{st} to be defined on all of X, with $k_{0t} = k_t$. Now if we restrict k_{st} to the three sides of the (s,t)-square where $t = 0$, $s = 1$, and $t = 1$, respectively, we get a homotopy $g_1 f \simeq \mathbb{1}$ rel A.

Since $g_1 \simeq g$, we have $f g_1 \simeq f g \simeq \mathbb{1}$, so the preceding argument can be repeated with the pair f, g replaced by g_1, f. The result is a map $f_1 : X \to X$ with $f_1 \,|\, A = \mathbb{1}$ and $f_1 g_1 \simeq \mathbb{1}$ rel A. Hence $f_1 \simeq f_1(g_1 f) = (f_1 g_1) f \simeq f$ rel A and so $f g_1 \simeq f_1 g_1 \simeq \mathbb{1}$ rel A. Thus g_1 is a homotopy inverse to f rel A. □

Corollary 0.5. *If (X, A) satisfies the homotopy extension property and the inclusion $A \hookrightarrow X$ is a homotopy equivalence, then A is a deformation retract of X.*

Proof: Apply the proposition to the inclusion $A \hookrightarrow X$. □

Corollary 0.6. *A map $f : X \to Y$ is a homotopy equivalence iff X is a deformation retract of the mapping cylinder M_f. Hence, two spaces X and Y are homotopy equivalent iff there is a third space containing both X and Y as deformation retracts.*

Proof: The pair (M_f, X) satisfies the homotopy extension property, so the preceding corollary can be applied to the inclusion $X \hookrightarrow M_f$. □

Exercises

1. Given positive integers v, e, and f satisfying $v - e + f = 2$, construct a cell structure on S^2 having v 0-cells, e 1-cells, and f 2-cells.

2. Enumerate all the subcomplexes of S^∞, with the cell structure described in this section, having two cells in each dimension.

3. Show that S^∞ is contractible.

4. Construct a 2-dimensional cell complex which contains both an annulus $S^1 \times I$ and a Möbius band as deformation retracts.

5. Show that $S^1 * S^1 = S^3$, and more generally $S^m * S^n = S^{m+n+1}$.

6. Show that the space obtained from S^2 by attaching n 2-cells along any collection of n circles in S^2 is homotopy equivalent to the wedge sum of $n + 1$ 2-spheres.

7. Let X be a finite graph lying in a half-plane $P \subset \mathbb{R}^3$ and intersecting the edge of P in a subset of the vertices of X. Describe the homotopy type of the 'surface of revolution' obtained by rotating X about the edge line of P.

8. Show that a CW complex is contractible if it is the union of two contractible subcomplexes whose intersection is also contractible.

9. Let X and Y be CW complexes with 0-cells x_0 and y_0. Show that $X * Y/(X * \{y_0\} \cup \{x_0\} * Y) = S(X \wedge Y)/S(\{x_0\} \wedge \{y_0\})$ and deduce that $X * Y \simeq S(X \wedge Y)$.

10. Show that for a CW complex X with components X_α, the suspension SX is homotopy equivalent to $Y \bigvee_\alpha SX_\alpha$ where Y is a graph. In the case that X is a finite graph, show that SX is homotopy equivalent to a wedge sum of circles and 2-spheres.

11. Use Corollary 0.5 to show that if (X, A) has the homotopy extension property, then $X \times I$ deformation retracts to $X \times \{0\} \cup A \times I$. Deduce from this that Proposition 0.3 holds more generally when (X, A) satisfies the homotopy extension property.

12. Given a pair (X, A) and a map $f : A \to B$, define X/f to be the quotient space of X obtained by identifying points in A having the same image in B. Show that the quotient map $X \to X/f$ is a homotopy equivalence if f is a surjective homotopy equivalence and (X, A) has the homotopy extension property. [Hint: Consider $X \cup M_f$ and use the preceding problem.] When B is a point this gives another proof of Proposition 0.2. Another interesting special case is when f is the projection $A \times I \to A$.

13. Show that if (X, A) satisfies the homotopy extension property, then so does $(Y \cup_f X, Y)$ for every map $f : A \to Y$.

Chapter 1
Fundamental Group and Covering Spaces

To a large extent, algebraic topology is concerned with processes for forming algebraic images of topological spaces. Most often these algebraic images are groups, but more elaborate structures such as rings, modules, and algebras also arise. The mechanisms which create these images — the 'lanterns' of algebraic topology, one might say — are known formally as *functors* and have the characteristic feature that they form images not only of spaces but also of maps. Thus, continuous maps between spaces are projected onto homomorphisms between their algebraic images, so topologically related spaces have algebraically related images. With suitably constructed lanterns one might hope to form a reasonably accurate picture of the shapes of all spaces, or at least of interesting classes of spaces. This is one of the main goals of algebraic topology, and to a surprising extent this goal is achieved. Of course, the lanterns necessary to do this are somewhat complicated pieces of machinery, but this machinery also has a certain intrinsic beauty.

This first chapter introduces one of the simplest and most important of the functors of algebraic topology, the fundamental group, which creates an algebraic picture of the loops, that is, paths starting and ending at the same point, in a given topological space.

The Idea of the Fundamental Group

To get a feeling for what the fundamental group is about, let us look at a few preliminary examples before giving the formal definitions.

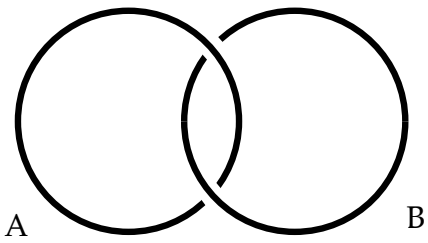

Consider two linked circles A and B in $\mathbb{R}^3$, as shown in the figure. The fact that the circles are linked means, intuitively, that B cannot be separated from A by any continuous motion of B, such as pushing, pulling, twisting, etc. We could even take B to be made of rubber or stretchable string and allow completely general continuous deformations of B, staying in the complement of A at all times, and it would still be impossible to pull B off A. At least that is what intuition suggests, and the fundamental group will give a way of making this intuition mathematically rigorous.

Instead of having B link with A just once, we could make it link A two or more times, as in the next figure.

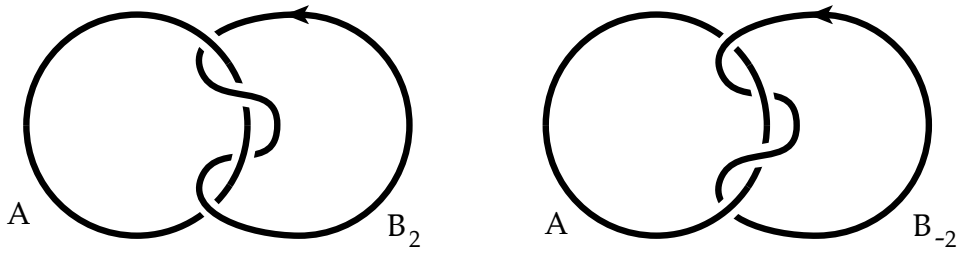

As a further variation, by assigning an orientation to B we can speak of B linking A a positive or a negative number of times, say positive when B comes forward through A and negative for the reverse direction. Thus for each nonzero integer n we have an oriented circle B_n linking A n times, where by 'circle' we mean a curve homeomorphic to a circle. To complete the scheme, we could let B_0 be a circle which doesn't link A at all.

Now, integers not only measure quantity, but they form a group under addition. Can the group operation be mimicked geometrically with some sort of addition operation on the oriented circles B linking A? An oriented circle B can be thought of as a path traversed in time, starting and ending at the same point x_0, which can be any point on the circle. Such a path starting and ending at the same point is called a *loop*. Two different loops B and B' both starting and ending at the same point x_0 can be 'added' to form a new loop $B + B'$ which travels first around B, then around B'. For example, if B_1 and B_1' are loops each linking A once in the positive direction, as in the next figure, then their sum $B_1 + B_1'$ is deformable to B_2, linking A twice.

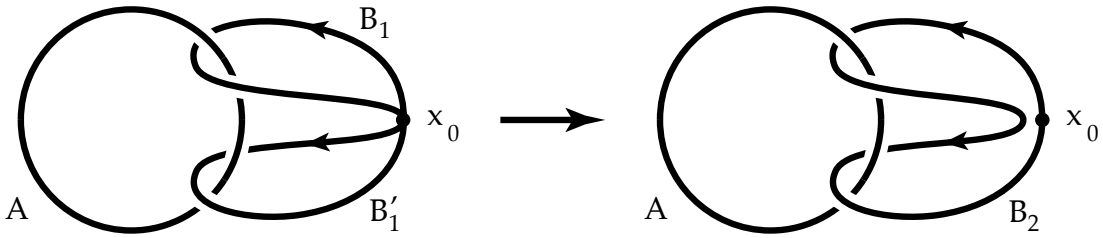

Similarly, $B_1 + B_{-1}$ can be deformed to the loop B_0, unlinked from A:

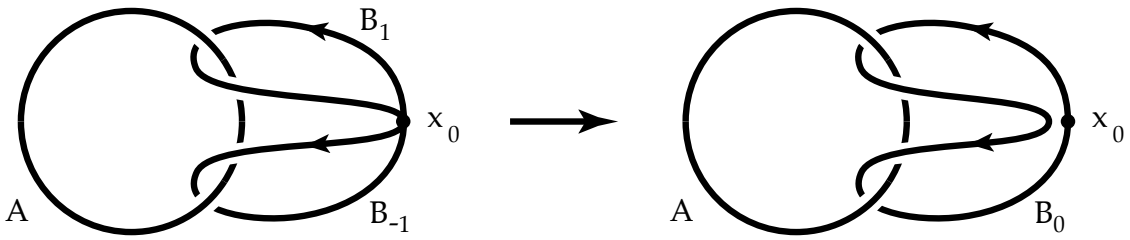

More generally, $B_m + B_n$ is deformable to B_{m+n} for arbitrary integers m and n, as one can easily convince oneself of by drawing a few pictures like those above.

Note that in forming sums of loops we produce loops which pass through the basepoint more than once. This is one reason why loops are defined merely as continuous paths, which are allowed to pass through the same point many times. So if one is thinking of a loop as something made of stretchable string, one has to give the string the magical power of being able to pass through itself unharmed. However, we must be sure not to allow our loops to intersect the fixed circle A at any time, otherwise we could always unlink them from A.

Next we consider a slightly more complicated sort of linking, involving three circles forming a configuration known as the Borromean rings, shown in the first figure below.

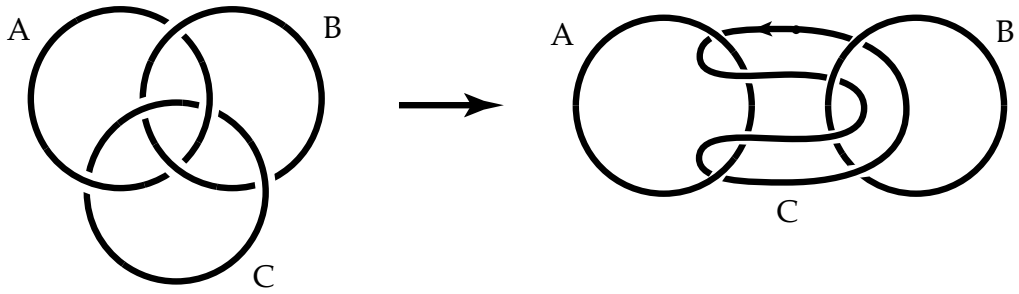

The interesting feature here is that if any one of the three circles is removed, the other two are not linked. It is only the combination of all three that is linked. In the same spirit as before, let us regard the third circle C as forming a loop in the complement of the other two, A and B, and we ask whether C can be continuously deformed to unlink it completely from A and B, always staying in the complement of A and B during the deformation.

We can redraw the picture by pulling A and B apart, dragging C along, and then we see C winding back and forth between A and B as shown in the second figure above. In this new position, if we start at the point of C indicated by the dot and proceed in the direction given by the arrow, then we pass in sequence: (1) forward through A, (2) forward through B, (3) backward through A, and (4) backward through B. If we measure the linking of C with A and B by two integers, then the 'forwards' and 'backwards' cancel and both integers are zero. This reflects the fact that C is not linked with A or B individually. To get a more accurate measure of how C links with A and B together, we regard the four parts (1)-(4) of C as an ordered sequence. Taking into account the directions in which these segments of C pass through A and B, we may deform C to the sum $a + b - a - b$ of four loops as in the next figure. We write the third and fourth

loops as the negatives of the first two since they can be deformed to the first two, but with the opposite orientations, and as we saw in the preceding example, the sum of two oppositely-oriented loops is deformable to a trivial loop, not linked with anything.

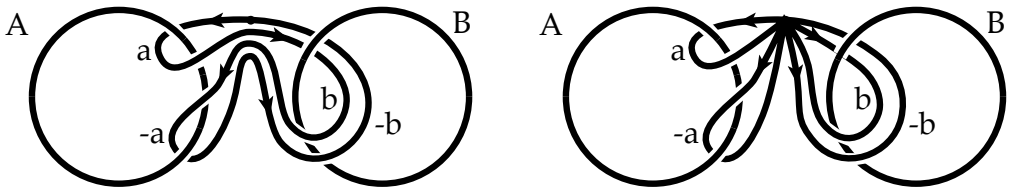

We would like to view the expression $a + b - a - b$ as lying in a nonabelian group, so that it is not automatically zero. Changing to the more usual multiplicative notation for nonabelian groups, it would be written $aba^{-1}b^{-1}$, the commutator of a and b.

To shed further light on this example, suppose we modify it slightly so that the circles A and B are now linked, as in the next figure.

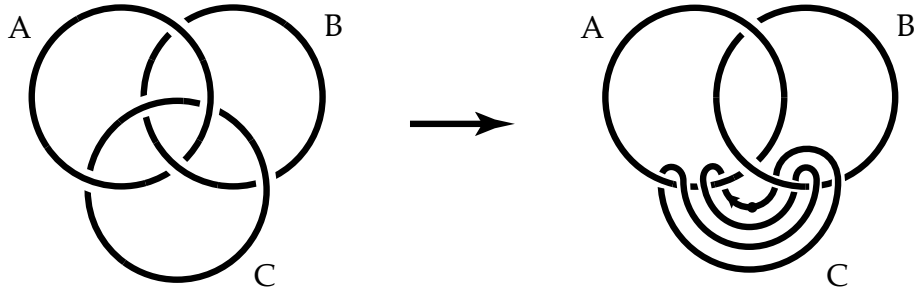

The circle C can then be deformed into the position shown at the right, where it again represents the composite loop $aba^{-1}b^{-1}$, where a and b are loops linking A and B. But it is apparent from the picture on the left that C can actually be unlinked completely from A and B. So in this case the product $aba^{-1}b^{-1}$ should be trivial, i.e., a and b should commute.

The fundamental group of a space X will be defined so that its elements are loops in X starting and ending at a fixed basepoint $x_0 \in X$, but two such loops are regarded as determining the same element of the fundamental group if one loop can be continuously deformed to the other within the space X. (All loops which occur during deformations must also start and end at x_0.) In the first example above, X is the complement of the circle A, while in the other two examples X is the complement of the two circles A and B. In the second section in this chapter we will show:

— The fundamental group of the complement of the circle A in the first example is $\mathbb{Z}$, with the loop B as a generator of this group. This amounts to saying that every loop

in the complement of A can be deformed to one of the loops B_n, and that B_n cannot be deformed to B_m if $n \neq m$.

— The fundamental group of the complement of the two unlinked circles A and B in the second example is the nonabelian free group on two generators, represented by the loops a and b linking A and B. In particular, the commutator $aba^{-1}b^{-1}$ is a nontrivial element of this group.

— The fundamental group of the complement of the two linked circles A and B in the third example is the free abelian group $\mathbb{Z} \times \mathbb{Z}$ on two generators, represented by the loops a and b linking A and B.

As a result of these calculations, we have two ways to tell when a pair of circles A and B is linked. The direct approach is given by the first example, where one circle is regarded as an element of the fundamental group of the complement of the other circle. An alternative and somewhat more subtle method is given by the second and third examples, where one distinguishes a pair of linked circles from a pair of unlinked circles by the fundamental group of their complement, which is abelian in one case and nonabelian in the other. This method is much more general — one can often show that two spaces are not homeomorphic by showing that their fundamental groups are not isomorphic, since it will be a trivial consequence of the definition of the fundamental group that homeomorphic spaces have isomorphic fundamental groups.

1. The Fundamental Group

Paths and Homotopy

There are two basic ingredients in the definition of the fundamental group: the idea of continuous deformation of paths, and an operation for composing paths to form new paths. We begin with the first of these.

Let X be a space and let I denote the unit interval $[0,1]$ in $\mathbb{R}$. By a **path** in X we mean a continuous map $f : I \to X$. The idea of continuously deforming a path, keeping its endpoints fixed, is made precise by the following definition. A **homotopy** of paths in X is a family $f_t : I \to X$, $0 \le t \le 1$, such that

(1) The endpoints $f_t(0) = x_0$ and $f_t(1) = x_1$ are independent of t.
(2) The associated map $F : I \times I \to X$ defined by $F(s,t) = f_t(s)$ is continuous.

When two paths f_0 and f_1 are connected in this way by a homotopy f_t, we will say they are **homotopic** and write $f_0 \simeq f_1$.

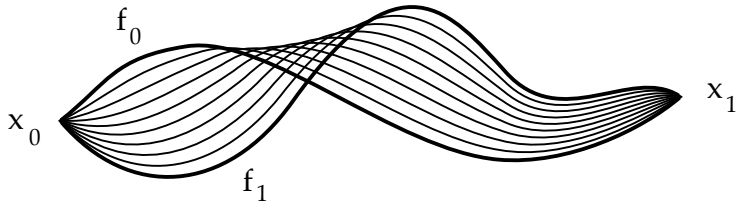

Example: Linear homotopies. If f_0 and f_1 are any two paths from x_0 to x_1 in $\mathbb{R}^n$, then $f_0 \simeq f_1$ via the homotopy $f_t(s) = (1-t)f_0(s) + tf_1(s)$. During this homotopy each point $f_0(s)$ travels along the line segment to $f_1(s)$ at constant speed. This is because the line through $f_0(s)$ and $f_1(s)$ is linearly parametrized as $f_0(s) + t[f_1(s) - f_0(s)] = (1-t)f_0(s) + tf_1(s)$, so as t goes from 0 to 1, $f_t(s)$ traces out the segment from $f_0(s)$ to $f_1(s)$. If $f_1(s) = f_0(s)$ then this segment degenerates to a point, so $f_t(s) = f_0(s)$ for all t. This happens in particular for $s = 0$ and $s = 1$, so f_t is a path from x_0 to x_1 for all t. Continuity of the homotopy f_t as a map $I \times I \to \mathbb{R}^n$ follows from continuity of f_0 and f_1 since the algebraic operations in the definition of f_t are continuous.

This construction shows more generally that all paths in a convex subspace $X \subset \mathbb{R}^n$, with given endpoints x_0 and x_1, are homotopic.

Before proceeding further we need to verify a technical property:

Proposition 1.1. *The relation of homotopy on paths with fixed endpoints in any space X is an equivalence relation.*

Proof: The constant homotopy $f_t = f$ shows that $f \simeq f$. If $f_0 \simeq f_1$ via f_t, then $f_1 \simeq f_0$ via f_{1-t}. For transitivity, if $f_0 \simeq f_1$ via f_t and $f_1 = g_0 \simeq g_1$ via g_t, then $f_0 \simeq g_1$ via the homotopy h_t which equals f_{2t} for $0 \le t \le \frac{1}{2}$ and g_{2t-1} for $\frac{1}{2} \le t \le 1$. Continuity of the associated map $H(s,t) = h_t(s)$ comes from the elementary fact, which will be used frequently without explicit mention, that a function defined on the union of two closed sets is continuous if its restrictions to both of the closed sets are continuous. In the case at hand, H is clearly continuous on each of $I \times [0, \frac{1}{2}]$ and $I \times [\frac{1}{2}, 1]$, and on $I \times \{\frac{1}{2}\}$ the two definitions of H agree since $f_1 = g_0$, so H is continuous on $I \times I$. □

The equivalence class of a path f under the equivalence relation of homotopy is called its **homotopy class**, written $[f]$.

Given two paths $f, g : I \to X$ such that $f(1) = g(0)$ we can define a 'composition' or 'product' path $f \cdot g$ which traverses first f then g, by the formula

$$f \cdot g(s) = \begin{cases} f(2s), & 0 \le s \le \frac{1}{2} \\ g(2s-1), & \frac{1}{2} \le s \le 1 \end{cases}$$

Thus the speed of traversal of f and g is doubled in order for $f \cdot g$ to be traversed in unit time. This product operation respects homotopy classes since if $f_0 \simeq f_1$ and $g_0 \simeq g_1$ via homotopies f_t and g_t, and if $f_0(1) = g_0(0)$ so that $f_0 \cdot g_0$ is defined, then $f_t \cdot g_t$ is defined and provides a homotopy $f_0 \cdot g_0 \simeq f_1 \cdot g_1$.

In particular, suppose we restrict attention to paths $f : I \to X$ with the same starting and ending point $f(0) = f(1) = x_0 \in X$. Such paths are called **loops**, and the common starting and ending point x_0 is referred to as the **basepoint**. The set $\pi_1(X, x_0)$ of homotopy classes of loops in X at the basepoint x_0 is called the **fundamental group** of X at the basepoint x_0. This terminology is justified by:

Proposition 1.2. *$\pi_1(X, x_0)$ is a group with respect to the product $[f] \cdot [g] = [f \cdot g]$.*

Proof: By restricting attention to loops with a fixed basepoint $x_0 \in X$ we guarantee that the product $f \cdot g$ of any two such loops is defined. We have already observed that the homotopy class of $f \cdot g$ depends only on the homotopy classes of f and g, so the product $[f] \cdot [g] = [f \cdot g]$ is well-defined. It remains then to verify the three axioms for a group.

As a preliminary step, define a **reparametrization** of a path f to be a composition $f\varphi$ where $\varphi : I \to I$ is any continuous map such that $\varphi(0) = 0$ and $\varphi(1) = 1$. Reparametrizing a path preserves its homotopy class since $f\varphi \simeq f$ via the homotopy $f\varphi_t$ where $\varphi_t(s) = (1-t)\varphi(s) + ts$ so that $\varphi_0 = \varphi$ and $\varphi_1(s) = s$. Note that $(1-t)\varphi(s) + ts$ lies between $\varphi(s)$ and s, hence in I, so $f\varphi_t$ is defined.

Given paths f, g, h with $f(1) = g(0)$ and $g(1) = h(0)$, then the composed paths $(f \cdot g) \cdot h$ and $f \cdot (g \cdot h)$ are reparametrizations of each other, differing only in the speeds at which f and h are traversed. Hence $(f \cdot g) \cdot h \simeq f \cdot (g \cdot h)$. Restricting attention to loops at the basepoint x_0, this says the product in $\pi_1(X, x_0)$ is associative.

Given a path $f : I \to X$, let c be the constant path at $f(1)$, defined by $c(s) = f(1)$ for all $s \in I$. Then $f \cdot c$ is a reparametrization of f via the function $\varphi(s)$ which equals $2s$ on $[0, \frac{1}{2}]$ and 1 on $[\frac{1}{2}, 1]$, so $f \cdot c \simeq f$. Similarly, $c \cdot f \simeq f$ where c is now the constant path at $f(0)$. Taking f to be a loop, we deduce that the homotopy class of the constant path at x_0 is a two-sided identity in $\pi_1(X, x_0)$.

For a path f from x_0 to x_1, the **inverse path** $\overline{f}$ from x_1 back to x_0 is defined by $\overline{f}(s) = f(1 - s)$. Consider the homotopy $h_t = f_t \cdot \overline{(f_t)}$ where f_t is the path which equals f on the interval $[0, t]$ and which is stationary at $f(t)$ on the interval $[t, 1]$. Thus f_0 is the constant path c at x_0 and $f_1 = f$, so h_t is a homotopy from $c \cdot c = c$ to $f \cdot \overline{f}$. Now specialize to the case that f is a loop at the basepoint x_0. We have shown that $f \cdot \overline{f} \simeq c$, and replacing f by $\overline{f}$ gives $\overline{f} \cdot f \simeq c$. So $[\overline{f}]$ is an inverse for $[f]$ in $\pi_1(X, x_0)$. □

In Chapter 4 we will see that $\pi_1(X, x_0)$ is just the first in a sequence of groups $\pi_n(X, x_0)$, called homotopy groups, which are defined in an entirely analogous fashion using the n-dimensional cube I^n in place of I.

Example. A convex set X in $\mathbb{R}^n$ has $\pi_1(X, x_0) = 0$, the trivial group, for every basepoint $x_0 \in X$, since any two loops f_0 and f_1 based at x_0 are homotopic via the linear homotopy $f_t(s) = (1 - t)f_0(s) + tf_1(s)$.

It is not so easy to show that a space has a nontrivial fundamental group since one must somehow demonstrate the nonexistence of homotopies between certain loops. A likely candidate for such a space is the punctured plane $\mathbb{R}^2 - \{0\}$ since intuition suggests that a loop like $f(s) = (\cos 2\pi s, \sin 2\pi s)$ which encircles the origin cannot be homotopic to a constant loop by any homotopy which avoids the origin. Later in this section we will prove this.

One may ask how $\pi_1(X, x_0)$ depends on the choice of the basepoint x_0. Since $\pi_1(X, x_0)$ involves only the path component of X containing x_0, it is clear that we can hope to find a relation between $\pi_1(X, x_0)$ and $\pi_1(X, x_1)$ for two basepoints $x_0, x_1 \in X$ only if x_0 and x_1 lie in the same path component of X. So let $h : I \to X$ be a path from x_0 to x_1, with the inverse path $\overline{h}(s) = h(1 - s)$ from x_1 back to x_0. We can then associate to each loop f based at x_1 the loop $h \cdot f \cdot \overline{h}$ based at x_0. Strictly speaking, we should choose an order of forming the product $h \cdot f \cdot \overline{h}$, either $(h \cdot f) \cdot \overline{h}$ or $h \cdot (f \cdot \overline{h})$, but the two choices are homotopic and we are only interested in homotopy classes here.

Proposition 1.3. *The map* $\beta_h : \pi_1(X, x_1) \to \pi_1(X, x_0)$ *defined by* $\beta_h[f] = [h \cdot f \cdot \overline{h}]$ *is an isomorphism.*

Proof: If f_t is a homotopy of loops based at x_1 then $h \cdot f_t \cdot \overline{h}$ is a homotopy of loops based at x_0, so β_h is well-defined. Further, β_h is a homomorphism since $\beta_h[f \cdot g] = [h \cdot f \cdot g \cdot \overline{h}] = [h \cdot f \cdot \overline{h} \cdot h \cdot g \cdot \overline{h}] = \beta_h[f]\beta_h[g]$. Finally, β_h is an isomorphism with inverse $\beta_{\overline{h}}$ since $\beta_h \beta_{\overline{h}}[f] = \beta_h[\overline{h} \cdot f \cdot h] = [h \cdot \overline{h} \cdot f \cdot h \cdot \overline{h}] = [f]$, and similarly $\beta_{\overline{h}}\beta_h[f] = [f]$. □

Thus if X is path-connected, then $\pi_1(X, x_0)$ is, up to isomorphism, independent of the choice of basepoint x_0. In this case the notation $\pi_1(X, x_0)$ is often abbreviated to $\pi_1(X)$.

Note that if h happens to be a loop at x_0, then β_h is conjugation by $[h]$, an inner automorphism of $\pi_1(X, x_0)$.

A path-connected space X with $\pi_1(X) = 0$ is said to be **simply-connected**. This terminology probably stems from the following fact:

Proposition 1.4. *For a space* X *the following are equivalent:*

(a) X *is path-connected and* $\pi_1(X) = 0$.

(b) For all pairs of points $x_0, x_1 \in X$ *there is a unique homotopy class of paths in* X *from* x_0 *to* x_1.

Proof: If we choose $x_1 = x_0$ in (b) then we obtain $\pi_1(X, x_0) = 0$. Conversely, let f_0 and f_1 be two paths joining given points x_0 and x_1. Then $f_0 \cdot \overline{f}_1$ is a loop at x_0, so if $\pi_1(X, x_0) = 0$ there is a homotopy from $f_0 \cdot \overline{f}_1$ to the constant loop. Hence $f_0 \cdot \overline{f}_1 \cdot f_1 \simeq f_1$. On the other hand, $\overline{f}_1 \cdot f_1$ is homotopic to the constant loop at x_1, so $f_0 \cdot \overline{f}_1 \cdot f_1 \simeq f_0$ and thus $f_0 \simeq f_1$. □

The next proposition can be used to show that a number of common spaces are simply-connected.

Proposition 1.5. *A space* X *is simply-connected if it can be expressed as the union of two simply-connected open subsets* A *and* B *such that* $A \cap B$ *is non-empty and path-connected.*

Proof: If A, B, and $A \cap B$ are path-connected, so is X. To show $\pi_1(X) = 0$, consider a loop $f : I \to X$ based at a point x_0 which we choose to lie in $A \cap B$. We can choose a partition $0 = s_0 < s_1 < \cdots < s_n = 1$ of I such that f maps each interval $[s_{k-1}, s_k]$ either to A or to B, by the following procedure. Each $s \in I$ is an interior point of an interval $[a_s, b_s] \subset \mathbb{R}$ for which $f([a_s, b_s] \cap I)$ is contained in either A or B. Finitely many of the

open intervals (a_s, b_s) cover I by compactness. The endpoints in I of these finitely many intervals then give the desired partition.

After discarding unnecessary s_k's we may assume that no two adjacent intervals $[s_{k-1}, s_k]$ and $[s_k, s_{k+1}]$ map both to A or both to B. Hence $f(s_k) \in A \cap B$ for each k. Since $A \cap B$ is path-connected we can choose a path f_k in $A \cap B$ from $f(s_{k-1})$ to $f(s_k)$. Since A and B are simply-connected, the restriction of f to $[s_{k-1}, s_k]$ is homotopic to f_k, so f is homotopic to the loop $f_1 \cdot \cdots \cdot f_n$. This loop lies in $A \cap B$, and in particular in A which is simply-connected, so $f_1 \cdot \cdots \cdot f_n$ is homotopic to the constant loop. □

Corollary 1.6. *The unit sphere S^n in $\mathbb{R}^{n+1}$ is simply-connected if $n \geq 2$.*

Proof: Take $A = S^n - \{(0, \cdots, 0, 1)\}$ and $B = S^n - \{(0, \cdots, 0, -1)\}$. Then A and B are each homeomorphic to $\mathbb{R}^n$, hence simply-connected. (It should be apparent from the definition that any space homeomorphic to a simply-connected space is itself simply-connected.) Furthermore, $A \cap B$ is homeomorphic to $S^{n-1} \times \mathbb{R}$ hence is path-connected if $n \geq 2$, and the result follows. □

Our list of simply-connected spaces can be expanded using:

Proposition 1.7. *$\pi_1(X \times Y)$ is isomorphic to $\pi_1(X) \times \pi_1(Y)$ if X and Y are path-connected.*

Proof: A basic property of the product topology is that a map $f : Z \to X \times Y$ is equivalent to a pair of maps $g : Z \to X$ and $h : Z \to Y$, where $f(z) = (g(z), h(z))$. In particular, a loop in $X \times Y$ based at (x_0, y_0) is equivalent to a pair of loops in X and Y based at x_0 and y_0 respectively. By the same token, a homotopy of a loop in $X \times Y$ is equivalent to a pair of homotopies of the corresponding loops in X and Y. Thus we obtain a bijection $\pi_1(X \times Y, (x_0, y_0)) \approx \pi_1(X, x_0) \times \pi_1(Y, y_0)$, and this is obviously a group isomorphism. □

Example. If we delete a point from $\mathbb{R}^n$, is the resulting space still simply-connected? Taking the deleted point to be the origin, $\mathbb{R}^n - \{0\}$ is homeomorphic to $S^{n-1} \times (0, \infty)$, and $(0, \infty)$ is simply-connected, being convex, so we obtain an isomorphism $\pi_1(\mathbb{R}^n - \{0\}) \approx \pi_1(S^{n-1})$, assuming $n \geq 2$ so that these spaces are path-connected. Thus $\mathbb{R}^n - \{0\}$ is simply-connected if $n \geq 3$.

More generally, the complement of any finite set $S \subset \mathbb{R}^n$ is simply-connected if $n \geq 3$. This can be proved by applying Proposition 1.5 and induction on the number of points in S, as follows. Assuming that S contains at least two points, it contains points differing in some coordinate, say in the coordinate x_j, so there is a number $c \in \mathbb{R}$ such that the plane

P defined by the equation $x_j = c$ is disjoint from S and separates S into two proper subsets. Let $U = \{x \in \mathbb{R}^n \mid x_j < c + \varepsilon\}$ and $V = \{x \in \mathbb{R}^n \mid x_j > c - \varepsilon\}$ with ε small enough so that $U \cap V$ contains no points of S. Since U and V are homeomorphic to $\mathbb{R}^n$, induction on the number of points in S implies that $U - S$ and $V - S$ are simply-connected if $n \geq 3$. Then Proposition 1.5 implies that $\mathbb{R}^n - S$ is simply-connected if $n \geq 3$.

For $n = 2$ we have $\pi_1(\mathbb{R}^2 - \{0\}) \approx \pi_1(S^1)$, and our next task will be to compute this group.

The Fundamental Group of the Circle

Our first real theorem will be the calculation $\pi_1(S^1) \approx \mathbb{Z}$. This is a very basic result. Besides its intrinsic interest, it has several immediate applications. Furthermore, it will be the starting point for many more calculations in the next section, and the ideas in its proof will underlie much of §1.3.

Theorem 1.8. *The map $\psi : \mathbb{Z} \to \pi_1(S^1)$ sending an integer n to the homotopy class of the loop $\omega_n(s) = (\cos 2\pi ns, \sin 2\pi ns)$ based at $(1, 0)$ is an isomorphism.*

Proof: The idea is to compare paths in S^1 with paths in $\mathbb{R}$ via the map $p : \mathbb{R} \to S^1$ given by $p(s) = (\cos 2\pi s, \sin 2\pi s)$. This map can be visualized geometrically by embedding $\mathbb{R}$ in $\mathbb{R}^3$ as the helix $(\cos 2\pi s, \sin 2\pi s, s)$, and then p is the restriction to the helix of the projection of $\mathbb{R}^3$ onto $\mathbb{R}^2$, $(x, y, z) \mapsto (x, y)$, as in the figure. Observe that the loop ω_n is the composition $p\tilde{\omega}_n$ where $\tilde{\omega}_n : I \to \mathbb{R}$ is the path $\tilde{\omega}_n(s) = ns$, starting at 0 and ending at n. The relation $\omega_n = p\tilde{\omega}_n$ is expressed by saying that $\tilde{\omega}_n$ is a **lift** of ω_n.

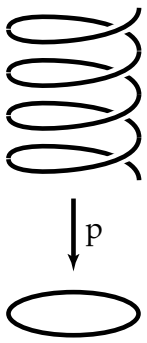

The definition of ψ can be restated by setting $\psi(n)$ equal to the homotopy class of the loop $p\tilde{f}$ where $\tilde{f}$ is any path in $\mathbb{R}$ from 0 to n. Such an $\tilde{f}$ is homotopic to $\tilde{\omega}_n$ since $\mathbb{R}$ is simply-connected, hence $p\tilde{f}$ is homotopic to $p\tilde{\omega}_n = \omega_n$ and the new definition of ψ agrees with the old one.

To verify that ψ is a homomorphism, let $\tau_m : \mathbb{R} \to \mathbb{R}$ be the translation $\tau_m(x) = x + m$. Then $\tilde{\omega}_m \cdot (\tau_m \tilde{\omega}_n)$ is a path in $\mathbb{R}$ from 0 to $m + n$, so $\psi(m + n)$ is the homotopy class of the loop in S^1 which is the image of this path under p. This image is just $\omega_m \cdot \omega_n$, so $\psi(m + n) = \psi(m) \cdot \psi(n)$.

To show that ψ is an isomorphism we shall use two facts:

(a) For each path $f : I \to S^1$ starting at a point $x_0 \in S^1$ and each $\tilde{x}_0 \in p^{-1}(x_0)$ there is a unique lift $\tilde{f} : I \to \mathbb{R}$ starting at $\tilde{x}_0$.

(b) For each homotopy $f_t : I \to S^1$ of paths starting at x_0 and each $\tilde{x}_0 \in p^{-1}(x_0)$ there is a unique lifted homotopy $\tilde{f}_t : I \to \mathbb{R}$ of paths starting at $\tilde{x}_0$.

Before proving these facts, let us see how they imply the theorem. To show that ψ is surjective, let $f : I \to S^1$ be a loop at the basepoint $(1, 0)$ representing a given element of $\pi_1(S^1)$. By (a) there is a lift $\tilde{f}$ starting at 0. This path $\tilde{f}$ ends at some integer n since $p\tilde{f}(1) = f(1) = (1, 0)$ and $p^{-1}(1, 0) = \mathbb{Z} \subset \mathbb{R}$. By the extended definition of ψ we then have $\psi(n) = [p\tilde{f}] = [f]$ and ψ is surjective.

To show that ψ is injective, suppose $\psi(m) = \psi(n)$, i.e., $\omega_m \simeq \omega_n$. Let f_t be a homotopy from $\omega_m = f_0$ to $\omega_n = f_1$. By (b) this homotopy lifts to a homotopy $\tilde{f}_t$ of paths starting at 0. The uniqueness part of (a) implies that $\tilde{f}_0 = \tilde{\omega}_m$ and $\tilde{f}_1 = \tilde{\omega}_n$. Since $\tilde{f}_t$ is a homotopy of paths, the endpoint $\tilde{f}_t(1)$ is independent of t. For $t = 0$ this endpoint is m and for $t = 1$ it is n, so $m = n$.

It remains to prove (a) and (b). These can both be deduced from a third statement:

(c) Given a map $F : Y \times I \to S^1$ and a map $\tilde{F} : Y \times \{0\} \to \mathbb{R}$ lifting $F|Y \times \{0\}$, then there is a unique map $\tilde{F} : Y \times I \to \mathbb{R}$ lifting F and restricting to the given $\tilde{F}$ on $Y \times \{0\}$.

Statement (a) is the special case that Y is a point, and (b) is obtained by applying (c) with $Y = I$ in the following way. The homotopy f_t in (b) gives $F : I \times I \to S^1$ by setting $F(s, t) = f_t(s)$ as usual. A unique lift $\tilde{F} : I \times \{0\} \to \mathbb{R}$ is obtained by an application of (a). Then (c) gives a unique $\tilde{F} : I \times I \to \mathbb{R}$. The restrictions $\tilde{F}|\{0\} \times I$ and $\tilde{F}|\{1\} \times I$ are paths lifting the constant path at x_0, hence they must also be constant by the uniqueness part of (a), so $\tilde{f}_t(s) = \tilde{F}(s, t)$ is a homotopy of paths lifting f_t.

We shall prove (c) using just one special property of the projection $p : \mathbb{R} \to S^1$, namely:

(∗) There is an open cover $\{U_i\}$ of S^1 such that for each i, $p^{-1}(U_i)$ is the disjoint union of open sets $\tilde{U}_{ij}$ each of which is mapped homeomorphically onto U_i by p.

For example, we could take the cover $\{U_i\}$ to consist of any two open arcs in S^1 whose union is S^1.

To prove (c) we will first construct $\tilde{F} : N \times I \to \mathbb{R}$ where N is some neighborhood of a given point $y_0 \in Y$. Since F is continuous, every point (y_0, t) has a product neighborhood $N_t \times V_t$ such that $F(N_t \times V_t) \subset U_i$ for some i. By compactness of $\{y_0\} \times I$, finitely many such products $N_t \times V_t$ cover $\{y_0\} \times I$, so we can choose a single neighborhood N of y_0 and a partition $0 = t_0 < t_1 < \cdots < t_m = 1$ of I so that for each k, $F(N \times [t_k, t_{k+1}])$ is contained in some U_i. Assume inductively that $\tilde{F}$ has been constructed on $N \times [0, t_k]$, and suppose that $F(N \times [t_k, t_{k+1}]) \subset U_i$ and $\tilde{F}(y_0, t_k) \in \tilde{U}_{ij}$. After replacing N by

$N \cap (\widetilde{F}\,|\,N \times \{t_k\})^{-1}(\widetilde{U}_{ij})$, a perhaps smaller neighborhood of y_0, we may assume that $\widetilde{F}(N \times \{t_k\}) \subset \widetilde{U}_{ij}$. Then we can define $\widetilde{F}$ on $N \times [t_k, t_{k+1}]$ to be the composition of the maps F and $p^{-1}: U_i \to \widetilde{U}_{ij}$. After finitely many repetitions of this induction step we eventually get a lift $\widetilde{F}: N \times I \to \mathbb{R}$.

Next we show a special case of the uniqueness part of (c), the case that Y is a point. In this case we can omit Y from the notation. So suppose $\widetilde{F}$ and $\widetilde{F}'$ are two lifts of $F: I \to S^1$ such that $\widetilde{F}(0) = \widetilde{F}'(0)$. As before, choose a partition $0 = t_0 < t_1 < \cdots < t_m = 1$ of I so that for each k, $F([t_k, t_{k+1}])$ is contained in some U_i. Assume inductively that $\widetilde{F}' = \widetilde{F}$ on $[0, t_k]$. Since $[t_k, t_{k+1}]$ is connected, so is $\widetilde{F}'([t_k, t_{k+1}])$, which must therefore lie in a single $\widetilde{U}_{ij}$, in fact in the $\widetilde{U}_{ij}$ containing $\widetilde{F}([t_k, t_{k+1}])$ since $\widetilde{F}'(t_k) = \widetilde{F}(t_k) \in \widetilde{U}_{ij}$. Since p is injective on $\widetilde{U}_{ij}$ it follows that $\widetilde{F}' = \widetilde{F}$ on $[t_k, t_{k+1}]$, finishing the induction.

The last step in the proof of (c) is to observe that since the $\widetilde{F}$'s constructed above on sets of the form $N \times I$ are unique when restricted to each segment $\{y\} \times I$, they must agree whenever two such sets $N \times I$ overlap. So we obtain a well-defined lift $\widetilde{F}$ on all of $Y \times I$. This $\widetilde{F}$ is continuous since it is continuous on each $N \times I$, and it is unique since it is unique on each segment $\{y\} \times I$. □

Although algebraic topology is usually 'algebra serving topology,' the roles are reversed in the following proof of the Fundamental Theorem of Algebra based on the calculation $\pi_1(S^1) \approx \mathbb{Z}$.

Theorem 1.9. *Every nonconstant polynomial with coefficients in* $\mathbb{C}$ *has a root in* $\mathbb{C}$.

Proof: We may assume the polynomial is monic, of the form $p(z) = z^n + a_1 z^{n-1} + \cdots + a_n$. If we assume that $p(z)$ has no roots in $\mathbb{C}$, then for each real number $r \geq 0$, the formula

$$f_r(s) = \frac{p(re^{2\pi is})/p(r)}{|p(re^{2\pi is})/p(r)|}$$

defines a loop in S^1 based at 1. As r varies, f_r is a homotopy of loops based at 1, and f_0 is the trivial loop, so $[f_r] \in \pi_1(S^1)$ is zero for all r. Now fix an $r > \max\{1, |a_1| + \cdots + |a_n|\}$. Then for $|z| = r$ we have $|z^n| > |a_1 z^{n-1} + \cdots + a_n|$, so for $0 \leq t \leq 1$ the polynomial $p_t(z) = z^n + t(a_1 z^{n-1} + \cdots + a_n)$ has no roots on the circle $|z| = r$. Replacing p by p_t in the formula for f_r above and letting t go from 1 to 0, we obtain a homotopy from the loop f_r to the loop $g(s) = e^{2ins}$. By Theorem 1.8, g represents the element $n \in \mathbb{Z} \approx \pi_1(S^1)$. Thus $[g] = [f_r] = 0$ implies $n = 0$, and we have shown that the only polynomials without roots in $\mathbb{C}$ are constants. □

For the next result we use the standard notation D^n for the closed unit disk in $\mathbb{R}^n$, all vectors x of length $|x| \leq 1$. Thus the boundary of D^n is the unit sphere S^{n-1}.

Theorem 1.10. *Every continuous map $h : D^2 \to D^2$ has a fixed point, i.e., a point $x \in D^2$ with $h(x) = x$.*

Proof: If on the contrary $h(x) \neq x$ for all $x \in D^2$, we can define a map $r : D^2 \to S^1$ by letting $r(x)$ be the point of S^1 where the ray in $\mathbb{R}^2$ starting at $h(x)$ and passing through x leaves D^2. Continuity of r is clear since small perturbations of x produce small perturbations of $h(x)$, hence small perturbations of the ray through these two points. The crucial property of r, besides continuity, is that $r(x) = x$ if $x \in S^1$.

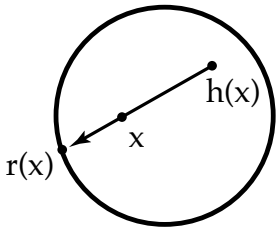

Let f_0 be a loop in S^1. In D^2 there is a homotopy of f_0 to a constant loop, for example the linear homotopy $f_t(s) = (1-t)f_0(s) + tx_0$ where x_0 is the basepoint of f_0. Since r is the identity on S^1, the composition rf_t is then a homotopy in S^1 from $rf_0 = f_0$ to the constant loop at x_0. But this contradicts the fact that $\pi_1(S^1) \neq 0$. □

This theorem was first proved by Brouwer around 1910, one of the early triumphs of algebraic topology. Brouwer in fact proved the corresponding result for D^n, and we shall obtain this generalization in Chapter 2 using homology groups in place of π_1. One could also use the higher homotopy group π_n. Brouwer's original proof, however, used neither homology nor homotopy groups, which had not been invented at the time. Instead it used the notion of degree for maps $S^n \to S^n$, which we shall define in §2.2 using homology but which Brouwer defined directly in more geometric terms.

All these versions of the proof of the fixed point theorem are arguments by contradiction, and so they show just the existence of a fixed point without giving any clue as to how to find it in explicit cases. Our proof of the fundamental theorem of algebra was similar in this regard. It is interesting that a few years after proving the fixed point theorem, Brouwer became very critical of such nonconstructive proofs and shifted his energy away from topology toward furthering the constructive view of mathematics.

Another application of the calculation $\pi_1(S^1) \approx \mathbb{Z}$ is the following:

Example: The Torus. Since $\pi_1(S^1 \times S^1) \approx \pi_1(S^1) \times \pi_1(S^1)$ by Proposition 1.7, we get $\pi_1(S^1 \times S^1) \approx \mathbb{Z} \times \mathbb{Z}$. Under this isomorphism a pair $(p, q) \in \mathbb{Z} \times \mathbb{Z}$ corresponds to a loop which winds p times around one S^1 factor of the torus and q times around the other S^1 factor, for example the loop $w_{pq}(s) = (w_p(s), w_q(s))$.

More generally, the n-dimensional torus, which is the product of n circles, has fundamental group the product of n copies of $\mathbb{Z}$. This follows by induction on n.

Induced Homomorphisms

Suppose $\varphi : X \to Y$ is a map taking the basepoint $x_0 \in X$ to the basepoint $y_0 \in Y$. For brevity we write $\varphi : (X, x_0) \to (Y, y_0)$ in this situation. Then φ induces a homomorphism $\varphi_* : \pi_1(X, x_0) \to \pi_1(Y, y_0)$, defined by composing loops $f : I \to X$ based at x_0 with φ, that is, $\varphi_*[f] = [\varphi f]$. This induced map φ_* is well-defined since a homotopy f_t of loops based at x_0 yields a composed homotopy φf_t of loops based at y_0, so $\varphi_*[f_0] = [\varphi f_0] = [\varphi f_1] = \varphi_*[f_1]$. Furthermore, φ_* is a homomorphism since $\varphi(f \cdot g) = (\varphi f) \cdot (\varphi g)$.

Two basic properties of induced homomorphisms are:

(i) $(\varphi\psi)_* = \varphi_*\psi_*$ for a composition $(X, x_0) \xrightarrow{\psi} (Y, y_0) \xrightarrow{\varphi} (Z, z_0)$.

(ii) $\mathbb{1}_* = \mathbb{1}$, where '$\mathbb{1}$' is a generic symbol for the identity map of a space or a group.

The first of these follows since composition of maps is associative, $(\varphi\psi)f = \varphi(\psi f)$, and the second is obvious. From (i) and (ii) it follows formally that if φ is a homeomorphism with inverse ψ, then φ_* is an isomorphism with inverse ψ_* since $\varphi_*\psi_* = (\varphi\psi)_* = \mathbb{1}_* = \mathbb{1}$ and similarly $\psi_*\varphi_* = \mathbb{1}$.

Properties (i) and (ii) of induced homomorphisms are what makes the fundamental group into a functor. The formal definition of functors requires the introduction of certain other preliminary concepts, however, so we postpone this until it is needed in §2.3.

Induced homomorphisms allow certain relations between spaces to be transformed into relations between their fundamental groups. For example:

Proposition 1.11. *If A is a retract of X, then the homomorphism $i_* : \pi_1(A, x_0) \to \pi_1(X, x_0)$ induced by the inclusion $i : A \hookrightarrow X$ is injective. If A is a deformation retract of X, then i_* is an isomorphism.*

Proof: If $r : X \to A$ is a retraction, then $ri = \mathbb{1}$, hence $r_*i_* = \mathbb{1}$, which implies that i_* is injective. If r is a deformation retraction via a homotopy $r_t : X \to X$ with $r_0 = \mathbb{1}$, $r_t|A = \mathbb{1}$, and $r_1(X) \subset A$, then composing r_t with any loop $f : I \to X$ based at $x_0 \in A$ gives a homotopy of f to a loop in A, so i_* is then also surjective. □

In particular, this gives another way of seeing that S^1 is not a retract of D^2 since the inclusion-induced map $\pi_1(S^1) \to \pi_1(D^2)$ is a homomorphism $\mathbb{Z} \to 0$, which cannot be injective.

Recall from Chapter 0 the general definition of a homotopy: a family $\varphi_t : X \to Y$, $t \in I$, such that the associated map $\Phi : X \times I \to Y, \Phi(x, t) = \varphi_t(x)$, is continuous. If φ_t takes a subspace $A \subset X$ to a subspace $B \subset Y$ for all t, then we speak of a homotopy of

maps of pairs, $\varphi_t : (X, A) \to (Y, B)$. In particular, a **basepoint-preserving homotopy** $\varphi_t : (X, x_0) \to (Y, y_0)$ is the case that $\varphi_t(x_0) = y_0$ for all t. A third basic property of induced homomorphisms is their invariance under such homotopies:

(iii) If $\varphi_t : (X, x_0) \to (Y, y_0)$ is a basepoint-preserving homotopy, then $\varphi_{0*} = \varphi_{1*}$.

This holds since $\varphi_{0*}[f] = [\varphi_0 f] = [\varphi_1 f] = \varphi_{1*}[f]$, the middle equality coming from the homotopy $\varphi_t f$.

There is a notion of homotopy equivalence for spaces with basepoints: $(X, x_0) \simeq (Y, y_0)$ if there are maps $f : (X, x_0) \to (Y, y_0)$ and $g : (Y, y_0) \to (X, x_0)$ with homotopies $fg \simeq \mathbb{1}$ and $gf \simeq \mathbb{1}$ through maps fixing the basepoints. In this case the induced maps on π_1 satisfy $f_* g_* = (fg)_* = \mathbb{1}_* = \mathbb{1}$ and likewise $g_* f_* = \mathbb{1}$, so f_* and g_* are inverse isomorphisms $\pi_1(X, x_0) \approx \pi_1(Y, y_0)$. This somewhat formal argument gives another proof that a deformation retraction induces an isomorphism on fundamental groups, since if X deformation retracts onto A then $(X, x_0) \simeq (A, x_0)$ for $x_0 \in A$.

One of the exercises below is to show that a homotopy equivalence $f : X \to Y$ induces an isomorphism $f_* : \pi_1(X, x_0) \to \pi_1(Y, f(x_0))$ even without any conditions on basepoints. This shouldn't be too surprising since two spaces are homotopy equivalent iff there is a third space containing them both as deformation retracts, as we showed in Corollary 0.6.

Exercises

When fundamental groups are written without basepoints, assume the spaces are path-connected.

1. Show that composition of paths satisfies the following cancellation property: If $f_0 \cdot g_0 \simeq f_1 \cdot g_1$ and $g_0 \simeq g_1$ then $f_0 \simeq f_1$.

2. For the change-of-basepoint homomorphisms β_h, show that $h \simeq h'$ implies $\beta_h = \beta_{h'}$.

3. Show that $\pi_1(X)$ is abelian iff all basepoint-change homomorphisms β_h depend only on the endpoints of the path h.

4. Let $X \subset \mathbb{R}^n$ be the union of convex open sets $X_1, \cdots, X_n$ such that $X_i \cap X_j \cap X_k \neq \varnothing$ for all i, j, k. Show that X is simply-connected.

5. A subspace $X \subset \mathbb{R}^n$ is said to be *star-shaped* if there is a point $x_0 \in X$ such that for all $x \in X$, the line segment from x_0 to x lies in X. Show that if a subspace $X \subset \mathbb{R}^n$ is locally star-shaped, i.e., every point of X has a star-shaped neighborhood in X, then every path in X is homotopic (in X) to a piecewise linear path, consisting of a finite number of straight line segments traversed at constant speed. Show this applies in particular when X is open or when X is a union of finitely many closed convex sets.

6. Show that every homomorphism $\pi_1(S^1) \to \pi_1(S^1)$ can be realized as the induced homomorphism φ_* of a map $\varphi : S^1 \to S^1$.

7. Given a space X and a path-connected subspace A containing the basepoint x_0, show that the map $\pi_1(A, x_0) \to \pi_1(X, x_0)$ induced by the inclusion $A \hookrightarrow X$ is surjective iff every path in X with endpoints in A is homotopic to a path in A.

8. Show there is no retraction of the Möbius band onto its boundary circle.

9. Show that for a space X, the following three conditions are equivalent:

(i) Every map $S^1 \to X$ is homotopic to a constant map, with image a point.

(ii) Every map $S^1 \to X$ extends to a map $D^2 \to X$.

(iii) $\pi_1(X, x_0) = 0$ for all $x_0 \in X$.

Deduce that a space X is simply-connected iff all maps $S^1 \to X$ are homotopic. [In this problem, 'homotopic' means 'homotopic without regard to basepoints.']

10. Let $[S^1, X]$ be the set of homotopy classes of maps $S^1 \to X$, with no conditions on basepoints. Thus there is a natural map $\Phi : \pi_1(X, x_0) \to [S^1, X]$ obtained by ignoring basepoints. Show that Φ is onto if X is path-connected, and that $\Phi([f]) = \Phi([g])$ iff $[f]$ and $[g]$ are conjugate in $\pi_1(X, x_0)$. Hence Φ induces a one-to-one correspondence between $[S^1, X]$ and the set of conjugacy classes in $\pi_1(X)$, when X is path-connected.

11. Define $f : S^1 \times I \to S^1 \times I$ by $f(\theta, s) = (\theta + 2\pi s, s)$, so f restricts to the identity on the two boundary circles of $S^1 \times I$. Show that f is homotopic to the identity by a homotopy f_t which is the identity on one of the boundary circles, but not by any homotopy f_t which is the identity on both boundary circles. [Consider what f does to the path $s \mapsto (\theta_0, s)$ for fixed $\theta_0 \in S^1$.]

12. Construct infinitely many nonhomotopic retractions $S^1 \vee S^1 \to S^1$.

13. Use the calculation of $\pi_1(\mathbb{R}^2 - \{0\})$ to show that $\mathbb{R}^2$ is not homeomorphic to $\mathbb{R}^n$ for $n > 2$.

14. Given (X, x_0) with X_0 the path-component of X containing x_0, show that the inclusion $X_0 \hookrightarrow X$ induces an isomorphism $\pi_1(X_0, x_0) \to \pi_1(X, x_0)$.

15. Let $\varphi_t : X \to Y$ be a homotopy and let h be the path $\varphi_t(x_0)$ traced out by a basepoint $x_0 \in X$. Show that $\beta_h \varphi_{1*} = \varphi_{0*}$ for $\varphi_{i*} : \pi_1(X, x_0) \to \pi_1(Y, \varphi_i(x_0))$ and $\beta_h : \pi_1(Y, \varphi_1(x_0)) \to \pi_1(Y, \varphi_0(x_0))$ the change-of-basepoint isomorphism. In particular, φ_{0*} is an isomorphism iff φ_{1*} is an isomorphism.

16. Using the previous problem, show that a homotopy equivalence $\varphi : X \to Y$ induces an isomorphism $\varphi_* : \pi_1(X, x_0) \to \pi_1(Y, \varphi(x_0))$ for any $x_0 \in X$.

17. If $X \subset \mathbb{R}^3$ is the union of the spheres of radius $1/n$ tangent to the yz-plane at the origin, $n = 1, 2, \cdots$, show that $\pi_1(X) = 0$.

18. Suppose $f_t : X \to X$ is a homotopy such that f_0 and f_1 are each the identity map. For a basepoint $x_0 \in X$, show that the loop $f_t(x_0)$ represents an element of the center of $\pi_1(X, x_0)$. [Hint: Consider the map $X \times S^1 \to X$ defined by the homotopy f_t.]

2. Van Kampen's Theorem

The van Kampen theorem gives a method for computing the fundamental groups of spaces which can be decomposed into simpler spaces whose fundamental groups are already known. The statement of the theorem involves the notion of free product of groups, so we begin by explaining this bit of algebra.

Free Products of Groups

Suppose one is given a collection of groups G_α and one wishes to construct a single group containing all these groups as subgroups. One way to do this would be to take the product group $\prod_\alpha G_\alpha$, whose elements can be regarded as the functions $\alpha \mapsto g_\alpha \in G_\alpha$. Or one could restrict to functions taking on non-identity values at most finitely often, forming the direct sum $\bigoplus_\alpha G_\alpha$. Both these constructions produce groups containing all the G_α's as subgroups, but with the property that elements of different subgroups G_α commute with each other. In the realm of nonabelian groups this commutativity is unnatural, and so one would like a 'nonabelian' version of $\prod_\alpha G_\alpha$ or $\bigoplus_\alpha G_\alpha$. Since the sum $\bigoplus_\alpha G_\alpha$ is smaller and presumably simpler than $\prod_\alpha G_\alpha$, it should be easier to construct a nonabelian version of $\bigoplus_\alpha G_\alpha$, and this is what the free product $*_\alpha G_\alpha$ achieves. Before explaining the general construction, let us look at two examples of the free product $G_\alpha * G_\beta$ of two groups.

Example 1: $\mathbb{Z}*\mathbb{Z}$. Let a and b denote generators of two copies of $\mathbb{Z}$. Then elements of $\mathbb{Z}*\mathbb{Z}$ are words like $a^2b^{-3}ab^5a^{-1}$, with powers of a and b alternating. Multiplication of such words is defined in the obvious way: $(a^3b^2a^{-2})(a^4b^{-1}) = a^3b^2a^2b^{-1}$. The identity element is the empty word e. Inverses are what they have to be: $(ab^2a^{-3}b)^{-1} = b^{-1}a^3b^{-2}a^{-1}$.

Example 2: $\mathbb{Z}_2*\mathbb{Z}_2$. This is like the preceding example but simpler since $a^2 = e = b^2$, so one doesn't need powers of a and b, and $\mathbb{Z}_2*\mathbb{Z}_2$ consists of just the alternating words in a and b: a, b, ab, ba, aba, bab, $abab$, $baba$, $ababa$, $\cdots$, together with the empty word again. The structure of $\mathbb{Z}_2*\mathbb{Z}_2$ can be elucidated by looking at the homomorphism $\varphi : \mathbb{Z}_2*\mathbb{Z}_2 \to \mathbb{Z}_2$ associating to each word its length $\mod 2$. Obviously φ is surjective, and its kernel consists of the words of even length. These form an infinite cyclic subgroup generated by ab since $ba = (ab)^{-1}$. In fact, $\mathbb{Z}_2 * \mathbb{Z}_2$ is the semi-direct product of the subgroups $\mathbb{Z}$ and $\mathbb{Z}_2$ generated by ab and a, with the conjugation relation $a(ab)a^{-1} = (ab)^{-1}$. This group is sometimes called the infinite dihedral group.

Now let us define the general free product $*_\alpha G_\alpha$ of groups G_α. As a set, $*_\alpha G_\alpha$ consists of all words $g_1g_2\cdots g_m$ of finite length $m \geq 0$, where each letter g_i belongs to a

group G_{α_i} and is not the identity element of G_{α_i}, and adjacent letters g_i and g_{i+1} belong to different groups G_α. Words satisfying these conditions are called *reduced*, the idea being that nonreduced words can always be simplified until they are reduced, by writing adjacent letters which lie in the same G_{α_i} as a single letter and by cancelling trivial letters. The empty word is allowed, and will be the identity element of $*_\alpha G_\alpha$. The group operation in $*_\alpha G_\alpha$ is juxtaposition, $(g_1 \cdots g_m)(h_1 \cdots h_n) = g_1 \cdots g_m h_1 \cdots h_n$. This product may not be reduced, however: If g_m and h_1 belong to the same G_α they should be combined into a single letter $(g_m h_1)$ according to the multiplication in G_α, and if this new letter $g_m h_1$ happens to be the identity of G_α, it should be cancelled from the product. This may allow g_{m-1} and h_2 to be combined, and possibly cancelled too, etc. For example, in the product $(g_1 \cdots g_m)(g_m^{-1} \cdots g_1^{-1})$ everything cancels and we get the identity element of $*_\alpha G_\alpha$, the empty word.

Verifying directly that this multiplication is associative would be rather tedious, but there is an indirect approach which avoids most of the work. Let W be the set of reduced words $g_1 \cdots g_m$ as above, including the empty word. To each $g \in G_\alpha$ we associate the function $L_g : W \to W$ given by multiplication on the left, $L_g(g_1 \cdots g_m) = g g_1 \cdots g_m$ where we combine g with g_1 if $g_1 \in G_\alpha$ to make $g g_1 \cdots g_m$ a reduced word. A key property of the association $g \mapsto L_g$ is the formula $L_{gg'} = L_g L_{g'}$ for $g, g' \in G_\alpha$, i.e., $g(g'(g_1 \cdots g_m)) = (gg')(g_1 \cdots g_m)$. This special case of associativity follows rather trivially from associativity in G_α. The formula $L_{gg'} = L_g L_{g'}$ implies that L_g is invertible with inverse $L_{g^{-1}}$. Therefore the association $g \mapsto L_g$ defines a homomorphism from G_α to the group $P(W)$ of all permutations of W. More generally, we can define $L : W \to P(W)$ by $L(g_1 \cdots g_m) = L_{g_1} \cdots L_{g_m}$ for each reduced word $g_1 \cdots g_m$. This function L is injective since the permutation $L(g_1 \cdots g_m)$ sends the empty word to $g_1 \cdots g_m$. The product operation in W corresponds under L to composition in $P(W)$; this follows from the relation $L_{gg'} = L_g L_{g'}$. Since composition in $P(W)$ is associative, we conclude that the product in W is associative.

Each group G_α is naturally identified with a subgroup of $*_\alpha G_\alpha$, the subgroup consisting of the empty word and the nonidentity one-letter words $g \in G_\alpha$. From this viewpoint, the empty word is the common identity element of all the subgroups G_α, which are otherwise disjoint. A consequence of associativity is that any product $g_1 \cdots g_m$ of elements g_i in the groups G_α has a unique reduced form, the element of $*_\alpha G_\alpha$ obtained by performing the multiplications in any order. In fact, any sequence of reduction operations on an unreduced product $g_1 \cdots g_m$, combining adjacent letters g_i and g_{i+1} which lie in the same G_α or cancelling a g_i which is the identity, can be viewed as a way of inserting parentheses into $g_1 \cdots g_m$ and performing the resulting sequence of multiplications. Thus

associativity implies the non-obvious fact that any two sequences of reduction operations performed on the same unreduced word always yield the same reduced word in the end.

A basic property of the free product $*_\alpha G_\alpha$ is that any collection of homomorphisms $\varphi_\alpha : G_\alpha \to H$ always extends uniquely to a homomorphism $\varphi : *_\alpha G_\alpha \to H$. Namely, for a word $g_1 \cdots g_n$ with $g_i \in G_{\alpha_i}$ we must have $\varphi(g_1 \cdots g_n) = \varphi_{\alpha_1}(g_1) \cdots \varphi_{\alpha_n}(g_n)$, and this gives a well-defined homomorphism since the process of reducing an unreduced product in $*_\alpha G_\alpha$ does not affect its image under φ. For example, for a free product $G * H$ the inclusions $G \hookrightarrow G \times H$ and $H \hookrightarrow G \times H$ induce a surjective homomorphism $G * H \to G \times H$.

The van Kampen Theorem

Suppose a space X is decomposed as the union of a collection of path-connected open subsets A_α, each containing the basepoint x_0. By the remarks made in the preceding paragraph, the homomorphisms $j_\alpha : \pi_1(A_\alpha) \to \pi_1(X)$ induced by the inclusions $A_\alpha \hookrightarrow X$ then extend to a homomorphism $\Phi : *_\alpha \pi_1(A_\alpha) \to \pi_1(X)$. We can expect Φ to have a nontrivial kernel in general, since if $i_{\alpha\beta} : \pi_1(A_\alpha \cap A_\beta) \to \pi_1(A_\alpha)$ is the homomorphism induced by the inclusion $A_\alpha \cap A_\beta \hookrightarrow A_\alpha$ then $j_\alpha i_{\alpha\beta} = j_\beta i_{\beta\alpha}$, both these compositions being induced by the inclusion $A_\alpha \cap A_\beta \hookrightarrow X$; hence the kernel of Φ contains all the elements of the form $i_{\alpha\beta}(\omega) i_{\beta\alpha}(\omega)^{-1}$ for $\omega \in \pi_1(A_\alpha \cap A_\beta)$.

Now we can state van Kampen's theorem:

Theorem 1.12. *If X is the union of path-connected open sets A_α each containing the basepoint $x_0 \in X$ and if each intersection $A_\alpha \cap A_\beta$ is path-connected, then $\Phi : *_\alpha \pi_1(A_\alpha) \to \pi_1(X)$ is surjective. If in addition each intersection $A_\alpha \cap A_\beta \cap A_\gamma$ is path-connected, then the kernel of Φ is the normal subgroup N generated by all elements of the form $i_{\alpha\beta}(\omega) i_{\beta\alpha}(\omega)^{-1}$, and so Φ induces an isomorphism $\pi_1(X) \approx *_\alpha \pi_1(A_\alpha)/N$.*

Example: Wedge sums. In §0.2 we defined the wedge sum $\bigvee_\alpha X_\alpha$ of a collection of spaces X_α with basepoints $x_\alpha \in X_\alpha$ to be the quotient space of the disjoint union $\coprod_\alpha X_\alpha$ in which all the basepoints x_α are identified to a single point. If each x_α is a deformation retract of an open neighborhood U_α in X_α, then $A_\alpha = X_\alpha \bigvee_{\beta \neq \alpha} U_\beta$ is an open neighborhood of X_α in $\bigvee_\alpha X_\alpha$ which deformation retracts onto X_α, and the intersection of any two or more A_α's is $\bigvee_\alpha U_\alpha$ which is contractible. Hence van Kampen's theorem implies that $\Phi : *_\alpha \pi_1(X_\alpha) \to \pi_1(\bigvee_\alpha X_\alpha)$ is an isomorphism.

In particular, for a wedge sum $\bigvee_\alpha S^1_\alpha$ of circles, $\pi_1(\bigvee_\alpha S^1_\alpha)$ is a **free group**, the free product of copies of $\mathbb{Z}$, one for each circle S^1_α. For instance, $\pi_1(S^1 \vee S^1) \approx \mathbb{Z} * \mathbb{Z}$.

Van Kampen's theorem is often applied when there are just two sets A_α and A_β in the cover of X, so the condition on triple intersections $A_\alpha \cap A_\beta \cap A_\gamma$ is vacuous and one

obtains the isomorphism $\pi_1(X) \approx \pi_1(A_\alpha) * \pi_1(A_\beta)/N$, under the assumption that $A_\alpha \cap A_\beta$ is path-connected. The proof in this special case is virtually identical with the proof in the general case, however.

One can see that the intersections $A_\alpha \cap A_\beta$ need to be path-connected by considering the example of S^1 decomposed as the union of two open arcs. In this case Φ is not surjective. For an example showing that triple intersections $A_\alpha \cap A_\beta \cap A_\gamma$ need to be path-connected, let X be the graph consisting of two vertices joined by three edges, i.e., the union of a circle with a diameter, and let A_α, A_β, and A_γ be the complements of points x_α, x_β, and x_γ in the interiors of the three edges. The theorem does apply to the covering $\{A_\alpha, A_\beta\}$, so $\pi_1(X) \approx \pi_1(A_\alpha) * \pi_1(A_\beta) \approx \mathbb{Z} * \mathbb{Z}$ since $A_\alpha \cap A_\beta$ is contractible. If we tried to use the covering $\{A_\alpha, A_\beta, A_\gamma\}$ which has each of the two-fold intersections path-connected but not the triple intersection, then we would get $\pi_1(X) \approx \mathbb{Z} * \mathbb{Z} * \mathbb{Z}$. But the groups $\mathbb{Z} * \mathbb{Z}$ and $\mathbb{Z} * \mathbb{Z} * \mathbb{Z}$ are not isomorphic since their abelianizations are $\mathbb{Z} \times \mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}$, which are not isomorphic.

Proof of van Kampen's theorem: First we show Φ is surjective. Given a loop $f : I \to X$ at the basepoint x_0, we may choose as in the proof of Proposition 1.5 a partition $0 = s_0 < s_1 < \cdots < s_m = 1$ so that each subinterval $[s_{i-1}, s_i]$ is mapped by f to a single A_α, which we denote A_i. Let f_i be the path $f|[s_{i-1}, s_i]$, so f is the composition $f_1 \cdot \cdots \cdot f_m$. Since $A_i \cap A_{i+1}$ is path-connected, we may choose a path g_i in $A_i \cap A_{i+1}$ from x_0 to the point $f(s_i) \in A_i \cap A_{i+1}$. Inserting $\overline{g}_i \cdot g_i$ into f at s_i for each i then yields a loop

$$(f_1 \cdot \overline{g}_1) \cdot (g_1 \cdot f_2 \cdot \overline{g}_2) \cdot (g_2 \cdot f_3 \cdot \overline{g}_3) \cdot \cdots \cdot (g_{m-1} \cdot f_m)$$

homotopic to f. This loop is a composition of loops each lying in a single A_i, so $[f]$ is in the image of Φ, and Φ is surjective.

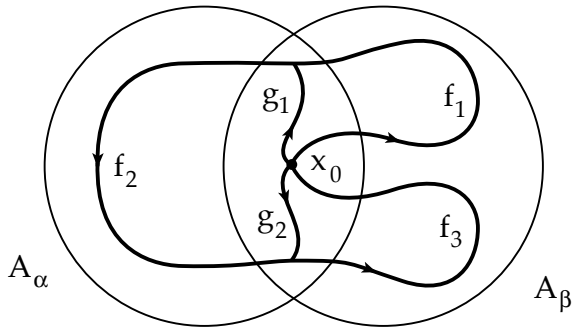

The harder part of the proof is to show that the kernel of Φ is N. Suppose we are given a reduced word $w = w_1 \cdots w_k$ in $*_\alpha \pi_1(A_\alpha)$ such that $\Phi(w)$ is trivial. If we can

show that w has trivial image in $*_\alpha \pi_1(A_\alpha)/N$, then this will imply that the kernel of Φ is contained in N. The opposite inclusion has already been noted, so this will finish the proof.

Corresponding to the word $w = w_1 \cdots w_k$ there is a product $f = f_1 \cdots f_k$ of loops f_i lying in A_α's. Since $\Phi(w)$ is trivial, there is a map $F : I \times I \to X$ defining a homotopy from f to the constant loop. There exist partitions $0 = s_0 < s_1 < \cdots < s_m = 1$ and $0 = t_0 < t_1 < \cdots < t_n = 1$ such that each rectangle $R_{ij} = [s_{i-1}, s_i] \times [t_{j-1}, t_j]$ is mapped by F into a single A_α, which we label A_{ij}. These partitions may be obtained by covering $I \times I$ by finitely many rectangles $[a, b] \times [c, d]$ each mapping to a single A_α, using a compactness argument, then partitioning $I \times I$ by the union of all the horizontal and vertical lines containing edges of these rectangles. We may assume the s_i-partition subdivides the partition decomposing f as $f_1 \cdots f_k$. Since F maps a neighborhood of R_{ij} to A_{ij}, we may perturb the vertical sides of the rectangles R_{ij} so that each point of $I \times I$ lies in at most three R_{ij}'s. We can do this perturbation just on the rectangles R_{ij} with $j > 1$, leaving the rectangles R_{i1} in the bottom row unchanged.

Let f_{ij} be the loop obtained by restricting F to the path in $I \times I$ which goes in linear segments from $(0, t_j)$ to the upper right corner of R_{ij}, then to the lower right corner of R_{ij}, then to $(1, t_{j-1})$.

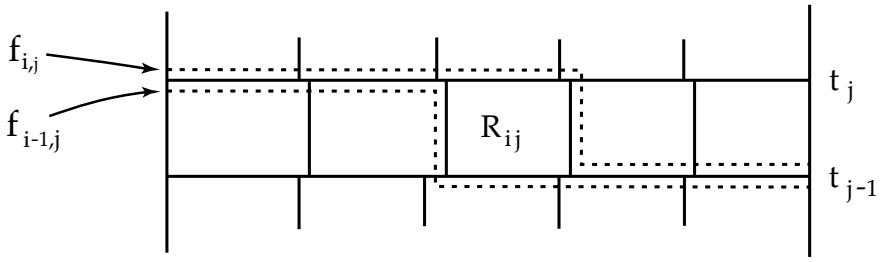

The restriction of F to the rectangle R_{ij} gives a homotopy from $f_{i-1,j}$ to f_{ij}, and in this way the homotopy F is broken up into a sequence of smaller homotopies:

$$f \simeq f_{01} \simeq f_{11} \simeq f_{21} \simeq \cdots \simeq f_{m1} \simeq f_{02} \simeq f_{12} \simeq f_{22} \simeq \cdots \simeq f_{mn}$$

Some of these homotopies, for example $f \simeq f_{01}$ and $f_{m1} \simeq f_{02}$, are just reparametrizations since the left and right edges of $I \times I$ map to the basepoint x_0. The final path f_{mn} is the constant loop.

Let us call the corners of the R_{ij}'s *vertices*. For each vertex v with $F(v) \neq x_0$, choose a path g_v from x_0 to $F(v)$ in the intersection of the two or three A_{ij}'s corresponding to the R_{ij}'s containing v. If we insert into f_{ij} the appropriate paths $\overline{g}_v g_v$ at successive vertices, as in the first paragraph of the proof, then we obtain a product $f_{ij}^1 f_{ij}^2 \cdots$ homotopic to

f_{ij}, with each f_{ij}^k a loop in a single A_α, namely the A_{pq} corresponding to one of the at most two R_{pq}'s containing the relevant horizontal or vertical segment. Making a choice of such an A_α for each f_{ij}^k gives a way of writing $[f_{ij}] \in \pi_1(X)$ as $\Phi(w_{ij})$ for a word $w_{ij} = w_{ij}^1 w_{ij}^2 \cdots$ in $*_\alpha \pi_1(A_\alpha)$, with $w_{ij}^k = [f_{ij}^k]$. Making a different choice A_β instead of A_α changes the corresponding w_{ij}^k from $w_{ij}^k = i_{\alpha\beta}(\omega) \in \pi_1(A_\alpha)$ to $w_{ij}^k = i_{\beta\alpha}(\omega) \in \pi_1(A_\beta)$. Hence w_{ij} changes from a product of the form $\omega_1 i_{\alpha\beta}(\omega)\omega_2$ to $\omega_1 i_{\beta\alpha}(\omega)\omega_2$. Modulo the normal subgroup N, these two w_{ij}'s are equivalent since

$$\omega_1 i_{\alpha\beta}(\omega)\omega_2 = [\omega_1 i_{\alpha\beta}(\omega) i_{\beta\alpha}(\omega)^{-1} \omega_1^{-1}]\omega_1 i_{\beta\alpha}(\omega)\omega_2$$

and the bracketed expression is in N.

The homotopy $f_{i-1,j} \simeq f_{ij}$ across R_{ij} changes the subword of $w_{i-1,j}$ corresponding to the left and bottom edges of R_{ij} to the subword of w_{ij} corresponding to the upper and right edges of R_{ij}. If we choose the terms in these subwords to lie in $\pi_1(A_{ij})$, then the homotopy across R_{ij} shows that these two subwords define the same element of $\pi_1(A_{ij})$, hence $w_{i-1,j} = w_{ij}$ in $*_\alpha \pi_1(A_\alpha)$. Thus we have shown that the image of w_{ij} in $*_\alpha \pi_1(A_\alpha)/N$ is independent of i and j.

We may assume that the homotopy $F(s,t)$ is independent of t in the bottom row of rectangles R_{i1}. Then since the s_i-partition subdivides the partition defining the factorization $f = f_1 \cdots f_k$, the bottom edge of each R_{i1} is contained in the domain of an f_ℓ and we may choose A_{i1} to be the A_α for which $[f_\ell] \in \pi_1(A_\alpha)$. The effect of this is that the word w_{01} reduces to $w = w_1 \cdots w_k$ in $*_\alpha \pi_1(A_\alpha)$. The final word w_{mn} is the identity element of $*_\alpha \pi_1(A_\alpha)$ since f_{mn} is the constant loop. Applying the conclusion of the preceding paragraph, we deduce that w has trivial image in $*_\alpha \pi_1(A_\alpha)/N$, finishing the proof. □

Example: Linking of Circles. We can apply van Kampen's theorem to calculate the fundamental groups of three spaces discussed in the introduction to this chapter, the complements in $\mathbb{R}^3$ of a single circle, two unlinked circles, and two linked circles.

(1) The complement $\mathbb{R}^3 - A$ of a single circle A deformation retracts onto a wedge sum $S^1 \vee S^2$ embedded in $\mathbb{R}^3 - A$ as shown in the first of the following figures. It is perhaps easier to see that $\mathbb{R}^3 - A$ deformation retracts onto the the union of S^2 with a diameter, as in the second figure: Points outside S^2 deformation retract onto S^2, and then points inside S^2 and not in A can be pushed away from A toward S^2 or the diameter.

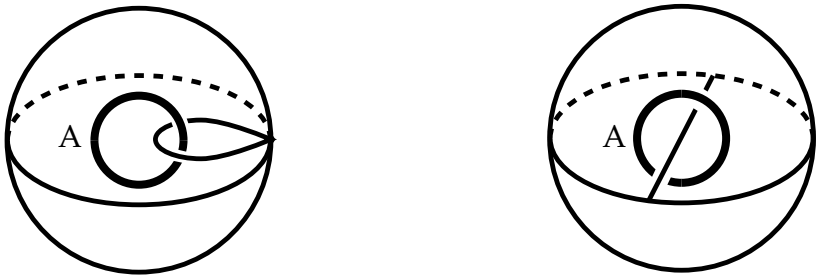

Having this deformation retraction in mind, one can then see how it must be modified if the two endpoints of the diameter are gradually moved towards each other along the equator until they coincide, forming the S^1. Another way of seeing the deformation retraction of $\mathbb{R}^3 - A$ onto $S^1 \vee S^2$ is to note first that an open ε-neighborhood of $S^1 \vee S^2$ obviously deformation retracts onto $S^1 \vee S^2$ if ε is sufficiently small. Then one observes that this neighborhood is homeomorphic to $\mathbb{R}^3 - A$ by a homeomorphism which is the identity on $S^1 \vee S^2$. In fact, the neighborhood can be gradually enlarged by homeomorphisms until it becomes all of $\mathbb{R}^3 - A$.

In any event, once we see that $\mathbb{R}^3 - A$ deformation retracts to $S^1 \vee S^2$, then we immediately obtain isomorphisms $\pi_1(\mathbb{R}^3 - A) \approx \pi_1(S^1 \vee S^2) \approx \mathbb{Z}$ since $\pi_1(S^2) = 0$.

(2) In similar fashion, the complement $\mathbb{R}^3 - (A \cup B)$ of two unlinked circles A and B deformation retracts onto $S^1 \vee S^1 \vee S^2 \vee S^2$, as in the figure on the left below. From this we get $\pi_1\big(\mathbb{R}^3 - (A \cup B)\big) \approx \mathbb{Z} * \mathbb{Z}$.

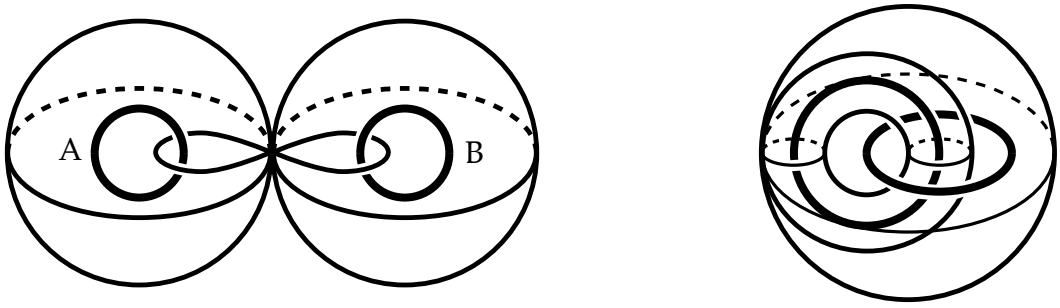

(3) If A and B are linked, then $\mathbb{R}^3 - (A \cup B)$ deformation retracts onto the wedge sum of S^2 and a torus $S^1 \times S^1$ separating A and B, as shown in the right-hand figure, so $\pi_1\big(\mathbb{R}^3 - (A \cup B)\big) \approx \mathbb{Z} \times \mathbb{Z}$ in this case.

Example: The Shrinking Wedge of Circles. Consider the subspace $X \subset \mathbb{R}^2$ which is the union of the circles C_n of radius $1/n$ and center $(1/n, 0)$ for $n = 1, 2, \cdots$. With the subspace topology from $\mathbb{R}^2$, an open neighborhood of the origin in X must contain all but finitely many of the circles C_n, so the subspace topology on X has fewer open sets than the usual wedge sum $\bigvee_\infty S^1$ of an infinite sequence of circles, which is a quotient space of $\coprod_\infty S^1$.

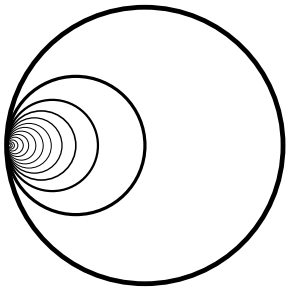

This difference in topologies has a big effect on $\pi_1(X)$, making it much larger than $\pi_1(\bigvee_\infty S^1)$. To see this, consider the retractions $r_n : X \to C_n$ collapsing all C_i's except C_n to the origin. Each r_n induces a surjection $\rho_n : \pi_1(X) \to \pi_1(C_n) \approx \mathbb{Z}$, where we take the origin as the basepoint. The product of the ρ_n's is a homomorphism $\rho : \pi_1(X) \to \prod_\infty \mathbb{Z}$ to the direct product (not the direct sum) of infinitely many copies of $\mathbb{Z}$, and ρ is surjective since for every sequence of integers k_n we can construct a loop $f : I \to X$ which wraps k_n times around C_n in the time interval $[1 - 1/n, 1 - 1/(n+1)]$. This infinite composition of loops is continuous since X has the subspace topology from $\mathbb{R}^2$. Since $\prod_\infty \mathbb{Z}$ is uncountable, so is $\pi_1(X)$. On the other hand, $\pi_1(\bigvee_\infty S^1)$ is free on countably many generators, hence is countable. Loops in $\bigvee_\infty S^1$ can go around only finitely many of the circles since these circles are not shrinking to a point as they are in X.

The group $\pi_1(X)$ is actually far more complicated than $\prod_\infty \mathbb{Z}$. For one thing, it is nonabelian, unlike $\prod_\infty \mathbb{Z}$, since the retraction $X \to C_1 \cup \cdots \cup C_n$ which collapses all the circles smaller than C_n to the basepoint induces a surjection from $\pi_1(X)$ to a free group on n generators. An algebraic description of $\pi_1(X)$ can be found in [Morgan-Morrison].

Applications to Cell Complexes

For the remainder of this section we shall be interested in 2-dimensional cell complexes, analyzing how the fundamental group is affected by attaching 2-cells. According to an exercise at the end of this section, attaching cells of higher dimension has no effect on π_1, so all the interest lies in how the 2-cells are attached.

Suppose we attach a collection of 2-cells e^2_α to a path-connected space X via maps $\varphi_\alpha : S^1 \to X$, producing a space Y. Let s_0 be a basepoint of S^1 and x_0 a basepoint of X. If $\varphi_\alpha(s_0) = x_0$ then φ_α determines an element of $\pi_1(X, x_0)$, but in general, if $\varphi_\alpha(s_0) \neq x_0$, then to make φ_α into a loop at x_0 we choose a path γ_α in X from x_0 to $\varphi_\alpha(s_0)$, and then $\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha$ is a loop at x_0. This loop may not be nullhomotopic in X, but it will certainly be nullhomotopic in Y, after the cell e^2_α is attached. Thus the

normal subgroup $N \subset \pi_1(X, x_0)$ generated by all the loops $\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha$ for varying α lies in the kernel of the map $\pi_1(X, x_0) \to \pi_1(Y, x_0)$ induced by the inclusion $X \hookrightarrow Y$.

Proposition 1.13. *The inclusion $X \hookrightarrow Y$ induces a surjection $\pi_1(X, x_0) \to \pi_1(Y, x_0)$ whose kernel is N. Thus $\pi_1(Y) \approx \pi_1(X)/N$.*

It follows that N is independent of the choice of the paths γ_α, but this can also be seen directly: If we replace γ_α by another path η_α having the same endpoints, then $\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha$ changes to $\eta_\alpha \varphi_\alpha \overline{\eta}_\alpha = (\eta_\alpha \overline{\gamma}_\alpha)\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha (\gamma_\alpha \overline{\eta}_\alpha)$, so $\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha$ and $\eta_\alpha \varphi_\alpha \overline{\eta}_\alpha$ define conjugate elements of $\pi_1(X, x_0)$.

Proof: Let us expand Y to a slightly larger space Z which deformation retracts onto Y and is more convenient for applying van Kampen's theorem. The space Z is obtained from Y by attaching rectangular strips $S_\alpha = I \times I$, with the lower edge $I \times \{0\}$ attached along γ_α, the right edge $\{1\} \times I$ attached along an arc in e^2_α, and all the left edges $\{0\} \times I$ of the different strips identified together. The top edges of the strips are not attached to anything, and this allows us to deformation retract Z onto Y.

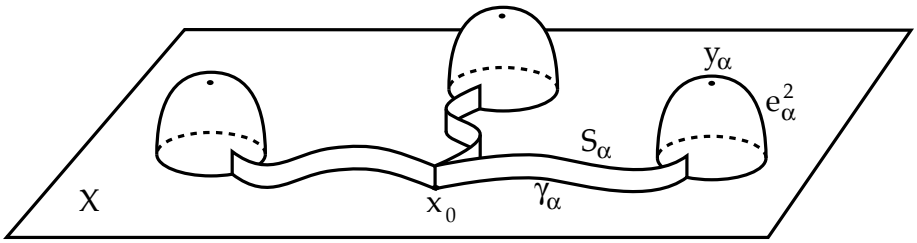

Next, choose a point y_α in each cell e^2_α and not in S_α. Let $A = Z - \bigcup_\alpha \{y_\alpha\}$ and $B = Z - X$. Then A deformation retracts onto X, and B is contractible. Since $\pi_1(B) = 0$, van Kampen's theorem applied to the cover $\{A, B\}$ says that $\pi_1(Z)$ is isomorphic to the quotient of $\pi_1(A)$ by the normal subgroup generated by the image of the map $\pi_1(A \cap B) \to \pi_1(A)$. So it remains only to see that $\pi_1(A \cap B)$ is generated by the loops $\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha$, or rather by loops in $A \cap B$ homotopic to these loops. This can be shown by another application of van Kampen's theorem, this time to the cover of $A \cap B$ by the open sets $A_\alpha = A \cap B - \bigcup_{\beta \neq \alpha} e^2_\beta$. Since A_α deformation retracts onto a circle in $e^2_\alpha - \{y_\alpha\}$, we have $\pi_1(A_\alpha) \approx \mathbb{Z}$ generated by a loop homotopic to $\gamma_\alpha \varphi_\alpha \overline{\gamma}_\alpha$, and the result follows. □

As a first application we compute the fundamental group of the orientable surface M_g of genus g. This has a cell structure with one 0-cell, $2g$ 1-cells, and one 2-cell, as we noted in § 0.2. The 1-skeleton is a wedge sum of $2g$ circles, with fundamental group free on $2g$ generators, and the 2-cell is attached along the loop given by the product of the

commutators of these generators, say $[a_1, b_1] \cdots [a_g, b_g]$. Therefore

$$\pi_1(M_g) \approx \langle\, a_1, b_1, \cdots, a_g, b_g \mid [a_1, b_1] \cdots [a_g, b_g] \,\rangle$$

where $\langle\, g_\alpha \mid r_\beta \,\rangle$ denotes the group with generators g_α and relators r_β, in other words the free group on the generators g_α modulo the normal subgroup generated by the words r_β in these generators.

Corollary 1.14. *The surface M_g is not homeomorphic, or even homotopy equivalent, to M_h if $g \neq h$.*

Proof: The abelianization of $\pi_1(M_g)$ is the direct sum of $2g$ copies of $\mathbb{Z}$. So if $M_g \simeq M_h$ then $\pi_1(M_g) \approx \pi_1(M_h)$, hence the abelianizations of these groups are isomorphic, which implies $g = h$. □

Non-orientable surfaces can be treated in the same way. If we attach a 2-cell to the wedge sum of g circles by the word $a_1^2 \cdots a_g^2$ we obtain a non-orientable surface N_g. For example, N_1 is the projective plane $\mathbb{R}\mathrm{P}^2$, the quotient of D^2 with antipodal points of ∂D^2 identified. And N_2 is the Klein bottle, though the more usual representation of the Klein bottle is as a square with opposite sides identified via the word $aba^{-1}b$. If one cuts the square along a diagonal and reassembles the resulting two triangles as shown in the figure, one obtains the other representation as a square with sides identified via the word a^2c^2.

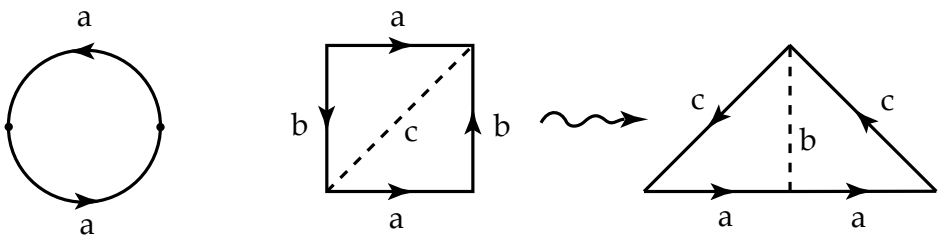

By the proposition, $\pi_1(N_g) \approx \langle\, a_1, \cdots, a_g \mid a_1^2 \cdots a_g^2 \,\rangle$. This abelianizes to the direct sum of $\mathbb{Z}_2$ with $g-1$ copies of $\mathbb{Z}$ since in the abelianization we can rechoose the generators to be $a_1, \cdots, a_{g-1}$ and $a_1 + \cdots + a_g$, with $2(a_1 + \cdots + a_g) = 0$. Hence N_g is not homotopy equivalent to N_h if $g \neq h$, nor is N_g homotopy equivalent to any orientable surface M_h.

Another consequence of the preceding proposition is:

Corollary 1.15. *For every group G there is a 2-dimensional cell complex X_G with $\pi_1(X_G) \approx G$.*

Proof: Choose a presentation $G = \langle\, g_\alpha \mid r_\beta \,\rangle$. This exists since every group is a quotient of a free group, so the g_α's can be taken to be the generators of this free group with the r_β's generators of the kernel of the map from the free group to G. Now construct X_G from $\bigvee_\alpha S^1_\alpha$ by attaching 2-cells e^2_β by the loops specified by the words r_β. □

Example. If $G = \langle\, a \mid a^n \,\rangle = \mathbb{Z}_n$ then X_G is S^1 with a cell e^2 attached by the map $z \mapsto z^n$, thinking of S^1 as the unit circle in $\mathbb{C}$. In case $n = 2$ we get $X_G = \mathbb{R}\mathrm{P}^2$, but for $n > 2$ the space X_G is not a surface since there are n 'sheets' of e^2 attached at each point of the circle $S^1 \subset X_G$. For example, when $n = 3$ one can construct a neighborhood N of S^1 in X_G by taking the product of the letter 'Y' with the interval I, then identifying the two ends of this product via a one-third twist as shown in the figure.

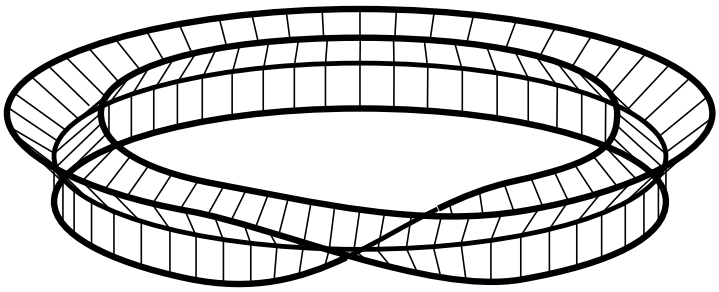

The boundary of N consists of a single circle, formed by the three endpoints of each 'Y' cross-section of N. To complete the construction of X_G from N one attaches a disk along the boundary circle of N. This cannot be done in $\mathbb{R}^3$, though it can in $\mathbb{R}^4$.

For $n = 4$ one would use the letter 'X' instead of 'Y,' with a one-quarter twist instead of a one-third twist. Similarly, for larger n one would use an n-pointed 'asterisk' and a $1/n$ twist.

Exercises

1. Show that the free product $G * H$ of nontrivial groups G and H has trivial center, and that the only elements of $G * H$ of finite order are the conjugates of finite-order elements of G and H.

2. Suppose a point x in a path-connected space X is a closed subset and has an open neighborhood homeomorphic to $\mathbb{R}^n$. Use the van Kampen theorem to show that the inclusion $X - \{x\} \subset X$ induces an isomorphism on π_1 if $n \geq 3$.

3. Suppose the space Y is obtained from the subspace X by attaching n-cells for a fixed $n \geq 3$. Show that the inclusion $X \hookrightarrow Y$ induces an isomorphism on π_1. [See the proof of 1.13.]

4. Let X be the quotient space of S^2 obtained by identifying the north and south poles to a single point. Put a cell complex structure on X and use this to compute $\pi_1(X)$.

5. Compute the fundamental group of the space obtained from two tori $S^1 \times S^1$ by identifying the two circles $S^1 \times \{x_0\}$ to a single circle.

6. Let C be a circle in the surface M_g of genus g which separates M_g into two compact subsurfaces M_h' and M_k' obtained from the closed surfaces M_h and M_k by deleting an open disk from each, as in the figure below.

(a) Show that M_h' does not retract onto its boundary circle C, and hence M_g does not retract onto C. [Hint: abelianize π_1.]

(b) Show that M_g does retract onto the circle C' shown in the figure.

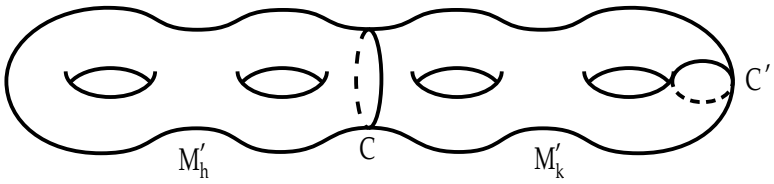

7. Consider two arcs α and β embedded in $D^2 \times I$ as shown in the figure below. The loop γ is obviously nullhomotopic in $D^2 \times I$, but show that there is no nullhomotopy of γ in the complement of $\alpha \cup \beta$.

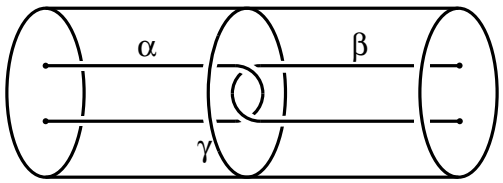

8. Let Y be the subspace of D^2 obtained by deleting the interiors of two disjoint subdisks in the interior of D^2, and let X be the quotient space of Y obtained by identifying its three boundary circles via homeomorphisms preserving the clockwise orientations of these circles. Construct a cell complex structure on X with one 0-cell, three 1-cells, and one 2-cell, and use this to compute $\pi_1(X)$ as a group with three generators and one relation. Observe that X is homeomorphic to a subspace of 'the house with two rooms' described in §0.1, namely the complement of the two open rectangles supporting the tunnels. Deduce that $\pi_1(X)$ is isomorphic to the fundamental group of the complement of the graph in $\mathbb{R}^3$ shown in the figure below, consisting of two circles joined by an arc.

[It can be shown that $\pi_1(X)$ is not isomorphic to the free group of two generators, which is the fundamental group of the complement of the more obvious embedding of the same graph, having image contained in a plane.]

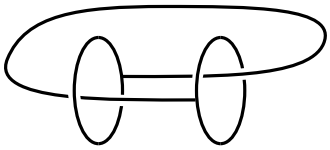

9. Consider the quotient space of a cube I^3 obtained by identifying each square face with the opposite square face via the right-handed screw motion consisting of a translation by one unit in the direction perpendicular to the face combined with a one-quarter twist of the face about its center point. Show this quotient space X is a cell complex with two 0-cells, four 1-cells, three 2-cells, and one 3-cell. Using this structure, show that $\pi_1(X)$ is the quaternion group $\{\pm 1, \pm i, \pm j, \pm k\}$, of order eight.

10. Given a space X with basepoint $x_0 \in X$, we may construct a cell complex $L(X)$ having a single 0-cell, a 1-cell for each loop in X based at x_0, and a 2-cell for each map of a standard triangle T into X taking the three vertices to the basepoint. Such a 2-cell is attached to the three 1-cells which are the loops obtained by restricting the map to the three edges of T. Show that $\pi_1(L(X))$ is isomorphic to $\pi_1(X, x_0)$ via an isomorphism induced by a natural map $L(X) \to X$.

11. Show that the fundamental group of the surface of infinite genus shown below is free on an infinite number of generators.

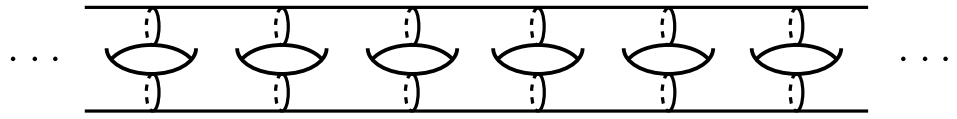

12. Show that $\pi_1(\mathbb{R}^2 - \mathbb{Q}^2)$ is uncountable.

13. Show that the join $X * Y$ of two non-empty spaces X and Y is simply-connected if X is path-connected.

3. Covering Spaces

We come now to the second main topic of this chapter, covering spaces. We have in fact already encountered one example of a covering space in our calculation of $\pi_1(S^1)$. This was the map $\mathbb{R} \to S^1$ which we pictured as the projection of a helix onto a circle. One thinks of the helix as lying above, or 'covering,' the circle, so the helix is said to be a 'covering space' of the circle. A number of things we proved for this helix example hold true for all covering spaces, and this allows covering spaces to serve as a useful tool for calculating fundamental groups in general. But there is much more to the story than this, and in fact the connection between fundamental group and covering spaces is so close that in many ways they can be regarded as two viewpoints toward the same thing. This means that algebraic features of the fundamental group can often be translated into the geometric language of covering spaces. One of the main results in this section provides a strong instance of this, as we shall see that the various covering spaces of a given space X correspond exactly to subgroups of $\pi_1(X)$. This is strikingly reminiscent of Galois theory, with its correspondence between field extensions and subgroups of the Galois group.

Let us begin with the definition. A **covering space** of a space X is a space $\widetilde{X}$ together with a map $p\colon \widetilde{X} \to X$ satisfying the following condition: There exists an open cover $\{U_\alpha\}$ of X such that for each α, $p^{-1}(U_\alpha)$ is a disjoint union of open sets in $\widetilde{X}$, each of which is mapped by p homeomorphically onto U_α.

In the helix example one has $p\colon \mathbb{R} \to S^1$ given by $p(t) = (\cos 2\pi t, \sin 2\pi t)$, and the cover $\{U_\alpha\}$ can be taken to consist of any two open arcs whose union is S^1. A related example is the helicoid surface $S \subset \mathbb{R}^3$ consisting of points $(s\cos 2\pi t, s\sin 2\pi t, t)$ for $(s,t) \in (0,\infty) \times \mathbb{R}$. This projects onto $\mathbb{R}^2 - \{0\}$ via $(x,y,z) \mapsto (x,y)$, and this projection defines a covering space $p\colon S \to \mathbb{R}^2 - \{0\}$ since for each open disk U in $\mathbb{R}^2 - \{0\}$, $p^{-1}(U)$ consists of countably many disjoint open disks in S each mapped homeomorphically onto U by p.

Another example is the map $p\colon S^1 \to S^1$, $p(z) = z^n$ where we view z as a complex number with $|z| = 1$ and n is any positive integer. This example, unfortunately, cannot be realized as a linear projection in 3-space analogous to the projection of the helix. The best one can do is draw a circle wrapping around a cylinder n times and intersecting itself in $n-1$ points which one just has to imagine are not really intersections. As our general theory will show, these examples for $n \geq 1$, together with the helix example, exhaust all the connected coverings spaces of S^1. There are many more disconnected examples, such as n disjoint

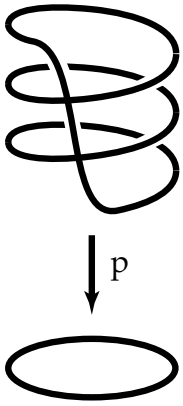

circles each mapped homeomorphically onto a single circle, but these are less interesting, and we will generally focus our attention on connected covering spaces.

The covering spaces of $S^1 \vee S^1$ form a remarkably rich family illustrating most of the general theory very concretely, so let us look at a small sample of these covering spaces to get a glimpse of what is going on.

To abbreviate notation, set $X = S^1 \vee S^1$. We view this as a graph with one vertex and two edges. If we label the edges a and b and choose orientations for a and b, then near the vertex we see the ends of four edges: an a edge entering the vertex, an a edge leaving the vertex, a b edge entering the vertex, and a b edge leaving the vertex. Suppose $\widetilde{X}$ is another graph with four edges meeting at each vertex, and suppose the edges of $\widetilde{X}$ are oriented and labeled a and b in such a way that the local picture near each vertex is the same as in X, i.e., an a-edge oriented toward the vertex, an a-edge oriented away from the vertex, a b-edge oriented toward the vertex, and a b-edge oriented away from the vertex. To give a name to this structure, let us call $\widetilde{X}$ a *2-oriented* graph. The table on the next page shows a number of examples. Given a 2-oriented graph $\widetilde{X}$ we can construct a map $p\colon \widetilde{X} \to X$ sending all vertices of $\widetilde{X}$ to the vertex of X and sending all edges homeomorphically to edges, preserving orientations and labeling. It is clear that the covering space condition is satisfied for p. The converse is also true: Every covering space of X is a graph which inherits a 2-orientation from X.

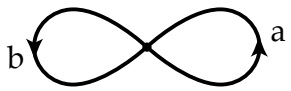

As the reader will discover by experimentation, it seems that every graph having four edges incident at each vertex can be 2-oriented. This can be proved for finite graphs as follows. A very classical and easily shown fact is that every finite graph with an even number of edges incident at each vertex has an Eulerian circuit: a loop traversing each edge exactly once. If there are four edges at each vertex, then labelling the edges of an Eulerian circuit alternately a and b produces a labelling with two a and two b edges at each vertex. The union of the a edges is then a collection of disjoint circles, as is the union of the b edges, and choosing orientations for all these circles gives a 2-orientation. It is a theorem in graph theory that infinite graphs with four edges incident at each vertex can also be 2-oriented; see [Koenig, chapter 13] for a proof. There is also a generalization to n-oriented graphs, which are covering spaces of the wedge sum of n circles.

(1)

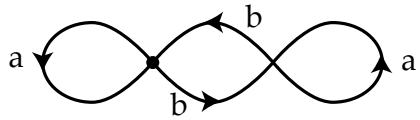

$\langle\, a,\ b^2,\ bab^{-1} \,\rangle$

(2)

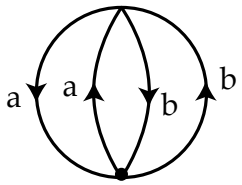

$\langle\, a^2,\ b^2,\ ab \,\rangle$

(3)

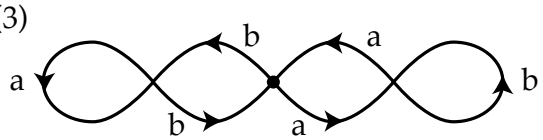

$\langle\, a^2,\ b^2,\ aba^{-1},\ bab^{-1} \,\rangle$

(4)

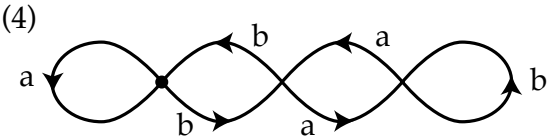

$\langle\, a,\ b^2,\ ba^2b^{-1},\ baba^{-1}b^{-1} \,\rangle$

(5)

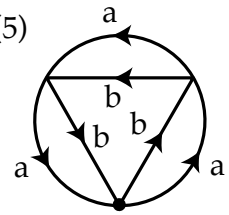

$\langle\, a^3,\ b^3,\ ab^{-1},\ b^{-1}a \,\rangle$

(6)

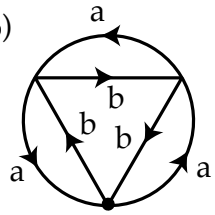

$\langle\, a^3,\ b^3,\ ab,\ ba \,\rangle$

(7)

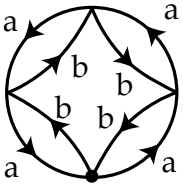

$\langle\, a^4,\ b^4,\ ab,\ ba,\ a^2b^2 \,\rangle$

(8)

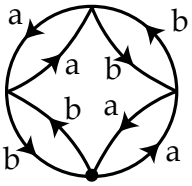

$\langle\, a^2,\ b^2,\ (ab)^2,\ (ba)^2,\ ab^2a \,\rangle$

(9)

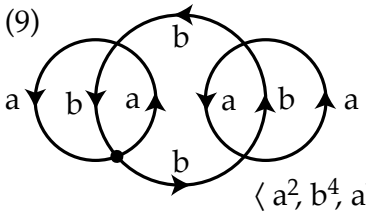

$\langle\, a^2,\ b^4,\ ab,\ ba^2b^{-1},\ bab^{-2} \,\rangle$

(10)

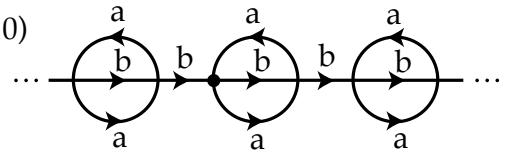

$\langle\, b^{2n}ab^{-2n-1},\ b^{2n+1}ab^{-2n} \mid n \in \mathbb{Z} \,\rangle$

(11)

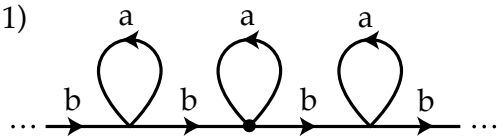

$\langle\, b^nab^{-n} \mid n \in \mathbb{Z} \,\rangle$

(12)

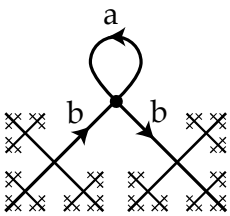

$\langle\, a \,\rangle$

(13)

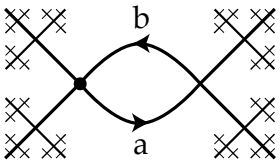

$\langle\, ab \,\rangle$

(14)

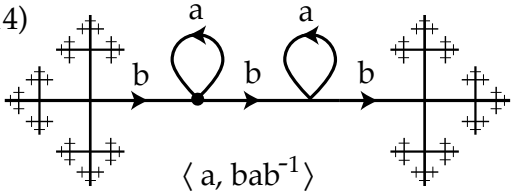

$\langle\, a,\ bab^{-1} \,\rangle$

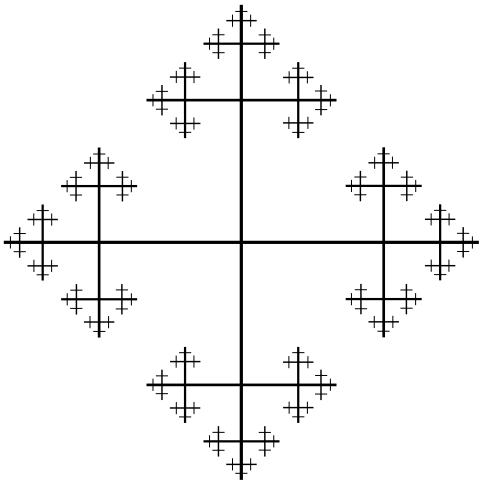

A portion of a rather special 2-oriented graph is shown in the figure to the right. This graph can be constructed concretely as a subspace of $\mathbb{R}^2$ by the following procedure. Start with the open intervals $(-1,1)$ in the x and y axes. Next, for a fixed number λ, $0 < \lambda < 1/2$, for example $\lambda = \frac{1}{3}$, adjoin four open segments of length 2λ, at distance λ from the ends of the previous segments and perpendicular to them, the new shorter segments being bisected by the older ones. For the third stage, add perpendicular segments of length $2\lambda^2$ at distance λ^2 from the endpoints of all the previous segments and bisected by them. The process is now repeated indefinitely, at the n^{th} stage adding segments of length $2\lambda^{n-1}$ at distance λ^{n-1} from all the previous endpoints. The union of all these segments is a graph, with vertices the intersection points of horizontal and vertical segments, and edges the subsegments between adjacent vertices. We can label all the horizontal edges a, oriented to the right, and the vertical edges b, oriented upward.

The special feature of this covering space is that it is simply-connected. It is called the *universal cover* of X because, as our general theory will show, it is a covering space of every other connected covering space of X.

The covering spaces (1)-(14) in the table are all nonsimply-connected. Their fundamental groups are free, namely, free products of $\mathbb{Z}$'s with generators the loops given by the listed words in a and b, starting at the basepoint $\tilde{x}_0$ indicated by the heavily shaded vertex; this can be proved in each case by applying van Kampen's theorem. One can also interpret the list of words as generators of the image subgroup $p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ in $\pi_1(X, x_0) = \langle a, b \rangle$. A general fact we shall prove about covering spaces is that the induced map $p_* : \pi_1(\widetilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$ is always injective. Thus we have the at-first-glance-paradoxical fact that the free group on two generators can contain as a subgroup a free group on any finite number of generators, or even on a countably infinite set of generators as in examples (10) and (11).

Changing the basepoint vertex changes the subgroup $p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ to a conjugate subgroup in $\pi_1(X, x_0)$, in fact conjugate by the element of $\pi_1(X, x_0)$ defined by a path joining one basepoint to the other. For example, the covering spaces (3) and (4) differ only in the choice of basepoints, and the corresponding subgroups of $\mathbb{Z} * \mathbb{Z}$ differ by conjugation by b.

The main classification theorem for covering spaces says that by associating the sub-

group $p_*(\pi_1(\widetilde{X},\widetilde{x}_0))$ to the covering space $p:\widetilde{X}\to X$ we obtain a one-to-one correspondence between all the different connected covering spaces of X and the conjugacy classes of subgroups of $\pi_1(X)$. If one keeps track of the basepoint vertex $\widetilde{x}_0\in\widetilde{X}$, then this is a one-to-one correspondence between covering spaces $p:(\widetilde{X},\widetilde{x}_0)\to(X,x_0)$ and actual subgroups of $\pi_1(X,x_0)$, not just conjugacy classes. Of course, for these statements to make sense one has to have a precise notion of when two covering spaces are the same, or 'isomorphic.' In the case at hand, two covering spaces of X are isomorphic if there is a homeomorphism between them which preserves the labeling and orientations of edges. Thus the covering spaces in (3) and (4) are isomorphic, but not by an isomorphism preserving basepoints, so the two corresponding subgroups of $\mathbb{Z}*\mathbb{Z}$ are distinct but conjugate. On the other hand, the two covering spaces in (5) and (6) are not isomorphic, though the graphs are homeomorphic, so the two corresponding subgroups of index 3 are isomorphic but not conjugate.

Another feature which might be observed is that some of the covering spaces are more symmetric than others, where by a 'symmetry' we mean a homeomorphism from the graph to itself which preserves the labelings and orientations. Specifically, one can ask in each case whether there are symmetries taking any one vertex onto any other. The examples (1), (2), (5)-(8), and (11) are the ones with this symmetry property. We shall see that having this symmetry property is equivalent to the corresponding subgroup of $\mathbb{Z}*\mathbb{Z}$ being a normal subgroup. And if the covering space does have this symmetry property, then its symmetries form a group isomorphic to the quotient group of $\mathbb{Z}*\mathbb{Z}$ by the normal subgroup. Since every group generated by two elements is a quotient group of $\mathbb{Z}*\mathbb{Z}$, this implies that every two-generator group is the symmetry group of some covering space of X.

Lifting Properties

Covering spaces are defined in fairly geometric terms, as maps $p:\widetilde{X}\to X$ which are local homeomorphisms in a rather strong sense. But from the viewpoint of algebraic topology, the distinctive feature of covering spaces is their behavior with respect to lifting of maps. Recall the terminology from §1.1: a **lift** of a map $f:Y\to X$ is a map $\widetilde{f}:Y\to\widetilde{X}$ such that $p\widetilde{f}=f$. We will describe three special lifting properties of covering spaces, and derive a few applications of these.

First we have the **homotopy lifting property**, sometimes called the **covering homotopy property**:

Proposition 1.16. *Given a covering space $p:\widetilde{X}\to X$, a homotopy $f_t:Y\to X$, and a*

map $\tilde f_0 : Y \to \tilde X$ *lifting* f_0, *then there exists a unique homotopy* $\tilde f_t : Y \to \tilde X$ *of* $\tilde f_0$ *which lifts* f_t.

Proof: For the covering space $p : \mathbb{R} \to S^1$ this is property (c) in Theorem 1.8, and the proof there was written so that it applies to any covering space. □

Taking Y to be a point gives the following **path lifting property** for a covering space $p : \tilde X \to X$: For each path $f : I \to X$ starting at a point $x_0 \in X$ and each $\tilde x_0 \in p^{-1}(x_0)$ there is a unique path $\tilde f : I \to \tilde X$ lifting f starting at $\tilde x_0$.

Here is a simple application:

Proposition 1.17. *The map* $p_* : \pi_1(\tilde X, \tilde x_0) \to \pi_1(X, x_0)$ *induced by a covering space* $p : (\tilde X, \tilde x_0) \to (X, x_0)$ *is injective. The image subgroup* $p_*(\pi_1(\tilde X, \tilde x_0))$ *in* $\pi_1(X, x_0)$ *consists of the homotopy classes of loops in* X *based at* x_0 *whose lifts to* $\tilde X$ *starting at* $\tilde x_0$ *are loops.*

Proof: An element of the kernel of p_* is represented by a loop $\tilde f_0 : I \to \tilde X$ with a homotopy $f_t : I \to X$ of $f_0 = p\tilde f_0$ to the trivial loop f_1. By the homotopy lifting property there is a lifted homotopy $\tilde f_t$ of $\tilde f_0$. The paths $\tilde f_t(0)$ and $\tilde f_t(1)$ traced out by the endpoints during this homotopy lie in the discrete set $p^{-1}(x_0)$, hence must be constant, so $\tilde f_t$ is a homotopy of loops at $\tilde x_0$. By the same reasoning, $\tilde f_1$ is the constant loop since f_1 is the constant loop. Hence $[\tilde f_0] = [\tilde f_1] = 0$ in $\pi_1(\tilde X, \tilde x_0)$ and p_* is injective. The other statement is clear from the definitions. □

If $p : \tilde X \to X$ is a covering space, then the cardinality of the set $p^{-1}(x)$ is locally constant over X, hence if X is connected, this cardinality is constant as x ranges over all of X. It is called the number of **sheets** of the covering.

Proposition 1.18. *The number of sheets of a covering space* $p : (\tilde X, \tilde x_0) \to (X, x_0)$ *with* X *and* $\tilde X$ *path-connected equals the index of* $p_*(\pi_1(\tilde X, \tilde x_0))$ *in* $\pi_1(X, x_0)$.

Proof: For a loop g in X based at x_0, let $\tilde g$ be its lift to $\tilde X$ starting at $\tilde x_0$. A product $h \cdot g$ with $[h] \in H = p_*(\pi_1(\tilde X, \tilde x_0))$ has the lift $\tilde h \cdot \tilde g$ ending at the same point as $\tilde g$ since $\tilde h$ is a loop. Thus we may define a function Φ from cosets $H[g]$ to $p^{-1}(x_0)$ by sending $[g]$ to $\tilde g(1)$. The path-connectedness of $\tilde X$ implies that Φ is surjective since $\tilde x_0$ can be joined to any point in $p^{-1}(x_0)$ by a path $\tilde g$ projecting to a loop g at x_0. To see that Φ is injective, observe that $\Phi([g_1]) = \Phi([g_2])$ implies that $g_1\overline{g}_2$ lifts to a loop in $\tilde X$ based at $\tilde x_0$, so $[g_1][g_2]^{-1}$ lies in H and $H[g_1] = H[g_2]$. □

It is useful to know about the existence and uniqueness of lifts not just of homotopies, but of general maps. For the existence question, there is the following **lifting criterion**:

Proposition 1.19. *Let $p:(\widetilde{X},\widetilde{x}_0)\to(X,x_0)$ be a covering space and $f:(Y,y_0)\to(X,x_0)$ a map. Assume Y is path-connected and locally path-connected. Then a lift $\widetilde{f}:(Y,y_0)\to(\widetilde{X},\widetilde{x}_0)$ of f exists iff $f_*(\pi_1(Y,y_0))\subset p_*(\pi_1(\widetilde{X},\widetilde{x}_0))$.*

When we say a space has a certain property locally, such as being locally path-connected, we shall mean that each point has arbitrarily small open neighborhoods with this property. Thus for Y to be locally path-connected means that for each point $y\in Y$ and each neighborhood U of y, there is an open neighborhood $V\subset U$ of y which is path-connected. Some authors weaken the requirement that V be path-connected to the condition that any two points in V be joinable by a path in U. This broader definition would work just as well for our purposes, necessitating only small adjustments in the proofs, but for simplicity we shall use the more restrictive definition.

Proof: The 'only if' statement is obvious since $f_*=p_*\widetilde{f}_*$. For the converse, let $y\in Y$ and let γ be a path in Y from y_0 to y. The path $f\gamma$ in X starting at x_0 has a unique lift $\widetilde{f\gamma}$ starting at $\widetilde{x}_0$. Define $\widetilde{f}(y)=\widetilde{f\gamma}(1)$. To show this is well-defined, independent of the choice of γ, let γ' be another path from y_0 to y. Then $(f\gamma')\cdot(\overline{f\gamma})$ is a closed loop h_0 at x_0 with $[h_0]\in f_*\pi_1(Y,y_0)\subset p_*\pi_1(\widetilde{X},\widetilde{x}_0)$. This means there is a homotopy h_t of h_0 to a loop h_1 which lifts to a loop $\widetilde{h}_1$ in $\widetilde{X}$ based at $\widetilde{x}_0$. Apply the covering homotopy property to h_t, to get a lifting $\widetilde{h}_t$. The paths $\widetilde{h}_t(0)$ and $\widetilde{h}_t(1)$ traced by the endpoints lie in the discrete set $p^{-1}(x_0)$, so must be constant. Therefore $\widetilde{h}_0$ is a closed loop at $\widetilde{x}_0$ lifting h_0. By the uniqueness of lifted paths, the first half of $\widetilde{h}_0$ is $\widetilde{f\gamma}'$ and the second half is $\widetilde{f\gamma}$ traversed backwards, with the common midpoint $\widetilde{f\gamma}(1)=\widetilde{f\gamma}'(1)=\widetilde{f}(y)$. This shows $\widetilde{f}$ is well-defined.

To see that $\widetilde{f}$ is continuous, let $U\subset X$ be an open neighborhood of $f(y)$ having a lift $\widetilde{U}\subset\widetilde{X}$ containing $\widetilde{f}(y)$ such that $p:\widetilde{U}\to U$ is a homeomorphism. Choose a path-connected open neighborhood V of y with $f(V)\subset U$. Then $\widetilde{f}(V)\subset\widetilde{U}$ and $\widetilde{f}|V=p^{-1}f$ since for paths from y_0 to $y'\in V$ we can take a fixed path γ from y_0 to y followed by paths η in V from y to y', and then the paths $(f\gamma)\cdot(f\eta)$ in X have lifts $(\widetilde{f\gamma})\cdot(\widetilde{f\eta})$ where $\widetilde{f\eta}=p^{-1}f\eta$. Since $\widetilde{f}|V=p^{-1}f$, it follows that $\widetilde{f}$ is continuous at y. □

An example showing the necessity of the local path-connectedness assumption on Y is described in the exercises at the end of the section.

Next we have the **unique lifting property**:

Proposition 1.20. *Given a covering space $p:\widetilde{X}\to X$ and a map $f:Y\to X$ with two lifts $\widetilde{f}_1,\widetilde{f}_2:Y\to\widetilde{X}$ which agree at one point of Y, then if Y is connected, these two lifts must agree on all of Y.*

Proof: For a point $y\in Y$, let U_1 and U_2 be neighborhoods of $\widetilde{f}_1(y)$ and $\widetilde{f}_2(y)$, respectively, mapping homeomorphically to a neighborhood U of Y by p, as in the definition of a covering space. By continuity there is a neighborhood N of y mapped into U_1 by $\widetilde{f}_1$ and into U_2 by $\widetilde{f}_2$. If $\widetilde{f}_1(y)\neq\widetilde{f}_2(y)$ then $U_1\neq U_2$, hence U_1 and U_2 are disjoint and $\widetilde{f}_1\neq\widetilde{f}_2$ throughout the neighborhood N. On the other hand, if $\widetilde{f}_1(y)=\widetilde{f}_2(y)$ then $U_1=U_2$ so $\widetilde{f}_1=\widetilde{f}_2$ on N since $p\widetilde{f}_1=p\widetilde{f}_2$ and p is injective on $U_1=U_2$. Thus the set of points where $\widetilde{f}_1$ and $\widetilde{f}_2$ agree is both open and closed in Y, hence if Y is connected this set is either empty or all of Y. □

These results can be applied to prove the two-dimensional case of the Borsuk-Ulam theorem:

Theorem 1.21. *If $f:S^n\to\mathbb{R}^n$ is a continuous map, then there is a point $x\in S^n$ with $f(x)=f(-x)$.*

It may be that there is only one pair of antipodal points $x,-x$ in S^n with the same image in $\mathbb{R}^n$, for example if f is orthogonal projection of the unit sphere $S^n\subset\mathbb{R}^{n+1}$ onto a hyperplane $\mathbb{R}^n$.

The proof for $n=1$ is easy since the difference $f(x)-f(-x)$ changes sign as x goes halfway around the circle, hence must be zero for some x. For $n\geq 2$ the theorem is certainly less obvious. Is it apparent, for example, that at every instant there must be a pair of antipodal points on the surface of the earth where the temperature is the same and also the barometric pressure?

Proof for $\mathbf{n=2}$: If the conclusion is false for $f:S^n\to\mathbb{R}^n$, we can define $g:S^n\to S^{n-1}$ by

$$g(x)=\frac{f(x)-f(-x)}{|f(x)-f(-x)|}$$

Then $g(-x)=-g(x)$, so g induces a quotient map $\overline{g}:\mathbb{R}\mathrm{P}^n\to\mathbb{R}\mathrm{P}^{n-1}$ sending each pair of antipodal points $\pm x$ to the pair of antipodal points $\pm g(x)$. The map $\overline{g}$ fits into a diagram

$$\begin{array}{ccc} S^n & \xrightarrow{\ g\ } & S^{n-1} \\ \big\downarrow{\scriptstyle p_n} & & \big\downarrow{\scriptstyle p_{n-1}} \\ \mathbb{R}\mathrm{P}^n & \xrightarrow{\ \overline{g}\ } & \mathbb{R}\mathrm{P}^{n-1} \end{array}$$

The compositions $p_{n-1}g$ and $\overline{g}p_n$ are equal since they both send $x \in S^n$ to the pair $\pm g(x)$. We may assume all four maps in the diagram are basepoint-preserving by choosing a basepoint in S^n and then taking its images in the other three spaces. When $n = 2$ the induced homomorphism $\overline{g}_* : \pi_1(\mathbb{R}\mathrm{P}^n) \to \pi_1(\mathbb{R}\mathrm{P}^{n-1})$ is necessarily trivial, being a homomorphism $\mathbb{Z}_2 \to \mathbb{Z}$. The lifting criterion then implies that a basepoint-preserving lift $\widetilde{g} : \mathbb{R}\mathrm{P}^n \to S^{n-1}$ of $\overline{g}$ exists. The maps g and $\widetilde{g}p_n$ are both lifts of $\overline{g}p_n$, and they are equal on the basepoint, so $g = \widetilde{g}p_n$ by the unique lifting property. Hence $g(-x) = \widetilde{g}p_n(-x) = \widetilde{g}p_n(x) = g(x)$, but this contradicts $g(-x) = -g(x)$. □

When $n > 2$ a more subtle argument is needed to show that $\overline{g}_*$ is trivial. With the machinery of cohomology in Chapter 3 this becomes rather easy, so we leave it as one of the exercises for §3.2.

Application. Whenever S^2 is expressed as the union of three closed sets A_1, A_2, and A_3, then one of these sets must contain a pair of antipodal points x and $-x$. To show this, let $d_i : S^2 \to \mathbb{R}$ measure the distance from points in S^2 to A_i, $d_i(x) = \inf_{y \in A_i} |x - y|$. This is certainly a continuous function, so we may apply the Borsuk-Ulam theorem to the map $S^2 \to \mathbb{R}^2$, $x \mapsto \big(d_1(x), d_2(x)\big)$, obtaining a pair of antipodal points x and $-x$ simultaneously equidistant from A_1 and from A_2. If either of these distances is zero, then x and $-x$ lie in A_1 or in A_2 since these are closed sets. On the other hand, if both distances are positive, then x and $-x$ lie in neither A_1 nor A_2, so they must lie in A_3.

The number 'three' in this result is best possible. For consider a sphere inscribed in a tetrahedron. Projecting the four faces of the tetrahedron radially onto the sphere, we obtain a cover of S^2 by four closed sets, none of which contains a pair of antipodal points.

The same arguments show that S^n cannot be covered by $n+1$ closed sets without antipodal pairs of points, though it can be covered by $n+2$ such sets. Even the case $n = 1$ is somewhat interesting: If the circle is covered by two closed sets, one of them must contain an antipodal pair of points.

The Classification of Covering Spaces

The idea in classifying all the different covering spaces of a given space X is to study the function which assigns to each covering space $p : (\widetilde{X}, \widetilde{x}_0) \to (X, x_0)$ the subgroup $p_*(\pi_1(\widetilde{X}, \widetilde{x}_0))$ of $\pi_1(X, x_0)$.

First we consider whether this function is surjective, i.e., whether every subgroup of $\pi_1(X, x_0)$ is realized as $p_*(\pi_1(\widetilde{X}, \widetilde{x}_0))$ for some covering space $p : (\widetilde{X}, \widetilde{x}_0) \to (X, x_0)$. A special case is whether the trivial subgroup is realized. Since p_* is always injective, this

amounts to asking whether X has a simply-connnected covering space. Answering this will take some work.

A necessary condition for X to have a simply-connected covering space is the following: Each point $x \in X$ has a neighborhood U such that the inclusion-induced map $\pi_1(U,x) \to \pi_1(X,x)$ is trivial; one says X is **semilocally simply-connected** if this holds. To see the necessity of this condition, suppose $p\colon \widetilde{X} \to X$ is a covering space with $\widetilde{X}$ simply-connected. Every point $x \in X$ has a neighborhood U having a lift $\widetilde{U} \subset \widetilde{X}$ projecting homeomorphically to U by p. Each loop in U lifts to a loop in $\widetilde{U}$, and the lifted loop is nullhomotopic in $\widetilde{X}$ since $\pi_1(\widetilde{X}) = 0$. So, composing this nullhomotopy with p, the original loop in U is nullhomotopic in X.

A locally simply-connected space is certainly semilocally simply-connected. For example, cell complexes have the much stronger property of being locally contractible, as we show in the Appendix. An example of a space which is not semilocally simply-connected is the subspace $X \subset \mathbb{R}^2$ consisting of the circles of radius $1/n$ centered at the point $(1/n, 0)$ for $n = 1, 2, \cdots$. For this space X, the cone $CX = (X \times I)/(X \times \{0\})$, is contractible and hence semilocally simply-connected, though it is not locally simply-connected.

We shall now show how to construct a simply-connected covering space of X if X is path-connected, locally path-connected, and semilocally simply-connected. To motivate the construction, suppose $p\colon (\widetilde{X}, \widetilde{x}_0) \to (X, x_0)$ is a simply-connected covering space. Each point $\widetilde{x} \in \widetilde{X}$ can then be joined to $\widetilde{x}_0$ by a unique homotopy class of paths, so we can view points of $\widetilde{X}$ as homotopy classes of paths starting at $\widetilde{x}_0$. The advantage of this is that, by the homotopy lifting property, homotopy classes of paths in $\widetilde{X}$ starting at $\widetilde{x}_0$ are the same as homotopy classes of paths in X starting at x_0. This gives a way of describing $\widetilde{X}$ purely in terms of X.

Thus, given a path-connected, locally path-connected, semilocally simply-connected space X with a basepoint $x_0 \in X$, we are led to define

$$\widetilde{X} = \{\, [\gamma] \mid \gamma \text{ is a path in } X \text{ starting at } x_0 \,\}$$

where, as usual, $[\gamma]$ denotes the homotopy class of γ with respect to homotopies which fix the endpoints $\gamma(0)$ and $\gamma(1)$. The function $p\colon \widetilde{X} \to X$ sending $[\gamma]$ to $\gamma(1)$ is then well-defined. Since X is path-connected, the endpoint $\gamma(1)$ can be any point of X, so p is surjective.

To define a topology on $\widetilde{X}$, note first that since X is locally path-connected and semilocally simply-connected, the path-connected open sets $U \subset X$ such that $\pi_1(U) \to \pi_1(X)$ is trivial form a basis for the topology of X. Given such a basic set U and a path

γ in X from x_0 to a point in U, let

$$U_{[\gamma]} = \{\, [\gamma \cdot \eta] \mid \eta \text{ is a path in } U \text{ with } \eta(0) = \gamma(1) \,\}.$$

As the notation indicates, $U_{[\gamma]}$ depends only on the homotopy class $[\gamma]$. Observe that $p : U_{[\gamma]} \to U$ is surjective since U is path-connected and injective since different choices of η joining $\gamma(1)$ to a fixed $x \in U$ are all homotopic in X, the map $\pi_1(U) \to \pi_1(X)$ being trivial. Further properties are:

(1) $U_{[\gamma]} = U_{[\gamma \cdot \eta]}$ if η is a path in U starting at $\gamma(1)$. The inclusion $U_{[\gamma]} \subset U_{[\gamma \cdot \eta]}$ holds since $\gamma \cdot \mu \simeq (\gamma \cdot \eta) \cdot (\overline{\eta} \cdot \mu)$ for paths μ in U, and the reverse inclusion holds since $(\gamma \cdot \eta) \cdot \mu \simeq \gamma \cdot (\eta \cdot \mu)$.

(2) If $U_{[\gamma]} \cap U_{[\gamma']} \neq \varnothing$, then $U_{[\gamma]} = U_{[\gamma']}$. This is because a point in $U_{[\gamma]} \cap U_{[\gamma']}$ has the form $[\gamma \cdot \eta] = [\gamma' \cdot \eta']$, and the points $[\gamma \cdot \eta \cdot \mu] = [\gamma' \cdot \eta' \cdot \mu]$ for varying paths μ in U fill up both $U_{[\gamma]}$ and $U_{[\gamma']}$, by (1).

(3) If for basic sets U and V an intersection $U_{[\gamma]} \cap V_{[\delta]}$ is nonempty, containing a point $[\varepsilon] = [\gamma \cdot \eta] = [\delta \cdot \mu]$, then for any basic set $W \subset U \cap V$ containing $\varepsilon(1)$ we have $W_{[\varepsilon]} \subset U_{[\gamma]} \cap V_{[\delta]}$. This follows as in (2) since $W_{[\varepsilon]}$ consists of points $[\varepsilon \cdot \omega] = [\gamma \cdot \eta \cdot \omega] = [\delta \cdot \mu \cdot \omega]$ with ω a path in W.

From (3) we deduce that the sets $U_{[\gamma]}$ form a basis for a topology on $\widetilde{X}$. Each projection $p : U_{[\gamma]} \to U$ is a homeomorphism since p gives a one-to-one correspondence between basic sets $V_{[\delta]} \subset U_{[\gamma]}$ and basic sets $V \subset U$. Property (2) then implies that $p : \widetilde{X} \to X$ is a covering space.

A natural basepoint for $\widetilde{X}$ is $[x_0]$, the homotopy class of the constant path at x_0. Any $[\gamma] \in \widetilde{X}$ can be joined to $[x_0]$ by a path in $\widetilde{X}$ by restricting γ to progressively shorter segments $[0, t] \subset [0, 1]$, so $\widetilde{X}$ is path-connected. To see that $\pi_1(\widetilde{X}) = 0$, let $f : I \to \widetilde{X}$ be a loop based at $[x_0]$. The composition pf is then a loop γ in X based at x_0. Let γ_t be the path in X obtained by restricting the loop γ to $[0, t]$. Then $[\gamma_t]$ for t varying from 0 to 1 forms a path in $\widetilde{X}$ lifting the loop γ. This lift $[\gamma_t]$ starts at $[\gamma_0] = [x_0] = f(0)$, so by the unique lifting property of the cover space $\widetilde{X} \to X$ we must have $[\gamma_t] = f(t)$ for all t. In particular, $[\gamma_1] = f(1) = [x_0]$. Since $\gamma_1 = \gamma$, this says the loop $\gamma = pf$ is nullhomotopic. Thus $p_*([f]) = 0$ in $\pi_1(X)$. Since p_* is injective, $[f] = 0$ and hence $\pi_1(\widetilde{X}) = 0$. This completes the construction.

The hypotheses for constructing a simply-connected covering space of X in fact suffice for constructing covering spaces realizing arbitrary subgroups of $\pi_1(X)$:

Proposition 1.22. *Suppose X is path-connected, locally path-connected, and semilocally*

simply-connected. Then for every subgroup $H \subset \pi_1(X, x_0)$ *there is a covering space* $p : X_H \to X$ *such that* $p_*(\pi_1(X_H, \tilde{x}_0)) = H$ *for a suitably chosen basepoint* $\tilde{x}_0 \in X_H$.

Proof: For points $[\gamma], [\gamma']$ in the simply-connected covering space $\tilde{X}$ constructed above, define $[\gamma] \sim [\gamma']$ to mean $\gamma(1) = \gamma'(1)$ and $[\gamma\overline{\gamma'}] \in H$. It is easy to see that this is an equivalence relation since H is a subgroup; namely, it is reflexive since H contains the identity element, symmetric since H is closed under inverses, and transitive since H is closed under multiplication. Let X_H be the quotient space of $\tilde{X}$ obtained by identifying $[\gamma]$ with $[\gamma']$ if $[\gamma] \sim [\gamma']$. Note that if $\gamma(1) = \gamma'(1)$, then $[\gamma] \sim [\gamma']$ iff $[\gamma[\eta] \sim [\gamma'\eta]$. This means that if any two points in basic neighborhoods $U_{[\gamma]}$ and $U_{[\gamma'']}$ are identified in X_H then the whole neighborhoods are identified. Hence the natural projection $X_H \to X$ induced by $[\gamma] \mapsto \gamma(1)$ is a covering space.

If we choose for the basepoint $\tilde{x}_0 \in X_H$ the equivalence class of the constant path c at x_0, then the image of $p_* : \pi_1(X_H, \tilde{x}_0) \to \pi_1(X, x_0)$ is exactly H. This is because for a loop γ in X based at x_0, its lift to $\tilde{X}$ starting at $[c]$ ends at $[\gamma]$, so the image of this loop in X_H is a closed loop iff $[\gamma] \sim [c]$, i.e., $[\gamma] \in H$. □

Having taken care of the existence of covering spaces of X corresponding to all subgroups of $\pi_1(X)$, we turn now to uniqueness up to isomorphism, where an **isomorphism** between covering spaces $p_1 : \tilde{X}_1 \to X$ and $p_2 : \tilde{X}_2 \to X$ is a homeomorphism $f : \tilde{X}_1 \to \tilde{X}_2$ such that $p_1 = p_2 f$. This condition means exactly that f preserves the covering space structures, taking $p_1^{-1}(x)$ to $p_2^{-1}(x)$ for each $x \in X$. The inverse f^{-1} is then also an isomorphism, and the composition of two isomorphisms is an isomorphism, so we have an equivalence relation.

Proposition 1.23. *If* X *is path-connected and locally path-connected, then two path-connected covering spaces* $p_1 : \tilde{X}_1 \to X$ *and* $p_2 : \tilde{X}_2 \to X$ *are isomorphic via an isomorphism* $f : \tilde{X}_1 \to \tilde{X}_2$ *taking a basepoint* $\tilde{x}_1 \in p_1^{-1}(x_0)$ *to a basepoint* $\tilde{x}_2 \in p_2^{-1}(x_0)$ *iff* $p_{1*}(\pi_1(\tilde{X}_1, \tilde{x}_1)) = p_{2*}(\pi_1(\tilde{X}_2, \tilde{x}_2))$.

Proof: If there is an isomorphism $f : (\tilde{X}_1, \tilde{x}_1) \to (\tilde{X}_2, \tilde{x}_2)$, then the relations $p_1 = p_2 f$ and $p_2 = p_1 f^{-1}$ imply that $p_{1*}(\pi_1(\tilde{X}_1, \tilde{x}_1)) = p_{2*}(\pi_1(\tilde{X}_2, \tilde{x}_2))$. Conversely, suppose that $p_{1*}(\pi_1(\tilde{X}_1, \tilde{x}_1)) = p_{2*}(\pi_1(\tilde{X}_2, \tilde{x}_2))$. By the lifting criterion, we may lift p_1 to a map $\tilde{p}_1 : (\tilde{X}_1, \tilde{x}_1) \to (\tilde{X}_2, \tilde{x}_2)$ with $p_2 \tilde{p}_1 = p_1$. Symmetrically, we obtain $\tilde{p}_2 : (\tilde{X}_2, \tilde{x}_2) \to (\tilde{X}_1, \tilde{x}_1)$ with $p_1 \tilde{p}_2 = p_2$. Then by the unique lifting property, $\tilde{p}_1 \tilde{p}_2 = \mathbb{1}$ and $\tilde{p}_2 \tilde{p}_1 = \mathbb{1}$ since these composed lifts fix the basepoints. Thus $\tilde{p}_1$ and $\tilde{p}_2$ are inverse isomorphisms. □

We have proved the first half of the following classification theorem:

Theorem 1.24. *If the space X is path-connected, locally path-connected, and semilocally simply-connected, then there is a bijection between the isomorphism classes of path-connected covering spaces $p:(\widetilde{X},\widetilde{x}_0)\to(X,x_0)$ and the subgroups of $\pi_1(X,x_0)$, obtained by associating the subgroup $p_*(\pi_1(\widetilde{X},\widetilde{x}_0))$ to the covering space $(\widetilde{X},\widetilde{x}_0)$. If basepoints are ignored, this correspondence gives a bijection between the set of isomorphism classes of path-connected covering spaces $p:\widetilde{X}\to X$ and conjugacy classes of subgroups of $\pi_1(X,x_0)$.*

Proof: It remains only to prove the last statement. For a covering space $p:(\widetilde{X},\widetilde{x}_0)\to(X,x_0)$, changing the basepoint $\widetilde{x}_0$ within $p^{-1}(x_0)$ changes $p_*(\pi_1(\widetilde{X},\widetilde{x}_0))$ to a conjugate subgroup of $\pi_1(X,x_0)$ since a new basepoint $\widetilde{x}_0'$ is joined to the old one $\widetilde{x}_0$ by a path $\widetilde{h}$ in $\widetilde{X}$ which projects to a loop h in X, so the change-of-basepoint isomorphism $\beta_{\widetilde{h}}$ corresponds, via the inclusions p_*, to β_h which is conjugation by $[h]$ in $\pi_1(X,x_0)$.

$$\begin{array}{ccc}
\pi_1(\widetilde{X},\widetilde{x}_0) & \xrightarrow[\beta_{\widetilde{h}}]{\approx} & \pi_1(\widetilde{X},\widetilde{x}_0') \\
\big\downarrow p_* & & \big\downarrow p_* \\
\pi_1(X,x_0) & \xrightarrow[\approx]{\beta_h} & \pi_1(X,x_0)
\end{array}$$

Conversely, every subgroup of $\pi_1(X,x_0)$ conjugate to $p_*(\pi_1(\widetilde{X},\widetilde{x}_0))$ arises from such a change of basepoint in $p^{-1}(x_0)$ since a loop h in X at x_0 lifts to a path $\widetilde{h}$ in $\widetilde{X}$ from $\widetilde{x}_0$ to a new basepoint in $p^{-1}(x_0)$. □

A consequence of the lifting criterion is that a simply-connected covering space of a path-connected, locally path-connected space X is a covering space of every other path-connected covering space of X. A simply-connected covering space of X is therefore called a **universal cover**. Such a covering space is unique up to isomorphism, so one is justified in saying *the* universal cover.

More generally, there is a partial ordering on the collection of covering spaces of X, according to which ones cover which others. This corresponds to the partial ordering by inclusion of the various subgroups of $\pi_1(X)$, or conjugacy classes of subgroups if basepoints are ignored.

Deck Transformations and Group Actions

For a covering space $p:\widetilde{X}\to X$ the isomorphisms $\widetilde{X}\to\widetilde{X}$ are called **deck transformations** or **covering transformations**. These form a group $G(\widetilde{X})$ under composition. For example, for the covering space $p:\mathbb{R}\to S^1$, $p(t)=(\cos 2\pi t,\sin 2\pi t)$, the deck transformations are the translations $x\mapsto x+n$ for $n\in\mathbb{Z}$, so $G(\widetilde{X})\approx\mathbb{Z}$ in this case. If we

think of this covering space as the projection of a vertical helix onto a circle, then the deck transformations are the vertical translations taking the helix onto itself.

By the unique lifting property, a deck transformation is completely determined by where it sends a single point, assuming $\widetilde{X}$ is path-connected. In particular only the identity deck transformation can fix a point of $\widetilde{X}$.

A **normal** covering space $p:\widetilde{X}\to X$ is one for which $G(\widetilde{X})$ acts **transitively** on $p^{-1}(x)$ for all $x\in X$, i.e., for any two lifts $\widetilde{x},\widetilde{x}'$ of x there is a deck transformation taking $\widetilde{x}$ to $\widetilde{x}'$. Intuitively, a normal covering space is one with 'maximal symmetry.'

Proposition 1.25. *Let $p:(\widetilde{X},\widetilde{x}_0)\to(X,x_0)$ be a path-connected covering space of the path-connected, locally path-connected space X, and let $H=p_*(\pi_1(\widetilde{X},\widetilde{x}_0))\subset\pi_1(X,x_0)$.*

(a) This covering space is normal iff H is a normal subgroup of $\pi_1(X,x_0)$.

(b) $G(\widetilde{X})$ is isomorphic to $N(H)/H$ where $N(H)$ is the normalizer of H in $\pi_1(X,x_0)$. In particular, $G(\widetilde{X})$ is isomorphic to $\pi_1(X,x_0)/H$ if $\widetilde{X}$ is a normal covering. Hence for the universal cover $\widetilde{X}\to X$ we have $G(\widetilde{X})\approx\pi_1(X)$.

Proof: We observed earlier in the proof of the classification theorem that changing the basepoint $\widetilde{x}_0\in p^{-1}(x_0)$ to $\widetilde{x}_1\in p^{-1}(x_0)$ corresponds precisely to conjugating H by a loop h which is the image of a path $\widetilde{h}$ from $\widetilde{x}_0$ to $\widetilde{x}_1$. Thus $[h]$ is in the normalizer $N(H)$ iff $p_*(\pi_1(\widetilde{X},\widetilde{x}_0))=p_*(\pi_1(\widetilde{X},\widetilde{x}_1))$, which by the lifting criterion is equivalent to the existence of a deck transformation taking $\widetilde{x}_0$ to $\widetilde{x}_1$. Hence the covering space is normal iff $N(H)=\pi_1(X,x_0)$, i.e., H is a normal subgroup of $\pi_1(X,x_0)$.

Define $\varphi:N(H)\to G(\widetilde{X})$ sending $[h]$ to the deck transformation τ taking $\widetilde{x}_0$ to $\widetilde{x}_1$, in the notation above. Then φ is a homomorphism: If h' is another loop corresponding to the deck transformation τ' taking $\widetilde{x}_0$ to $\widetilde{x}_1'$ then $h\cdot h'$ lifts to $\widetilde{h}\cdot(\tau(\widetilde{h}'))$, a path from $\widetilde{x}_0$ to $\tau(\widetilde{x}_1')=\tau\tau'(\widetilde{x}_0)$, so $\tau\tau'$ is the deck transformation corresponding to $[h][h']$. By the preceding paragraph φ is surjective. Its kernel consists of classes $[h]$ lifting to loops in $\widetilde{X}$, which are exactly the elements of $p_*(\pi_1(\widetilde{X},\widetilde{x}_0))=H$. □

The group of deck transformations is a special case of the general notion of 'groups acting on spaces.' Given a group G and a space Y, then an **action** of G on Y is a homomorphism ρ from G to the group $Homeo(Y)$ of homeomorphisms of Y. Thus to each $g\in G$ is associated a homeomorphism $\rho(g):Y\to Y$, which for notational simplicity we write simply as $g:Y\to Y$. For ρ to be a homomorphism amounts to requiring that $g_1(g_2(y))=(g_1g_2)(y)$ for all $g_1,g_2\in G$ and $y\in Y$. If ρ is injective then it identifies G with a subgroup of $Homeo(Y)$, and in practice not much is lost in assuming ρ is an inclusion $G\hookrightarrow Homeo(Y)$ since in any case the subgroup $\rho(G)\subset Homeo(Y)$ contains all the topological information about the action.

We shall be interested in actions satisfying the following condition:

(∗) Each point $y \in Y$ has a neighborhood U such that all the images $g(U)$ for varying $g \in G$ are disjoint, i.e., $g_1(U) \cap g_2(U) \neq \varnothing$ implies $g_1 = g_2$.

For deck transformations, one can see the action of $G(\widetilde{X})$ on $\widetilde{X}$ satisfies $(*)$ as follows. Suppose $\widetilde{U} \subset \widetilde{X}$ projects homeomorphically to $U \subset X$. If $g_1(\widetilde{U}) \cap g_2(\widetilde{U}) \neq \varnothing$ for $g_1, g_2 \in G(\widetilde{X})$, then we have $g_1(\widetilde{x}_1) = g_2(\widetilde{x}_2)$ for points $\widetilde{x}_1, \widetilde{x}_2 \in \widetilde{U}$. Since $\widetilde{x}_1$ and $\widetilde{x}_2$ must lie in the same set $p^{-1}(x)$, which intersects $\widetilde{U}$ in only one point, we must have $\widetilde{x}_1 = \widetilde{x}_2$, and $g_1^{-1}g_2$ fixes this point so $g_1^{-1}g_2 = \mathbb{1}$ and $g_1 = g_2$.

Note that in $(*)$ it suffices to take g_1 to be the identity since $g_1(U) \cap g_2(U) \neq \varnothing$ is equivalent to $U \cap g_1^{-1}g_2(U) \neq \varnothing$. Thus we have the equivalent condition that $U \cap g(U) \neq \varnothing$ only for g the identity.

Given an action of a group G on a space Y, we can form a space Y/G, the quotient space of Y in which each point y is identified with all its images $g(y)$ as g ranges over G. The points of Y/G are thus the **orbits** $Gy = \{\, g(y) \mid g \in G \,\} \subset Y$, and Y/G is called the **orbit space** of the action. For example, for a normal covering space $\widetilde{X} \to X$, the orbit space $\widetilde{X}/G(\widetilde{X})$ is just X.

Proposition 1.26. *If the action of G on Y satisfies the condition $(*)$, then the quotient map $p : Y \to Y/G$, $p(y) = Gy$, is a covering space. This is a normal covering space, and G is its group of deck transformations if Y is path-connected.*

Proof: Given an open set $U \subset Y$ as in condition $(*)$, the quotient map p simply identifies all the disjoint homeomorphic sets $\{\, g(U) \mid g \in G \,\}$ to a single open set $p(U)$ in Y/G. By the definition of the quotient topology on Y/G, p restricts to a homeomorphism $g(U) \approx p(U)$ for each $g \in G$ so we have a covering space. Each element of G acts as a deck transformation, and the action is transitive by the defiinition of Y/G, so the covering space is normal. G is contained in the group of deck transformations, and equals this group if Y is path-connected, since if f is any deck transformation, then for any point $y \in Y$, $p(y) = pf(y)$ so y and $f(y)$ are in the same orbit and there is a $g \in G$ with $g(y) = f(y)$, hence $f = g$ since deck transformations of a path-connected covering space are uniquely determined by where they send a point. □

In view of the preceding proposition, we shall call an action satisfying $(*)$ a **covering space action**. This is not standard terminology, but there does not seem to be a universally accepted name for actions satisfying $(*)$. Sometimes these are called 'properly discontinuous' actions, but more often this rather unattractive term means something

weaker: Every point $x \in X$ has a neighborhood U such that $U \cap g(U)$ is nonempty for only finitely many $g \in G$. Many symmetry groups have this proper discontinuity property without satisfying $(*)$, for example the group of symmetries of the familiar tiling of $\mathbb{R}^2$ by regular hexagons. The reason why the action of this group on $\mathbb{R}^2$ fails to satisfy $(*)$ is that there are **fixed points**: points y for which there is a nontrivial element $g \in G$ with $g(y) = y$. For example, the vertices of the hexagons are fixed by the 120 degree rotations about these points, and the midpoints of edges are fixed by 180 degree rotations. An action without fixed points is called a **free** action. Thus for a free action of G on Y, only the identity element of G fixes any point of Y. This is equivalent to requiring that all the images $g(y)$ of each $y \in Y$ are distinct, i.e., $g_1(y) = g_2(y)$ implies $g_1 = g_2$, since $g_1(y) = g_2(y)$ is equivalent to $g_1^{-1}g_2(y) = y$. Though condition $(*)$ implies freeness, the converse is not always true. An example is the action of $\mathbb{Z}$ on S^1 in which a generator of $\mathbb{Z}$ acts by rotation through an angle α which is an irrational multiple of 2π. In this case each orbit $\mathbb{Z}y$ is dense in S^1, so condition $(*)$ cannot hold since it implies that orbits are discrete subspaces. An exercise at the end of the section is to show that for actions on Hausdorff spaces, freeness plus proper discontinuity implies condition $(*)$. Note that proper discontinuity is automatic for actions by a finite group.

Here are two 2-dimensional examples.

Example 1. Take Y to be $\mathbb{R}^2$ and G the group of all translations $(x, y) \mapsto (x+m, y+n)$ for $m, n \in \mathbb{Z}$, so G is isomorphic to $\mathbb{Z} \times \mathbb{Z}$. The condition $(*)$ is satisfied by taking U to be an open disk of radius $1/2$ or less centered at a given point of $\mathbb{R}^2$. The quotient space Y/G is homeomorphic to the torus $S^1 \times S^1$. This can be seen by considering a unit square in $\mathbb{R}^2$ with edges parallel to the coordinate axes. This square is disjoint from all its translates by elements of G, except that the lower edge is translated to the upper edge and the left edge is translated to the right edge. This means that these opposite pairs of edges are identified in the quotient space, and there are no other identifications, so we obtain the familar description of the torus as a square with opposite edges glued together.

Example 2. Let Y be the closed orientable surface of genus 11, an '11-hole torus,' as in the figure below. This has a 5-fold rotational symmetry, generated by a rotation of angle $2\pi/5$. Thus we have the cyclic group $\mathbb{Z}_5$ acting on Y, and the condition $(*)$ is obviously satisfied. The quotient space $Y/\mathbb{Z}_5$ is a surface of genus 3, obtained from one of the five subsurfaces of Y cut off by the circles $C_1, \cdots, C_5$ by identifying its two boundary circles C_i and C_{i+1} to form the circle C as shown. Thus we have a covering space $M_{11} \to M_3$ where M_g denotes the closed orientable surface of genus g. This example obviously generalizes by replacing the two holes in each 'arm' of M_{11} by m holes, $m \geq 0$, and the 5-fold symmetry

by n-fold symmetry, $n \geq 2$. Thus there is a covering space $M_{mn+1} \to M_{m+1}$. An exercise in §2.2 is to show, by an Euler characteristic argument, that if there is a covering space $M_g \to M_h$ then $g = mn + 1$ and $h = m + 1$ for some m and n.

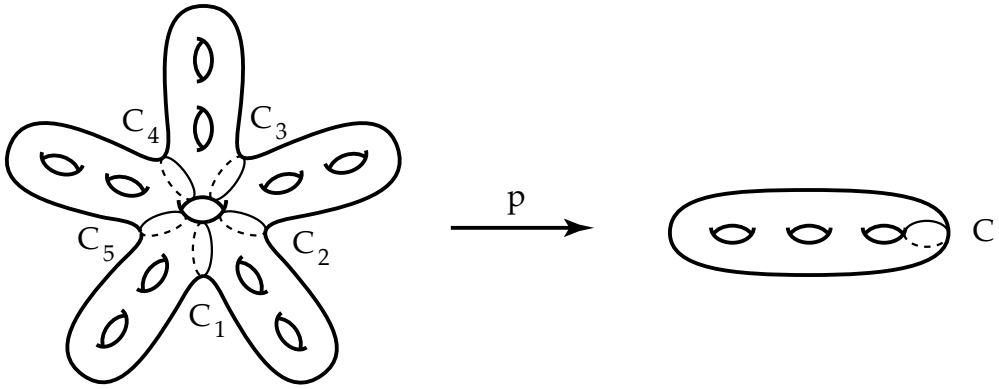

Proposition 1.27. *For a covering space action of a group G on a simply-connected space Y, $\pi_1(Y/G)$ is isomorphic to G.*

Proof: This follows from the preceding proposition and (b) of Proposition 1.25, at least when Y is locally path-connected. But it is not hard to give a direct proof without this additional hypothesis, along the lines of Theorem 1.8 which calculated $\pi_1(S^1)$. Let $p: Y \to Y/G$ be the quotient projection, a covering space by hypothesis. Choosing a basepoint $y_0 \in Y$, we can define a map $\psi: G \to \pi_1(Y/G, p(y_0))$, by sending $g \in G$ to the homotopy class of the loop pf where f is any path in Y from y_0 to $g(y_0)$. Since Y is simply-connected, any two choices of such an f are homotopic, and composing such a homotopy with p gives a homotopy between the loops pf in Y/G, so ψ is well-defined.

To see that ψ is a homomorphism, let f_1 and f_2 be paths in Y from y_0 to $g_1(y_0)$ and $g_2(y_0)$, for $g_1, g_2 \in G$. Then $f_1 \cdot (g_1 f_2)$ is a path from y_0 to $g_1 g_2(y_0)$, so $\psi(g_1 g_2) = [p(f_1 \cdot (g_1 f_2))] = [pf_1 \cdot pg_1 f_2] = [pf_1 \cdot pf_2] = \psi(g_1)\psi(g_2)$.

To show ψ is surjective, let $f: I \to Y/G$ be a loop at the basepoint $p(y_0)$ representing a given element of $\pi_1(Y/G)$. By the path lifting property there is a lift $\widetilde{f}$ starting at y_0. This path $\widetilde{f}$ ends at some point $g(y_0)$ since $p\widetilde{f}(1) = f(1) = p(y_0)$. So $\psi(g) = [f]$.

To show ψ is injective, suppose $\psi(g_1) = \psi(g_2)$ and let f_1 and f_2 be paths in Y from y_0 to $g_1(y_0)$ and $g_2(y_0)$. Since $\psi(g_1) = \psi(g_2)$, the loops pf_1 and pf_2 are homotopic. The homotopy lifting property then implies that f_1 and f_2 are homotopic, as paths with fixed endpoints, so the endpoints $g_1(y_0)$ and $g_2(y_0)$ are the same. Condition $(*)$ then implies $g_1 = g_2$. □

This can be used to calculate fundamental groups without using van Kampen's theo-

rem, as the following example illustrates.

Example: $\mathbb{R}\mathrm{P}^n$. The antipodal map of S^n, $x \mapsto -x$, generates an action of $\mathbb{Z}_2$ on S^n with orbit space $\mathbb{R}\mathrm{P}^n$, real projective n-space, as defined in Chapter 0. The action is a covering space action since each open hemisphere in S^n is disjoint from its antipodal image. By Corollary 1.6, S^n is simply-connected if $n \geq 2$, so from the covering space $S^n \to \mathbb{R}\mathrm{P}^n$ we deduce that $\pi_1(\mathbb{R}\mathrm{P}^n) \approx \mathbb{Z}_2$ for $n \geq 2$. For $n = 1$, a circle with antipodal points identified is again a circle, so $\mathbb{R}\mathrm{P}^1 \approx S^1$ and $\pi_1(\mathbb{R}\mathrm{P}^1) \approx \mathbb{Z}$. The natural inclusions $S^1 \subset S^2 \subset \cdots \subset S^n$ coming from $\mathbb{R}^2 \subset \mathbb{R}^3 \subset \cdots \subset \mathbb{R}^{n+1}$ induce inclusions $\mathbb{R}\mathrm{P}^1 \subset \mathbb{R}\mathrm{P}^2 \subset \cdots \subset \mathbb{R}\mathrm{P}^n$, and the circle $\mathbb{R}\mathrm{P}^1 \subset \mathbb{R}\mathrm{P}^n$ is a loop generating $\pi_1(\mathbb{R}\mathrm{P}^n)$ since it lifts to a path in $S^1 \subset S^n$ joining a pair of antipodal points. One can see explicitly that this loop $\mathbb{R}\mathrm{P}^1 \subset \mathbb{R}\mathrm{P}^n$ has order two in $\pi_1(\mathbb{R}\mathrm{P}^n)$ for $n \geq 2$ since the composition of this loop with itself lifts to the circle $S^1 \subset S^n$ which can be homotoped to the trivial loop since $\pi_1(S^n) = 0$, and the projection of this homotopy into $\mathbb{R}\mathrm{P}^n$ gives a nullhomotopy of twice the loop $\mathbb{R}\mathrm{P}^1$.

Which other finite groups act freely on S^n? We will show in §2.2 that $\mathbb{Z}_2$ is the only possibility when n is even, but for odd n the question is much more difficult. To begin with, it is easy to construct a free action of any cyclic group $\mathbb{Z}_m$ on S^{2k-1}, generated by the rotation $v \mapsto e^{2\pi i/m}v$ of the unit sphere S^{2k-1} in $\mathbb{C}^k = \mathbb{R}^{2k}$. The action is free since an equation $v = e^{2\pi i\ell/m}v$ with $0 < \ell < m$ implies $v = 0$, but 0 is not a point of S^{2k-1}. The orbit space $S^{2k-1}/\mathbb{Z}_m$ is one of a family of *lens spaces* studied in more detail in Example 7 of §2.2.

There are also noncyclic finite groups which act freely as rotations of S^n for odd $n > 1$. These actions are classified quite explicitly in [Wolf]. Examples in the simplest case $n = 3$ can be produced as follows.

View $\mathbb{R}^4$ as the quaternion algebra $\mathbb{H}$. Multiplication of quaternions satisfies $|ab| = |a||b|$ where $|a|$ denotes the usual Euclidean length of a vector $a \in \mathbb{R}^4$. Thus if a and b are unit vectors, so is ab, so quaternion multiplication defines a map $S^3 \times S^3 \to S^3$. This in fact makes S^3 into a group, though associativity is all we need now, since it implies that any subgroup G of S^3 acts on S^3 by left-multiplication, $g(x) = gx$. This action is free since an equation $x = gx$ in the division algebra $\mathbb{H}$ implies $g = 1$ or $x = 0$. As a concrete example, G could be the familiar quaternion group $Q_8 = \{\pm 1, \pm i, \pm j, \pm k\}$ from group theory. More generally, for a positive integer m, let Q_{4m} be the subgroup of S^3 generated by the two quaternions $a = e^{\pi i/m} \in \mathbb{C} \subset \mathbb{H}$ and $b = j$. Thus a has order $2m$ and b has order 4. The easily-verified relations $a^m = b^2 = -1$ and $bab^{-1} = a^{-1}$ imply that the subgroup $\mathbb{Z}_{2m}$ generated by a is normal and of index 2 in Q_{4m}. Hence Q_{4m} is a group of order $4m$, called the *generalized quaternion group*. Another common name for

this group is the *binary dihedral group* D^*_{4m} since its quotient by the subgroup $\{\pm 1\}$ is the ordinary dihedral group D_{2m} of order $2m$.

Besides the groups $Q_{4m} = D^*_{4m}$ there are just three other noncyclic finite subgroups of S^3: the binary tetrahedral, octahedral, and icosahedral groups T^*_{24}, O^*_{48}, and I^*_{120}, of orders indicated by the subscripts. These project two-to-one onto the groups of rotational symmetries of a regular tetrahedron, octahedron (or cube), and icosahedron (or dodecahedron). In fact, it is not hard to see that the homomorphism $S^3 \to SO(3)$ sending $u \in S^3 \subset \mathbb{H}$ to the isometry $v \to u^{-1}vu$ of $\mathbb{R}^3$, viewing $\mathbb{R}^3$ as the 'pure imaginary' quaternions $v = ai + bj + ck$, is surjective with kernel $\{\pm 1\}$. Then the groups D^*_{4m}, T^*_{24}, O^*_{48}, I^*_{120} are the preimages in S^3 of the groups of rotational symmetries of a regular polygon or polyhedron in $\mathbb{R}^3$.

There are two conditions which a finite group G acting freely on S^n must satisfy:

(a) Every abelian subgroup of G is cyclic. Equivalently, G contains no subgroup $\mathbb{Z}_p \times \mathbb{Z}_p$ with p prime.
(b) G contains at most one element of order 2.

Statement (a) is an exercise for §4.1, and for a proof of (b) the original source [Milnor57] is recommended reading. The groups satisfying (a) have been completely classified, as explained for example in [K.Brown, section VI.9]. An example of a group satisfying (a) but not (b) is the dihedral group D_{2m} for odd $m > 1$.

There is also a much more difficult converse: A finite group satisfying (a) and (b) acts freely on S^n for some n. See [Madsen-Thomas-Wall] or [Davis-Milgram]. There is also almost complete information about which n's are possible for a given group.

A Family of Examples: Cayley Complexes. There is a very classical method for viewing groups geometrically as graphs, which fits very nicely into covering space theory. Recall from §1.2 how we associated to each group presentation $G = \langle\, g_\alpha \mid r_\beta \,\rangle$ a 2-dimensional cell complex X_G with $\pi_1(X_G) \approx G$ by taking a wedge-sum of circles, one for each generator g_α, and then attaching a 2-cell for each relator r_β. A simply-connected cell complex $\widetilde{X}_G$ which is the universal cover of X_G can be constructed explicitly as follows. Let the vertices of $\widetilde{X}_G$ be the elements of G themselves. Then, at each vertex $g \in G$, insert an edge joining g to gg_α for each of the chosen generators g_α. The resulting graph is called the **Cayley graph** of G with respect to the generators g_α. This graph is connected since every element of G is a product of g_α's, so there is a path in the graph joining each vertex to the identity vertex e. Each relation r_β determines a loop in the graph, starting at any vertex g, and we attach a 2-cell for each such loop. The resulting cell complex $\widetilde{X}_G$ is the **Cayley complex** of G.

To see that $\widetilde{X}_G$ is simply-connected, represent a given element of $\pi_1(\widetilde{X}_G)$ by a loop of edges in $\widetilde{X}_G$, based at the identity vertex e. This loop corresponds to a word w in the g_α's and their inverses, and the element of G determined by w is e since we are dealing with a loop. Since $\langle\, g_\alpha \mid r_\beta \,\rangle$ is a presentation of G, the word w is a product of conjugates of the r_β's. Hence the loop corresponding to w can be homotoped, in the Cayley graph, to a product of loops corresponding to these conjugates. Each of the latter loops is nullhomotopic by the definition of the 2-cells of $\widetilde{X}_G$, the effect of the conjugations being just to change basepoints.

The group G acts on $\widetilde{X}_G$ by multiplication on the left, an element $g \in G$ sending a vertex $g' \in G$ to the vertex gg'. The action is obviously a covering space action, and the orbit space is just X_G.

Example 1. When G is the free group on two generators a and b, X_G is $S^1 \vee S^1$ and $\widetilde{X}_G$ is the Cayley graph of $\mathbb{Z} * \mathbb{Z}$ pictured at the right. The action of a on this graph is a rightward shift along the central horizontal axis, while b acts by an upward shift along the central vertical axis. The composition ab of these two shifts then takes the vertex e to the vertex ab. Similarly, the action of any $w \in \mathbb{Z} * \mathbb{Z}$ takes e to the vertex w.

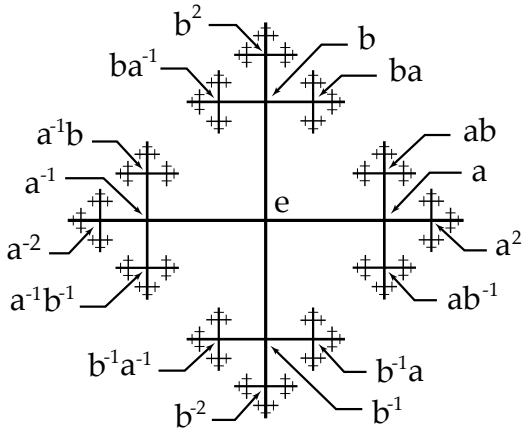

Example 2. The group $G = \mathbb{Z} \times \mathbb{Z}$ with presentation $\langle\, x, y \mid xyx^{-1}y^{-1} \,\rangle$ has $X_G = S^1 \times S^1$, and $\widetilde{X}_G$ is $\mathbb{R}^2$ with vertices the integer lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$ and edges the horizontal and vertical segments between these lattice points. The action of G is by translations $(x, y) \mapsto (x + m, y + n)$.

Example 3. For $G = \mathbb{Z}_2 = \langle\, x \mid x^2 \,\rangle$, X_G is $\mathbb{R}\mathrm{P}^2$ and $\widetilde{X}_G = S^2$. More generally, for $\mathbb{Z}_n = \langle\, x \mid x^n \,\rangle$, X_G is S^1 with a disk attached by the map $z \mapsto z^n$ and $\widetilde{X}_G$ consists of n disks $D_1, \cdots, D_n$ with their boundary circles identified. A generator of $\mathbb{Z}_n$ acts on this union of disks by sending D_i to D_{i+1} via a $2\pi/n$ rotation, the subscript i being taken mod n. The common boundary circle of the disks is rotated by $2\pi/n$.

Example 4. If $G = \mathbb{Z}_2 * \mathbb{Z}_2 = \langle\, a, b \mid a^2, b^2 \,\rangle$ then the Cayley graph is a union of an infinite sequence of circles each tangent to its two neighbors.

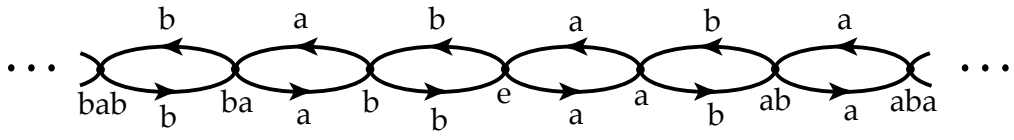

We obtain $\widetilde{X}_G$ from this graph by making each circle the equator of a 2-sphere, yielding an infinite sequence of tangent 2-spheres. Elements of the index-two normal subgroup $\mathbb{Z} \subset \mathbb{Z}_2 * \mathbb{Z}_2$ generated by ab act on $\widetilde{X}_G$ as translations by an even number of units, while each of the remaining elements of $\mathbb{Z}_2 * \mathbb{Z}_2$ acts as the antipodal map on one of the spheres and flips the whole chain of spheres end-for-end about this sphere. The orbit space X_G is $\mathbb{R}\mathrm{P}^2 \vee \mathbb{R}\mathrm{P}^2$.

It is not hard to see the generalization of this example to $\mathbb{Z}_m * \mathbb{Z}_n = \langle\, a, b \mid a^m, b^n \,\rangle$, when $\widetilde{X}_G$ consists of an infinite union of copies of the Cayley complexes for $\mathbb{Z}_m$ and $\mathbb{Z}_n$ constructed in Example 3, arranged in a tree-like pattern. The case of $\mathbb{Z}_2 * \mathbb{Z}_3$ is pictured below.

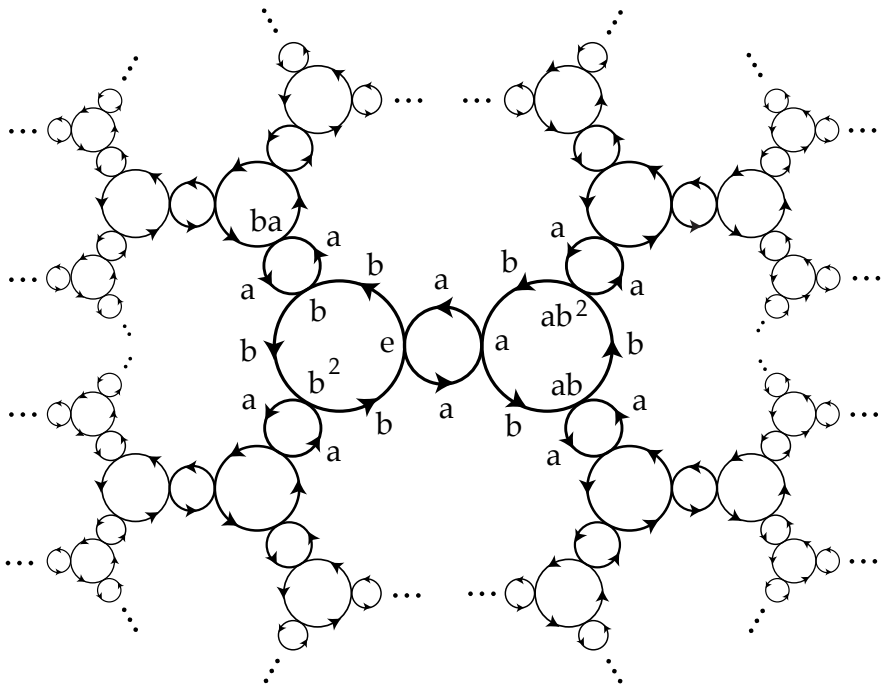

Exercises

1. Show that if $p_1 : \widetilde{X}_1 \to X_1$ and $p_2 : \widetilde{X}_2 \to X_2$ are covering spaces, so is their product $p_1 \times p_2 : \widetilde{X}_1 \times \widetilde{X}_2 \to X_1 \times X_2$.

2. Let $p : \widetilde{X} \to X$ be a covering space with $p^{-1}(x)$ finite for all $x \in X$. Show that $\widetilde{X}$ is Hausdorff iff X is, and similarly with 'compact' in place of 'Hausdorff.'

3. Let X be the subspace of $\mathbb{R}^2$ consisting of the four sides of the square $[0,1] \times [0,1]$ together with the segments of the vertical lines $x = 1/n$, $n \in \mathbb{Z}$, inside the square. Show that X has no simply-connected covering space.

4. Let X be the subspace of $\mathbb{R}^2$ consisting of the circles of radius $1/n$ and center $(1/n, 0)$, $n = 1, 2, \cdots$ and let $\widetilde{X}$ be its covering space shown in the figure below. Construct a two-sheeted covering space $Y \to \widetilde{X}$ such that the composition $Y \to \widetilde{X} \to X$ of the two covering spaces is not a covering space. Note that a composition of two covering spaces does have the unique path lifting property, however.

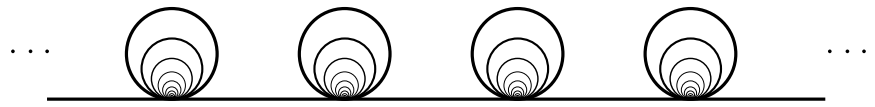

5. Does the Borsuk-Ulam theorem hold for the torus, i.e., for every map $f : S^1 \times S^1 \to \mathbb{R}^2$ must there exist $(x, y) \in S^1 \times S^1$ such that $f(x, y) = f(-x, -y)$?

6. Let A_1, A_2, A_3 be compact sets in $\mathbb{R}^3$. Use the Borsuk-Ulam theorem to show that there is one plane $P \subset \mathbb{R}^3$ which simultaneously divides each A_i into halves of equal measure. [For a unit vector $v \in S^2$ let $P_i(v)$ be such a bisecting plane for A_i orthogonal to v. Show (or perhaps assume, to avoid measure theoretic questions) that $P_i(v)$ can be chosen to vary continuously with v.]

7. Find all the connected 2-sheeted and 3-sheeted covering spaces of $S^1 \vee S^1$, up to isomorphism of covering spaces without basepoints.

8. Construct finite graphs X_1 and X_2 having a common finite-sheeted covering space $\widetilde{X}_1 = \widetilde{X}_2$, but such that there is no space having both X_1 and X_2 as covering spaces.

9. Let a and b be the generators of $\pi_1(S^1 \vee S^1)$ corresponding to the two S^1 summands. Draw pictures of the covering spaces of $S^1 \vee S^1$ corresponding to the normal subgroups $N(p, q, r)$ generated by a^p, b^q, and $(ab)^r$ for integers $p, q, r \geq 2$ satisfying $p^{-1} + q^{-1} + r^{-1} \geq 1$. [First enumerate the various possibilities for p, q, r. These covering spaces have very nice symmetry. The covering space is compact iff $p^{-1} + q^{-1} + r^{-1} > 1$.]

10. Determine the covering space of $S^1 \vee S^1$ corresponding to the subgroup of $\pi_1(S^1 \vee S^1)$ generated by the cubes of all elements. This subgroup is normal and contains the subgroup

$N(3,3,3)$ in the preceding exercise. The covering space is 27-sheeted and can be drawn on a torus so that the complementary regions are 9 triangles with edges labelled aaa, 9 triangles with edges labelled bbb, and 9 hexagons with edges labelled $ababab$. [If the problem were modified by replacing 'cubes' with 'sixth powers,' the resulting covering space would have $2^{28}3^{25}$ sheets! And for k^{th} powers with k sufficiently large, the covering space would have infinitely many sheets. The underlying group theory question here is known as Burnside's problem. It can also be asked for a free group on n generators.]

11. Let Y be the 'quasi-circle' shown in the figure, a closed subspace of $\mathbb{R}^2$ consisting of a portion of the graph of $y = \sin(1/x)$, the segment $[-1,1]$ in the y-axis, and an arc connecting these two pieces. Collapsing the segment of Y in the y-axis to a point gives a quotient map $f:Y \to S^1$. Show that f does not lift to the covering space $\mathbb{R} \to S^1$, even though $\pi_1(Y) = 0$. Thus local path-connectedness of Y is a necessary hypothesis in the lifting criterion.

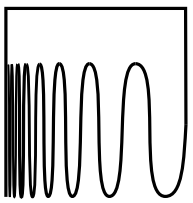

12. Let $\widetilde{X}$ and $\widetilde{Y}$ be simply-connected covering spaces of the path-connected, locally path-connected spaces X and Y. Show that if $X \simeq Y$ then $\widetilde{X} \simeq \widetilde{Y}$. [Hint: See exercise 9 in §0.1.]

13. Find all the connected covering spaces of $\mathbb{R}\mathrm{P}^2 \vee \mathbb{R}\mathrm{P}^2$.

14. Given covering space actions of groups G_1 on X_1 and G_2 on X_2, show that the action of $G_1 \times G_2$ on $X_1 \times X_2$ defined by $(g_1, g_2)(x_1, x_2) = (g_1(x_1), g_2(x_2))$ is a covering space action, and that $(X_1 \times X_2)/(G_1 \times G_2)$ is homeomorphic to $X_1/G_1 \times X_2/G_2$.

15. Show that if a group G acts freely and properly discontinuously on a Hausdorff space X, then the action is a covering space action. (Here 'properly discontinuously' means each $x \in X$ has a neighborhood U such that $\{\, g \in G \mid U \cap g(U) \neq \varnothing \,\}$ is finite.) In particular, a free action of a finite group on a Hausdorff space is a covering space action.

16. Let $\varphi:\mathbb{R}^2 \to \mathbb{R}^2$ be the linear transformation $\varphi(x,y) = (2x, y/2)$. This generates an action of $\mathbb{Z}$ on $X = \mathbb{R}^2 - \{0\}$. Show this action is a covering space action, describe the orbit space $X/\mathbb{Z}$ (it is non-Hausdorff), and compute $\pi_1(X/\mathbb{Z})$.

17. Let Y be path-connected, locally path-connected, and simply-connected, and let G_1 and G_2 be subgroups of $Homeo(Y)$ defining covering space actions on Y. Show that the orbit spaces Y/G_1 and Y/G_2 are homeomorphic iff G_1 and G_2 are conjugate subgroups of $Homeo(Y)$.

18. For integers $m, n \geq 2$, let $X_{m,n}$ be the quotient space of a cylinder $S^1 \times I$ with the identifications $(z,0) \sim (e^{2\pi i/m}z, 0)$ and $(z,1) \sim (e^{2\pi i/n}z, 1)$. Show that $\pi_1(X_{m,n}) =$

$\langle a, b \mid a^m b^{-n} \rangle$, a free product of two copies of $\mathbb{Z}$ with subgroups $m\mathbb{Z}$ and $n\mathbb{Z}$ amalgamated. Also, describe the universal cover of $X_{m,n}$ and the action of $\pi_1(X_{m,n})$. Note that $X_{2,2}$ is the Klein bottle.

19. For relatively prime positive integers m and n consider the **torus knot** $K_{m,n} \subset \mathbb{R}^3$ which is the image of the embedding $f : S^1 \to S^1 \times S^1 \subset \mathbb{R}^3$, $f(z) = (z^m, z^n)$, where the torus $S^1 \times S^1$ is embedded in $\mathbb{R}^3$ in the standard way. The knot $K_{m,n}$ then winds m times around the torus in the longitudinal direction and n times around in the meridional direction, as shown in the figure below for the cases $(m, n) = (3, 2)$ and $(4, 3)$. Viewing S^3 as the one-point compactification of $\mathbb{R}^3$, show that $S^3 - K_{m,n}$ deformation retracts onto a copy of the space $X_{m,n}$ in the previous exercise. [The torus $S^1 \times S^1$ separates S^3 into two solid tori $S^1 \times D^2$ and $D^2 \times S^1$. Formally, $S^3 = \partial D^4 = \partial(D^2 \times D^2) = \partial D^2 \times D^2 \cup D^2 \times \partial D^2$. Embed $X_{m,n}$ in S^3 so that the two halves of $X_{m,n}$ which are the quotients of $S^1 = [0, \frac{1}{2}]$ and $S^1 \times [\frac{1}{2}, 1]$ are the intersections of $X_{m,n}$ with the two solid tori $S^1 \times D^2$ and $D^2 \times S^1$.]

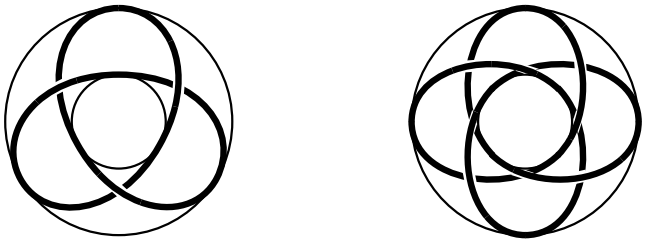

20. Draw the Cayley graph of the group $\mathbb{Z} * \mathbb{Z}_2 = \langle a, b \mid b^2 \rangle$.

21. Show that the normal covering spaces of $S^1 \vee S^1$ are precisely the graphs which are Cayley graphs of groups with two generators. More generally, the normal covering spaces of the wedge sum of n circles are the Cayley graphs of groups with n generators.

22. Consider covering spaces $p : \widetilde{X} \to X$ with $\widetilde{X}$ and X connected CW complexes, the cells of $\widetilde{X}$ projecting homeomorphically onto cells of X. Restricting p to the 1-skeleton then gives a covering space $\widetilde{X}^1 \to X^1$ over the 1-skeleton of X. Show:

(a) Two such covering spaces $\widetilde{X}_1 \to X$ and $\widetilde{X}_2 \to X$ are isomorphic iff the restrictions $\widetilde{X}_1^1 \to X^1$ and $\widetilde{X}_2^1 \to X^1$ are isomorphic.

(b) $\widetilde{X} \to X$ is a normal covering space iff $\widetilde{X}^1 \to X^1$ is.

(c) The groups of deck transformations of the coverings $\widetilde{X} \to X$ and $\widetilde{X}^1 \to X^1$ are isomorphic, via the restriction map.

23. Construct non-normal covering spaces of the Klein bottle by a Klein bottle and by a torus.

4. Additional Topics

Since all groups can be realized as fundamental groups of spaces, this opens the way for using topology to study algebraic properties of groups. The topics in this section give some illustrations of this principle. In particular, we shall see that covering spaces can provide much illumination.

We remind the reader that this section, like the concluding sections of the subsequent chapters, is to be regarded as 'extra material,' not a necessary part of the basic core of the book. Readers who are eager to move on to new topics in later chapters should feel free to skip this section.

a. Graphs and Free Groups

By definition, a **graph** is a 1-dimensional CW complex, i.e., a space X obtained from a discrete set X^0 by attaching a collection of 1-cells e_α. Thus X is constructed from the disjoint union $X^0 \coprod_\alpha I_\alpha$, where I_α is a copy of the interval I, by identifying the two endpoints of each I_α with points of X^0. The points of X^0 are the **vertices** and the 1-cells the **edges** of X. The two endpoints of an edge can be the same vertex, so the closure $\overline{e}_\alpha$ of an edge e_α is homeomorphic either to I or to S^1.

Since X has the quotient topology from $X^0 \coprod_\alpha I_\alpha$, a subset of X is open (or closed) iff it intersects the closure $\overline{e}_\alpha$ of each edge e_α in an open (or closed) set in $\overline{e}_\alpha$. One says X has the **weak topology** with respect to the subspaces $\overline{e}_\alpha$. In this topology a sequence of points in the interiors of distinct edges forms a closed subset, hence never converges. This is true in particular if the edges containing the sequence all have a common vertex and one tries to choose the sequence so that it gets 'closer and closer' to the vertex. Thus if there is a vertex which is the endpoint of infinitely many edges, then the weak topology cannot be a metric topology. An exercise at the end of this subsection is to show the converse, that the weak topology is a metric topology if each vertex is an endpoint of only finitely many edges.

A basis for the topology of X consists of the open segments in the edges e_α together with the path-connected neighborhoods of the vertices v, i.e., the sets U which are unions of connected open neighborhood U_α of v in $\overline{e}_\alpha$ for all $\overline{e}_\alpha$ containing v. In particular, a graph is locally path-connected, so it is connected iff it is path-connected.

If X has only finitely many vertices and edges, then X is compact, being the continuous image of the compact space $X^0 \coprod_\alpha I_\alpha$. The converse is also true, and more generally, a compact subset C of a graph X can meet only finitely many vertices and edges e_α of X. For let $D \subset C$ consist of the vertices in C together with one point in each edge e_α

which C meets. Then D is a closed subset of X since it meets each $\overline{e}_\alpha$ in a closed set. For the same reason, any subset of D is closed, so D has the discrete topology. But D is compact, being a closed subset of the compact space C, so D must be finite, and hence C can meet only finitely many vertices and edges e_α.

A **subgraph** of a graph X is a subspace $Y \subset X$ which is a union of vertices and edges of X, such that $e_\alpha \subset Y$ implies $\overline{e}_\alpha \subset Y$. The latter condition just says that Y is a closed subspace of X. A **tree** is a connected graph containing no subgraph homeomorphic to a circle. By a tree in a graph X we mean a subgraph which is a tree. We will be interested in **maximal** trees in X, i.e., trees not properly contained in any larger tree in X.

Proposition 1.28. *Let X be a connected graph. Then X contains a maximal tree, and in fact any tree in X is contained in a maximal tree. A tree in X is maximal iff it contains all the vertices of X. A tree deformation retracts onto any subtree, and in particular onto any vertex, so trees are contractible.*

Proof: Given a subgraph $X_0 \subset X$, we construct a sequence of subgraphs $X_0 \subset X_1 \subset X_2 \subset \cdots$ inductively, letting X_{i+1} be obtained from X_i by adjoining the closures $\overline{e}_\alpha$ of all edges $e_\alpha \subset X - X_i$ having at least one endpoint in X_i. The union $\bigcup_i X_i$ is open in X since a neighborhood of a point in X_i is contained in X_{i+1}, and $\bigcup_i X_i$ is closed since it is a union of closed edges. So $X = \bigcup_i X_i$ if X is connected. Now let X_0 be any tree T_0 in X, and construct inductively a sequence of trees $T_0 \subset T_1 \subset T_2 \subset \cdots$ such that $T_i \subset X_i$ and T_i contains all the vertices of X_i, by letting T_{i+1} be obtained from T_i by adjoining one edge connecting each vertex of $X_{i+1} - X_i$ to X_i. Clearly T_{i+1} is a tree if T_i is. Hence $\bigcup_i T_i$ is a tree since any subgraph homeomorphic to S^1 would be compact, containing only finitely many edges and therefore contained in some T_i. Because $\bigcup_i T_i$ contains all vertices of X it is maximal, since adding another edge would create a circle subgraph. If T_0 were itself maximal, we would have $T_0 = \bigcup_i T_i$ so T_0 would contain all vertices of X.

If X is a tree then it is a maximal subtree of itself, so $X = \bigcup_i T_i$. It is evident that T_{i+1} deformation retracts to T_i, and we may obtain a deformation retraction of X to T_0 by performing the deformation retraction of T_{i+1} to T_i during the time interval $[1/2^{i+1}, 1/2^i]$. Thus a point $x \in T_{i+1} - T_i$ is stationary until this interval, when it moves into T_i and thereafter continues moving until it reaches T_0. Continuity of the resulting homotopy $h_t : X \to X$ follows from the fact that X has the weak topology. □

Given a maximal tree $T \subset X$ and a base vertex $x_0 \in T$, then each edge e_α of $X - T$ determines a loop f_α in X which goes first from x_0 to one endpoint of e_α by a path in T,

then across e_α, then back to x_0 by a path in T. Strictly speaking, we should first orient the edge e_α in order to specify which direction to cross it. Note that the homotopy class of f_α is independent of the choice of the paths in T since T is simply-connected.

Proposition 1.29. *For a connected graph X, $\pi_1(X, x_0)$ is a free group with basis the classes $[f_\alpha]$ corresponding to the edges e_α of $X - T$.*

Proof: For each edge e_α of $X - T$, let A_α be an open neighborhood of $T \cup e_\alpha$ which deformation retracts to $T \cup e_\alpha$, obtained by deleting a closed interval from the interior of each edge $e_\beta \neq e_\alpha$ of $X - T$. Then $T \cup e_\alpha$ deformation retracts to a circle C_α, the union of e_α with the arc in T joining the endpoints of e_α, since the closures of the components of $(T \cup e_\alpha) - C_\alpha$ are trees which deformation retract to their intersection points with C_α. Thus A_α deformation retracts onto C_α, and it follows that $\pi_1(A_\alpha) \approx \mathbb{Z}$, generated by $[f_\alpha]$. Since the intersection of at least two A_α's is contractible, deformation retracting onto T, we obtain $\pi_1(X) \approx *_\alpha \pi_1(A_\alpha)$. □

This implies in particular that every loop in a graph is homotopic to an **edgepath**: a path consisting of a finite sequence of edges, each traversed monotonically. One can assume the edgepath has no backtracking, i.e., does not cross an edge and then immediate retrace this edge in the reverse direction. For an edgepath without backtracking we can read off the corresponding reduced word in the free group $\pi_1(X)$ from the sequence of edges of $X - T$ which the loop traverses, taking orientations of these edges into account.

Another way to prove the proposition would be to show that the quotient map $X \to X/T$ is a homotopy equivalence, a special case of Proposition 0.2. The quotient X/T is a graph with only one vertex, hence is a wedge sum of circles, whose fundamental group we showed in §1.2 to be free with basis the loops given by the edges of X/T, which are the images of the loops f_α in X.

The next result will be a key step in the proof of the theorem that follows.

Lemma 1.30. *Every covering space of a graph is also a graph, with vertices and edges the lifts of the vertices and edges in the base graph.*

Proof: Let $p : \widetilde{X} \to X$ be the covering space. For the vertices of $\widetilde{X}$ we take the discrete set $\widetilde{X}^0 = p^{-1}(X^0)$. Writing $X = X^0 \coprod_\alpha I_\alpha / \sim$ as in the definition of a graph and applying the path lifting property to the maps $I_\alpha \to X$, we get a unique lift $I_\alpha \to \widetilde{X}$ passing through each point in $p^{-1}(x)$, for $x \in e_\alpha$. These lifts define the edges of a graph structure on $\widetilde{X}$. The resulting topology on $\widetilde{X}$ is the same as its original topology since both topologies have the same basic open sets, the covering projection $\widetilde{X} \to X$ being a local homeomorphism. □

We can apply what we know about fundamental groups and graphs to prove a basic fact of group theory:

Theorem 1.31. *Every subgroup of a free group is free.*

Proof: Given a free group F and a subgroup G, there is a graph X with $\pi_1(X) \approx F$, for example a wedge of circles corresponding to a basis for F. By Proposition 1.22 there is a covering space $p : \widetilde{X} \to X$ with $p_*(\pi_1(\widetilde{X})) = G$, hence $\pi_1(\widetilde{X}) \approx G$ since p_* is injective by Proposition 1.17. Since $\widetilde{X}$ is a graph by the preceding lemma, $\pi_1(\widetilde{X}) \approx G$ is a free group. □

Exercises

1. Let X be a graph in which each vertex is an endpoint of only finitely many edges. Show the weak topology on X is a metric topology.

2. Show that a connected graph retracts onto any connected subgraph.

3. Construct a connected graph X and maps $f, g : X \to X$ such that $fg = \mathbb{1}$ but f and g do not induce isomorphisms on π_1. [Note that $f_*g_* = \mathbb{1}$ implies that f_* is surjective and g_* is injective.]

4. Let F be the free group on 2 generators, F' its commutator subgroup. Find a set of free generators for F' by considering the covering space of the graph $S^1 \vee S^1$ corresponding to F'.

5. If F is a finitely generated free group and N is a nontrivial normal subgroup of infinite index, show, using covering spaces, that N is not finitely generated.

6. Show that a finitely generated group has only a finite number of subgroups of a given finite index. [First do the case of free groups, using covering spaces of graphs. The general case then follows since every group is a quotient group of a free group.]

7. Using covering spaces, show that an index n subgroup H of a group G has at most n conjugate subgroups gHg^{-1} in G. Apply this to show that there exists a normal subgroup $K \subset G$ of finite index with $K \subset H$. [Consider the intersection of all the conjugate subgroups gHg^{-1}.]

8. Let X be the wedge sum of n circles, with its natural graph structure, and let $\widetilde{X} \to X$ be a covering space with $Y \subset \widetilde{X}$ a finite connected subgraph. Show there is a finite graph $Z \supset Y$ having the same vertices as Y, such that the projection $Y \to X$ extends to a covering space $Z \to X$.

9. Apply the two preceding problems to show that if F is a finitely generated free group and $x \in F$ is not the identity element, then there is a normal subgroup $H \subset F$ of finite

index such that $x \notin H$. Hence x has nontrivial image in a finite quotient group of F. In this situation one says F is *residually finite*.

10. Let F be a finitely generated free group, $H \subset F$ a finitely generated subgroup, and $x \in F - H$. Show there is a subgroup K of finite index in F such that $K \supset H$ and $x \notin K$. [Apply exercise 8.]

11. If X is a finite graph and $f : S^1 \to X$ is an embedding whose image contains the basepoint x_0, show that $\pi_1(X, x_0)$ has a basis in which one element is represented by the loop f.

12. Let x be a nontrivial element of a finitely generated free group F. Show there is a finite-index subgroup $H \subset F$ in which x is one element of a basis.

13. Show that the existence of maximal trees is equivalent to the axiom of choice.

b. K(G,1) Spaces

In this subsection we introduce a class of spaces whose homotopy type depends only on their fundamental group. These spaces arise many places in topology, especially where topology and group theory intermingle.

A path-connected space whose fundamental group is isomorphic to a given group G and which has a contractible universal covering space is called a **K(G,1) space**. The '1' here refers to π_1. More general $K(G,n)$ spaces are studied in §4.1. These are usually called Eilenberg-MacLane spaces, though in the case $n = 1$ they were studied by Hurewicz before Eilenberg and MacLane took up the general case. Here are some examples:

1. S^1 is a $K(\mathbb{Z},1)$. More generally, a connected graph is a $K(G,1)$ with G a free group, since by the results of the preceding subsection its universal cover is a tree, hence contractible.
2. Closed surfaces with infinite π_1, i.e., closed surfaces other than S^2 and $\mathbb{R}\mathrm{P}^2$, are $K(G,1)$'s. This will be shown in the next subsection. It also follows from the theorem in surface theory that the only simply-connected surfaces without boundary are S^2 and $\mathbb{R}^2$, so the universal cover of a closed surface with infinite fundamental group must be noncompact, hence $\mathbb{R}^2$. Non-closed surfaces deformation retract onto graphs, so are $K(G,1)$'s with G free.
3. The infinite-dimensional projective space $\mathbb{R}\mathrm{P}^\infty$ is a $K(\mathbb{Z}_2,1)$ since its universal cover is S^∞ which is contractible. To show the latter fact, a homotopy from the identity map of S^∞ to a constant map can be constructed in two stages as follows. First, define $f_t : \mathbb{R}^\infty \to \mathbb{R}^\infty$ by $f_t(x_1, x_2, \cdots) = (1-t)(x_1, x_2, \cdots) + t(0, x_1, x_2, \cdots)$. This takes nonzero vectors to nonzero vectors for all $t \in [0,1]$, so $f_t/|f_t|$ gives a homotopy from the identity map of S^∞ to the map $(x_1, x_2, \cdots) \mapsto (0, x_1, x_2, \cdots)$. A homotopy from this to a constant map is given by $g_t/|g_t|$ where $g_t(x_1, x_2, \cdots) = (1-t)(0, x_1, x_2, \cdots) + t(1, 0, 0, \cdots)$.
4. Generalizing the preceding example, we can construct a $K(\mathbb{Z}_m,1)$ as an infinite dimensional lens space $S^\infty/\mathbb{Z}_m$, where $\mathbb{Z}_m$ acts on S^∞, regarded as the unit sphere in $\mathbb{C}^\infty$, by scalar multiplication by m^{th} roots of unity, a generator of this action being the map $(z_1, z_2, \cdots) \mapsto e^{2\pi i/m}(z_1, z_2, \cdots)$. It is not hard to check that this is a covering space action.
5. A product $K(G,1) \times K(H,1)$ is a $K(G \times H, 1)$ since its universal cover is the product of the universal covers of $K(G,1)$ and $K(H,1)$. By taking products of the examples in (1) and (4) we therefore get $K(G,1)$'s for arbitrary finitely generated abelian groups G. For example the n-dimensional torus T^n, the product of n circles, is a $K(\mathbb{Z}^n,1)$.

It is not hard to construct a $K(G,1)$ for an arbitrary group G, using the notion of a Δ-complex defined at the beginning of §2.1. Let EG be the Δ-complex whose n-simplices are the ordered $(n+1)$-tuples $[g_0,\cdots,g_n]$ of elements of G. Such an n-simplex attaches to the $(n-1)$-simplices $[g_0,\cdots,\widehat{g}_i,\cdots,g_n]$ in the obvious way, just as a standard simplex attaches to its faces. (The notation $\widehat{g}_i$ means that this vertex is deleted.) The complex EG is contractible by the homotopy h_t which deforms each point in the simplex $[g_0,\cdots,g_n]$ linearly in the simplex $[e,g_0,\cdots,g_n]$ to the vertex $[e]$, where e is the identity element of G. This is well defined in EG since when we restrict to a face $[g_0,\cdots,\widehat{g}_i,\cdots,g_n]$ we have the linear homotopy to $[e]$ in $[e,g_0,\cdots,\widehat{g}_i,\cdots,g_n]$. Note that h_t carries $[e]$ around the loop $[e,e]$, so h_t is not actually a deformation retraction of EG onto $[e]$.

The group G acts on EG by left multiplication, an element $g\in G$ taking the simplex $[g_0,\cdots,g_n]$ linearly onto the simplex $[gg_0,\cdots,gg_n]$. Only the identity e takes any simplex to itself, so by an exercise at the end of this subsection, the action of G on EG is a covering space action. Hence the quotient map $EG\to EG/G$ is the universal cover of the orbit space $BG=EG/G$, and BG is a $K(G,1)$.

Since G acts on EG by freely permuting simplices, BG inherits a Δ-complex structure from EG. The action of G on EG identifies all the vertices of EG, so BG has just one vertex. To describe the Δ-complex structure on BG explicitly, note first that every n-simplex of EG can be written uniquely in the form $[g_0,g_0g_1,g_0g_1g_2,\cdots,g_0g_1\cdots g_n]=g_0[e,g_1,g_1g_2,\cdots,g_1\cdots g_n]$. The image of this simplex in BG may be denoted unambiguously by the symbol $[g_1|g_2|\cdots|g_n]$. In this 'bar' notation the g_i's and their ordered products can be viewed as labels for edges, thinking of edge labels as ratios between the group elements corresponding to the endpoints of the edge, as indicated in the figure below.

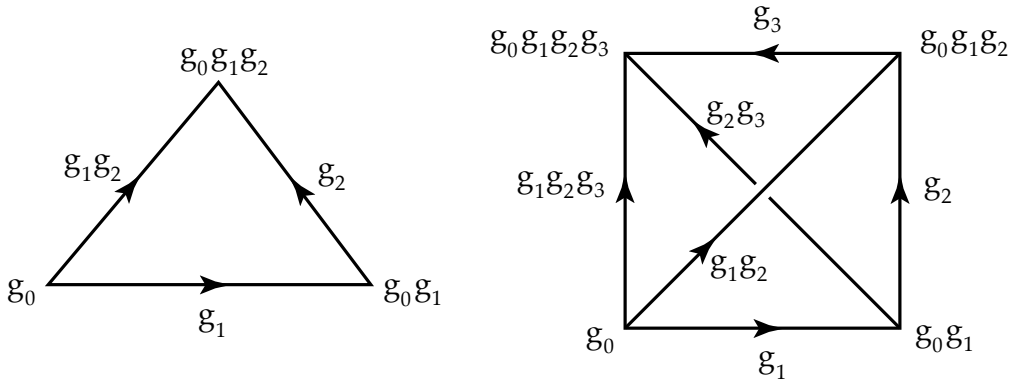

With this notation, the boundary of a simplex $[g_1|\cdots|g_n]$ of BG consists of the simplices $[g_2|\cdots|g_n]$, $[g_1|\cdots|g_{n-1}]$, and $[g_1|\cdots|g_ig_{i+1}|\cdots|g_n]$ for $i=1,\cdots,n-1$.

This construction of a $K(G,1)$ produces a rather large space, since BG is always infinite-dimensional, and if G is infinite, BG has an infinite number of cells in each positive

dimension. For example, $B\mathbb{Z}$ is much bigger than S^1, the most efficient $K(\mathbb{Z},1)$. On the other hand, BG has the virtue of being functorial: A homomorphism $f:G\to H$ induces a map $Bf:BG\to BH$ sending a simplex $[g_1|\cdots|g_n]$ to the simplex $[f(g_1)|\cdots|f(g_n)]$. A different construction of a $K(G,1)$ is given in §4.1. Here one starts with any 2-dimensional complex having fundamental group G, for example the complex X_G associated to a presentation of G, and then one attaches cells of dimension 3 and higher to make the universal cover contractible without affecting π_1. In general, it is hard to get any control on the number of higher-dimensional cells needed in this construction, so it too can be rather inefficient. Indeed, it is often a difficult problem to construct an efficient $K(G,1)$ for a given group G. A curious and almost paradoxical fact which is relevant to this efficiency question is that if G contains any elements of finite order, then every $K(G,1)$ CW complex must be infinite-dimensional. This is shown in Example 7 in §2.2. Thus the infinite-dimensional lens space $K(\mathbb{Z}_m,1)$'s cannot be replaced by any finite-dimensional complex.

In spite of the great latitude possible in the construction of $K(G,1)$'s, there is a very nice uniqueness statement, which accounts for much of the interest in $K(G,1)$'s:

Theorem 1.32. *The homotopy type of a CW complex $K(G,1)$ is uniquely determined by G.*

Having a unique homotopy type of $K(G,1)$'s associated to each group G means that algebraic invariants of spaces which depend only on homotopy type, such as homology and cohomology groups, become invariants of groups. This has proved to be a quite fruitful idea, and has been much studied. A standard reference is [K.Brown], and other books at a more advanced level are [Adem-Milgram] and [Benson 1992].

The theorem will follow easily from:

Proposition 1.33. *Let X be a connected CW complex and Y a $K(G,1)$. Then any homomorphism $\pi_1(X,x_0)\to\pi_1(Y,y_0)$ is induced by a map $(X,x_0)\to(Y,y_0)$ which is unique up to homotopy fixing x_0.*

To deduce the theorem from this, let X and Y be CW complex $K(G,1)$'s with isomorphic G's. By the proposition, there are maps $f:(X,x_0)\to(Y,y_0)$ and $g:(Y,y_0)\to(X,x_0)$ inducing inverse isomorphisms $\pi_1(X,x_0)\approx\pi_1(Y,y_0)$. Then fg and gf induce the identity on π_1 hence are homotopic to the identity maps.

Proof of 1.33: Let us first consider the case that X has a single 0-cell, x_0. Given a homomorphism $\varphi:\pi_1(X,x_0)\to\pi_1(Y,y_0)$, the construction of a map $f:(X,x_0)\to(Y,y_0)$ with $f_*=\varphi$ is done inductively over the skeleta of X, starting with $f(x_0)=y_0$. Each 1-cell e^1_α of X has closure a circle determining an element $[e^1_\alpha]\in\pi_1(X,x_0)$, and we let f

on the closure of e^1_α be a map representing $\varphi([e^1_\alpha])$. If $i : X^1 \hookrightarrow X$ denotes the inclusion, then $\varphi i_* = f_*$ since $\pi_1(X^1, x_0)$ is generated by the elements $[e^1_\alpha]$.

$$\begin{array}{ccc} \pi_1(X^1, x_0) & \xrightarrow{\quad f_* \quad} & \pi_1(Y, y_0) \\ & {\scriptstyle i_*}\searrow \quad \nearrow{\scriptstyle \varphi} & \\ & \pi_1(X, x_0) & \end{array}$$

To extend f over a cell e^2_β with attaching map $\psi_\beta : S^1 \to X^1$, all we need is for the composition $f\psi_\beta$ to be nullhomotopic. Choosing a basepoint $s_0 \in S^1$ and a path in X^1 from $\psi_\beta(s_0)$ to x_0, ψ_β determines an element $[\psi_\beta] \in \pi_1(X^1, x_0)$, and the existence of a nullhomotopy of $f\psi_\beta$ is equivalent to $f_*([\psi_\beta])$ being zero in $\pi_1(Y, y_0)$. But in the diagram above we see that $i_*([\psi_\beta]) = 0$ since the cell e^2_β provides a nullhomotopy of ψ_β in X, hence $f_*([\psi_\beta]) = 0$ and f can be extended over e^2_β.

Extending f inductively over cells e^n_γ with $n > 2$ is possible since the attaching maps $\psi_\gamma : S^{n-1} \to X^{n-1}$ have nullhomotopic compositions $f\psi_\gamma : S^{n-1} \to Y$. This is because $f\psi_\gamma$ lifts to the universal cover of Y if $n > 2$, and this cover is contractible by hypothesis, so the lift of $f\psi_\gamma$ is nullhomotopic, hence also $f\psi_\gamma$.

Turning to the uniqueness statement, if two maps $f_0, f_1 : (X, x_0) \to (Y, y_0)$ induce the same homomorphism on π_1, this immediately gives a homotopy $f_t : (X^1, x_0) \to (Y, y_0)$. Extending the resulting map $X^1 \times I \cup X \times I \to Y$ over cells $e^n \times (0, 1)$ for $n > 1$ can be done just as in the previous paragraph since these cells have dimension $n + 1 > 2$.

The case that X has more than one 0-cell can be treated by a small elaboration on this argument. Choose a maximal tree $T \subset X$. The construction of a map f realizing a given φ starts by letting $f(T) = y_0$. Then each edge e^1_α in $X - T$ determines an element $[e^1_\alpha] \in \pi_1(X, x_0)$, and we let f on the closure of e^1_α be a map representing $\varphi([e^1_\alpha])$. Extending f over higher-dimensional cells then proceeds just as before. Constructing a homotopy f_t joining two given maps f_0 and f_1 with $f_{0*} = f_{1*}$ also has an extra step. Let $h_t : X^1 \to X^1$ be a homotopy, starting with $h_0 = \mathbb{1}$, which restricts to a deformation retraction of T to x_0. (It's easy to extend such a deformation retraction to a homotopy defined on all of X^1.) To construct a homotopy $f_0|X^1 \simeq f_1|X^1$, one first deforms $f_0|X^1$ and $f_1|X^1$ to take T to y_0 by composing with h_t. Then the earlier argument can be used to give a homotopy $f_0|X^1 \simeq f_1|X^1$ and to extend this homotopy over all of X. □

The first part of the preceding proof also works for the 2-dimensional complexes X_G associated to presentations of groups. Thus every homomorphism $G \to H$ is realized as the induced homomorphism of some map $X_G \to X_H$. However, there is no uniqueness

statement for this map, and it can easily happen that different presentations of a group G give X_G's which are not homotopy equivalent.

Exercises

1. Suppose a group G acts simplicially on a Δ-complex X, i.e., each element of G takes each n-simplex of X onto another n-simplex by a linear homeomorphism. If the action is free, show it is a covering space action.

c. Graphs of Groups

As an illustration of how $K(G,1)$ spaces can be useful in group theory, we shall describe a procedure for assembling a collection of $K(G,1)$'s together into a $K(G,1)$ for a larger group G. Group-theoretically, this gives a method for assembling smaller groups together to form a larger group, generalizing the notion of free products.

Let Γ be a graph which is connected and oriented, that is, its edges are viewed as arrows, each edge having a specified direction. Suppose that at each vertex v of Γ we place a group G_v and along each edge e of Γ we put a homomorphism φ_e from the group at the tail of the edge to the group at the head of the edge. We call this data a **graph of groups**. Now build a space $B\Gamma$ by putting the space BG_v at each vertex v of Γ and then filling in a mapping cylinder $M_{B\varphi_e}$ along each edge e of Γ, identifying the two ends of the mapping cylinder with the two BG_v's at the ends of e. The resulting space $B\Gamma$ is then a CW complex since the maps $B\varphi_e$ take n-cells homeomorphically onto n-cells. In fact, the cell structure on $B\Gamma$ can be canonically subdivided into a Δ-complex structure using the prism construction from § 2.1, but we will not need to do this here.

More generally, instead of BG_v one could take any CW complex $K(G_v,1)$ at the vertex v, and then along edges put mapping cylinders of maps realizing the homomorphisms φ_e. We leave it for the reader to check that the resulting space $K\Gamma$ is homotopy equivalent to the $B\Gamma$ constructed above.

Example. Suppose Γ consists of one central vertex with a number of edges radiating out from it, and the group G_v at this central vertex is trivial, hence also all the edge homomorphisms. Then van Kampen's theorem implies that $\pi_1(K\Gamma)$ is the free product of the groups at all the outer vertices.

In view of this example, we shall call $\pi_1(K\Gamma)$ for a general graph of groups Γ the **graph product** of the vertex groups G_v with respect to the edge homomorphisms φ_e. (The more standard name for $\pi_1(K\Gamma)$ is 'the fundamental group of the graph of groups,' which seems rather awkward.)

Here is the main result we shall prove about graphs of groups:

Theorem 1.34. *If all the edge homomorphisms φ_e are injective, then $K\Gamma$ is a $K(G,1)$ and the inclusions $K(G_v,1) \hookrightarrow K\Gamma$ induce injective maps on π_1.*

Before giving the proof, let us look at some interesting special cases:

Free Products with Amalgamation. Suppose the graph of groups is $A \leftarrow C \rightarrow B$, with the two maps monomorphisms. One can regard this data as specifying an embedding of C as a subgroup of both A and B. Applying van Kampen's theorem to the decomposition of

$K\Gamma$ into its two mapping cylinders, we see that $\pi_1(K\Gamma)$ is the quotient of $A * B$ obtained by identifying the subgroup $C \subset A$ with the subgroup $C \subset B$. The standard notation for this group is $A *_C B$, the free product of A and B **amalgamated** along the subgroup C. According to the theorem, $A *_C B$ contains both A and B as subgroups.

For example, a free product with amalgamation $\mathbb{Z} *_{\mathbb{Z}} \mathbb{Z}$ can be realized by mapping cylinders of the maps $S^1 \leftarrow S^1 \rightarrow S^1$ which are m-sheeted and n-sheeted covering spaces, respectively. An exercise for §1.3 was to show that the complex $K\Gamma$ in this case is a deformation retract of the complement of a torus knot in S^3, if m and n are relatively prime. Hence the complements of torus knots are $K(G,1)$'s. This is true more generally for all smoothly embedded knots in S^3, and in fact these knot complements can be built up by iterated graph of groups constructions, starting with free groups.

HNN Extensions. This is the case that the graph of groups consists of two monomorphisms $\varphi, \psi : C \rightarrow A$. This is analogous to the previous case $A *_C B$, but with the two groups A and B coalesced to a single group. Accordingly, $\pi_1(K\Gamma)$ is denoted $A*_C$. To see what this group looks like, regard $K\Gamma$ as a quotient of $K(A,1) \amalg K(C,1) \times I$ with the two ends of $K(C,1) \times I$ attached to $K(A,1)$ via maps realizing the monomorphisms φ and ψ. Using a $K(C,1)$ with a single 0-cell, we see that $K\Gamma$ can be obtained from $K(A,1) \vee S^1$ by attaching cells of dimension two and greater, so $\pi_1(K\Gamma)$ is a quotient of $A * \mathbb{Z}$, and it is not hard to figure out that the relations defining this quotient are of the form $t\varphi(c)t^{-1} = \psi(c)$ where t is a generator of the $\mathbb{Z}$ factor and c ranges over C, or a set of generators for C. We leave the verification of this for the Exercises.

As a very special case, taking $\varphi = \psi = \mathbb{1}$ gives $A*_A = A \times \mathbb{Z}$ since we can take $K\Gamma = K(A,1) \times S^1$ in this case. More generally, taking $\varphi = \mathbb{1}$ with ψ an arbitrary automorphism of A, we realize any semi-direct product of A and $\mathbb{Z}$ as $A*_A$. For example, the Klein bottle occurs this way, with φ realized by the identity map of S^1 and ψ by a reflection. In these cases when $\varphi = \mathbb{1}$ we could realize the same group $\pi_1(K\Gamma)$ using a slightly simpler graph of groups, with a single vertex labelled A and a single edge labelled ψ.

Closed Surfaces. A closed orientable surface M of genus two or greater can be cut along a circle into two surfaces-with-boundary M_1 and M_2, such that the closed surfaces obtained from M_1 and M_2 by filling in their boundary circle with a disk have smaller genus than M. Each of M_1 and M_2 is the mapping cylinder of a map from S^1 to a finite graph inducing an injection on π_1. Since graphs are $K(G,1)$'s, the theorem then says that M is a $K(G,1)$. A similar argument works for closed non-orientable surfaces other than $\mathbb{R}\mathrm{P}^2$. For example, the Klein bottle is obtained from two Möbius bands by identifying their

boundary circles, and a Möbius band is the mapping cylinder of the two-sheeted covering space $S^1 \to S^1$.

Proof of 1.34: We shall construct a covering space $\widetilde{K} \to K\Gamma$ by gluing together copies of the universal covering spaces of the various mapping cylinders in $K\Gamma$ in such a way that $\widetilde{K}$ will be contractible. Hence $\widetilde{K}$ will be the universal cover of $K\Gamma$, which will therefore be a $K(G,1)$.

First a preliminary observation: Given a universal covering space $p\colon \widetilde{X} \to X$ and a connected, locally path-connected subspace $A \subset X$ such that the inclusion $A \hookrightarrow X$ induces an injection on π_1, then each component $\widetilde{A}$ of $p^{-1}(A)$ is a universal cover of A. For the restriction $p\colon \widetilde{A} \to A$ is certainly a covering space, so the induced map $\pi_1(\widetilde{A}) \to \pi_1(A)$ is injective, and this map factors through $\pi_1(\widetilde{X}) = 0$, hence $\pi_1(\widetilde{A}) = 0$. For example if X is the torus $S^1 \times S^1$ and A is the circle $S^1 \times \{x_0\}$, then $p^{-1}(A)$ consists of infinitely many parallel lines in $\mathbb{R}^2$, each of which is a universal cover of A.

In particular, consider the universal cover $p\colon \widetilde{M}_f \to M_f$ of the mapping cylinder of a map $f\colon A \to B$ between connected CW complexes. Then $\widetilde{M}_f$ is itself the mapping cylinder of a map $\widetilde{f}\colon p^{-1}(A) \to p^{-1}(B)$ since the line segments in the mapping cylinder structure on M_f lift to line segments in $\widetilde{M}_f$ defining a mapping cylinder structure. Since $\widetilde{M}_f$ is a mapping cylinder, it deformation retracts onto $p^{-1}(B)$, so $p^{-1}(B)$ is also simply-connected, hence is the universal cover of B. If f induces an injection on π_1, then the remarks in the preceding paragraph apply, and the components of $p^{-1}(A)$ are universal covers of A. If we assume further that A and B are $K(G,1)$'s, then $\widetilde{M}_f$ and the components of $p^{-1}(A)$ are contractible, and we claim that $\widetilde{M}_f$ deformation retracts onto each component $\widetilde{A}$ of A. Namely, the inclusion $\widetilde{A} \hookrightarrow \widetilde{M}_f$ is a homotopy equivalence since both spaces are contractible, and then Corollary 0.5 implies that $\widetilde{M}_f$ deformation retracts onto $\widetilde{A}$ since the pair $(\widetilde{M}_f, \widetilde{A})$ satisfies the homotopy extension property, as a consequence of the fact that $\widetilde{A}$ has a neighborhood in $\widetilde{M}_f$ of the form $\widetilde{A} \times I$.

Now we can describe the construction of the covering space $\widetilde{K}$ of $K\Gamma$. It will be the union of an increasing sequence of spaces $\widetilde{K}_1 \subset \widetilde{K}_2 \subset \cdots$. For the first stage, let $\widetilde{K}_1$ be the universal cover of one of the mapping cylinders M_f of $K\Gamma$. By the preceding remarks, this contains various disjoint copies of universal covers of the two $K(G_v,1)$'s at the ends of M_f. We build $\widetilde{K}_2$ from $\widetilde{K}_1$ by attaching to each of these universal covers of $K(G_v,1)$'s a copy of the universal cover of each mapping cylinder M_g of $K\Gamma$ meeting M_f at the end of M_f in question. Now repeat the process to construct $\widetilde{K}_3$ by attaching universal covers of mapping cylinders at all the universal covers of $K(G_v,1)$'s created in the previous step. In the same way, we construct $\widetilde{K}_{n+1}$ from $\widetilde{K}_n$ for all n, and then we set $\widetilde{K} = \bigcup_n \widetilde{K}_n$.

Note that $\widetilde{K}_{n+1}$ deformation retracts onto $\widetilde{K}_n$ since it is formed by attaching pieces to $\widetilde{K}_n$ which deformation retract onto the subspaces along which they attach, by our earlier remarks. It follows that $\widetilde{K}$ is contractible since we can deformation retract $\widetilde{K}_{n+1}$ onto $\widetilde{K}_n$ during the time interval $[1/2^{n+1}, 1/2^n]$, and then finish with a contraction of $\widetilde{K}_1$ to a point during the time interval $[1/2, 1]$.

The natural projection $\widetilde{K} \to K\Gamma$ is clearly a covering space, so this finishes the proof that $K\Gamma$ is a $K(G, 1)$.

The remaining statement that each inclusion $K(G_v, 1) \hookrightarrow K\Gamma$ induces an injection on π_1 can easily be deduced from the preceding constructions. For suppose a loop $\gamma : S^1 \to K(G_v, 1)$ is nullhomotopic in $K\Gamma$. By the lifting criterion for covering spaces, there is a lift $\widetilde{\gamma} : S^1 \to \widetilde{K}$. This has image contained in one of the copies of the universal cover of $K(G_v, 1)$, so $\widetilde{\gamma}$ is nullhomotopic in this universal cover, and hence γ is nullhomotopic in $K(G_v, 1)$. □

The various mapping cylinders which make up the universal cover of $K\Gamma$ are arranged in a tree-like pattern. The tree in question, call it $T\Gamma$, has one vertex for each copy of a universal cover of a $K(G_v, 1)$ in $\widetilde{K}$, and two vertices are joined by an edge whenever the two universal covers of $K(G_v, 1)$'s corresponding to these vertices are connected by a line segment lifting a line segment in the mapping cylinder structure of a mapping cylinder of $K\Gamma$. The inductive construction of $\widetilde{K}$ is reflected in an inductive construction of $T\Gamma$ as a union of an increasing sequence of subtrees $T_1 \subset T_2 \subset \cdots$. Corresponding to $\widetilde{K}_1$ is a subtree $T_1 \subset T\Gamma$ consisting of a central vertex with a number of edges radiating out from it, an 'asterisk' with possibly an infinite number of edges. When we enlarge $\widetilde{K}_1$ to $\widetilde{K}_2$, T_1 is correspondingly enlarged to a tree T_2 by attaching a similar asterisk at the end of each outer vertex of T_1. And so on at each stage. The action of $\pi_1(K\Gamma)$ on $\widetilde{K}$ as deck transformations induces an action on $T\Gamma$, permuting its vertices and edges, and the orbit space of $T\Gamma$ under this action is just the original graph Γ. The action on $T\Gamma$ will not be a free action in general since the elements of a subgroup $G_v \subset \pi_1(K\Gamma)$ fix a vertex corresponding to one of the universal covers of $K(G_v, 1)$. (Basepoints need to be chosen to make this precise.)

There is in fact an exact correspondence between graphs of groups and groups acting on trees. See [Scott-Wall] for an exposition of this rather nice theory. From the viewpoint of groups acting on trees, the definition of a graph of groups is usually taken to be slightly more restrictive than the one we have given here, namely, one considers only oriented graphs obtained from an unoriented graph by subdividing each edge by adding a vertex at its midpoint, then orienting the two resulting edges outward, away from the new vertex.

Exercises

1. Show that every graph product of trivial groups is free.

2. Use van Kampen's theorem to compute $A*_C$ as a quotient of $A * \mathbb{Z}$, as stated in the text.

3. Consider the graph of groups Γ having one vertex, $\mathbb{Z}$, and one edge, the map $\mathbb{Z} \to \mathbb{Z}$ which is multiplication by 2, realized by the 2-sheeted covering space $S^1 \to S^1$. Show that $\pi_1(K\Gamma)$ has presentation $\langle\, a, b \mid bab^{-1}a^{-2} \,\rangle$ and describe the universal cover of $K\Gamma$ explicitly. [This is the first in a family of groups called Baumslag-Solitar groups, having presentations of the form $\langle\, a, b \mid ba^m b^{-1} a^{-n} \,\rangle$. These are HNN extensions $\mathbb{Z}*_{\mathbb{Z}}$.]

4. Show that every graph product of groups can be realized by a graph whose vertices are partitioned into two subsets, with every oriented edge going from a vertex in the first subset to a vertex in the second subset.

5. Show that a finite graph product of finitely generated groups in finitely generated, and similarly for finitely presented groups.

6. Show that a finite graph product of finite groups has a free subgroup of finite index, by constructing a finite-sheeted covering space of $K\Gamma$ from universal covers of the mapping cylinders of $K\Gamma$.) [The converse is also true for finitely generated groups; see [Scott-Wall].]

Chapter 2
Homology

The fundamental group has the twin virtues of an intuitive definition and easy computability from geometry, for example from a cell structure. The natural n-dimensional generalization of the fundamental group is the homotopy group π_n, but unfortunately this lacks the geometric computability of the fundamental group. For example, a space as simple as the 2-sphere has such complicated higher homotopy groups $\pi_n(S^2)$ that it seems unlikely they will ever be completely known. What is needed then is a 'more geometric' substitute for the higher homotopy groups, retaining the computability properties of π_1. The homology groups $H_n(X)$ turn out to fill the bill quite well. However, the computability of homology groups comes at a price. First, just the definition is considerably more sophisticated than for homotopy groups. And then, before one can do any real computations it is necessary to prove a number of basic properties which are not exactly trivial.

In the exposition below we approach the definition of $H_n(X)$ by two preliminary stages, first giving a few motivating examples nonrigorously, then constructing a restricted model of homology theory called simplicial homology, before digging into the general theory. Once the definition of homology has been assimilated, the real work of establishing its basic properties begins. This takes close to 20 pages, and there is no getting around the fact that it requires a substantial effort. This takes up most of the first section of the chapter, with only a couple small digressions for applications, namely two classical theorems of Brouwer, the fixed point theorem and the 'invariance of dimension' theorem.

The second section gives more applications, including the homology definition of Euler characteristic and Brouwer's notion of degree for maps $S^n \to S^n$, but this section is mainly devoted to developing techniques for calculating homology groups efficiently. The maximally efficient method is 'cellular homology' whose power perhaps comes from the fact that it is 'homology squared' — homology defined in terms of homology. Another quite useful tool we develop is Mayer-Vietoris sequences, the analog for homology of van Kampen's theorem for the fundamental group.

An interesting feature of homology which begins to emerge after one has worked with it for awhile is that it is the basic properties of homology which are used most often, and not the actual definition itself. This suggests that an axiomatic approach to homology might be possible. This is indeed the case, and in the third section of this chapter we list axioms which completely characterize homology groups for CW complexes. One could

take the viewpoint that these rather algebraic axioms are all that really matters about homology groups, that the geometry involved in the definition of homology is secondary, needed only to show that the axiomatic theory is not vacuous. The extent to which one adopts this viewpoint is a matter of taste, and the route taken here of postponing the axioms until the theory is well-established is just one of several possible approaches.

The Idea of Homology

Let us see what homology is about by looking at some examples.

Example 1. Consider the graph X_1 shown is the first figure below, consisting of two vertices joined by four edges.

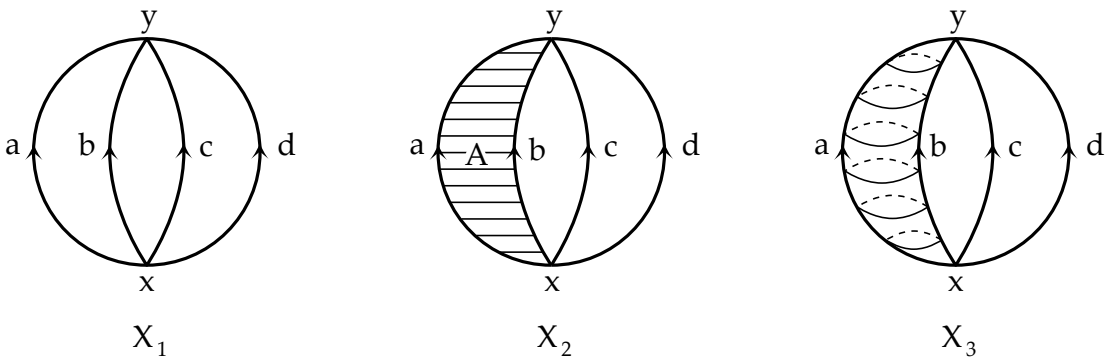

When studying the fundamental group of X_1, we consider loops formed by sequences of edges, starting and ending at a fixed basepoint. For example, at the basepoint x, the loop ab^{-1} travels forward along the edge a then backward along b, as indicated by the exponent -1. A more complicated loop would be $ac^{-1}bd^{-1}ca^{-1}$. A salient feature of the fundamental group is that it is generally nonabelian, which both enriches and complicates the theory. Suppose we simplify matters by abelianizing. Thus for example the two loops ab^{-1} and $b^{-1}a$ are to be regarded as equal if we make a commute with b^{-1}. These two loops ab^{-1} and $b^{-1}a$ are really the same circle, just with a different choice of starting and ending point: x for ab^{-1} and y for $b^{-1}a$. The same thing happens for all loops: Rechoosing the basepoint in a loop just permutes its letters cyclically, so a byproduct of abelianizing is that we no longer have to pin all our loops down to a fixed basepoint. Thus loops become *cycles*, without a chosen basepoint.

Having abelianized, let us switch to additive notation, so cycles become linear combinations of edges with integer coefficients, e.g., $a - b + c - d$. Let us call these linear combinations *chains* of edges. Some chains can be decomposed into cycles in several different ways, for example $(a - c) + (b - d) = (a - d) + (b - c)$, and if we adopt an algebraic viewpoint then we do not wish to distinguish between these different decompositions. Thus

we broaden the meaning of the term 'cycle' to be simply any linear combination of edges for which at least one decomposition into cycles in the previous more geometric sense exists.

What is the condition for a chain to be a cycle in this more algebraic sense? A geometric cycle, thought of as a path traversed in time, is distinguished by the property that it enters each vertex the same number of times that it leaves the vertex. For an arbitrary chain $ka+\ell b+mc+nd$, the net number of times this chain enters y is $k+\ell+m+n$ since each of a, b, c, and d enters y once. Similarly, each of the four edges leaves x once, so the net number of times the chain $ka+\ell b+mc+nd$ enters x is $-k-\ell-m-n$. Thus the condition for $ka+\ell b+mc+nd$ to be a cycle is simply $k+\ell+m+n=0$.

To describe this result in a way which would generalize to all graphs, let C_1 be the free abelian group with basis the edges a, b, c, d and let C_0 be the free abelian group with basis the vertices x, y. Elements of C_1 are chains of edges, or 1-dimensional chains, and elements of C_0 are linear combinations of vertices, or 0-dimensional chains. Define a homomorphism $\partial : C_1 \to C_0$ by sending each basis element a, b, c, d to $y-x$, the vertex at the head of the edge minus the vertex at the tail. Thus $\partial(ka+\ell b+mc+nd) = (k+\ell+m+n)y-(k+\ell+m+n)x$, and the cycles are precisely the kernel of ∂.

It is a simple calculation to verify that $a-b$, $b-c$, and $c-d$ form a basis for this kernel. Thus every cycle in X_1 is a unique linear combination of these three most obvious cycles. By means of these three basic cycles we convey the geometric information that the graph X_1 has three visible 'holes,' the empty spaces between the four edges.

Example 2. Let us enlarge the preceding graph X_1 by attaching a 2-cell A along the cycle $a-b$, producing a 2-dimensional cell complex X_2 shown in the second figure above. If we think of A as being oriented clockwise, then we can regard its boundary as the cycle $a-b$. This cycle is now homotopically trivial since we can contract it to a point by sliding over A. In other words, it no longer encloses a hole in X_2. This suggests that we form a quotient of the group of cycles in the preceding example by factoring out the subgroup generated by $a-b$. In this quotient the cycles $a-c$ and $b-c$, for example, become equivalent, consistent with the fact that they are homotopic in X_2.

Algebraically, we can define now a pair of homomorphisms

$$C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0$$

where C_2 is the infinite cyclic group generated by A and $\partial_2(A) = a-b$. The map ∂_1 is the boundary homomorphism in the previous example. The quotient group we are interested in is $Ker(\partial_1)/Im(\partial_2)$, i.e., the 1-dimensional cycles modulo those which are boundaries, the multiples of $a-b$. This quotient group is the *homology group* $H_1(X_2)$. The previous

example can be fit into this scheme too by taking C_2 to be zero since there are no 2-cells in X_1, so $H_1(X_1) = Ker(\partial_1)/Im(\partial_2) = Ker(\partial_1)$, which as we saw was free abelian on three generators. In the present example, $H_1(X_2)$ is free abelian on two generators, $b-c$ and $c-d$, expressing the geometric fact that by filling in the 2-cell A we have reduced the number of 'holes' in our space from three to two.

Example 3. Suppose we build a space X_3 by enlarging X_2 further by attaching a second 2-cell B along the same cycle $a-b$. This gives a 2-dimensional chain group C_2 consisting of linear combinations of A and B, and the boundary homomorphism $\partial_2 : C_2 \to C_1$ sends both A and B to $a-b$. The homology group $H_1(X_3) = Ker(\partial_1)/Im(\partial_2)$ is the same as for X_2, but now ∂_2 has a nontrivial kernel, the infinite cyclic group generated by $A-B$. We view $A-B$ as a 2-dimensional cycle, generating the homology group $H_2(X_3) = Ker(\partial_2) \approx \mathbb{Z}$. Topologically, the cycle $A-B$ is the sphere formed by the cells A and B together with their common boundary circle. This spherical cycle detects the presence of a 'hole' in X_3, the missing interior of the sphere. However, being enclosed by a sphere rather than a circle, this hole is of a different sort from those detected by $H_1(X_3) \approx \mathbb{Z} \times \mathbb{Z}$, which are detected by the cycles $b-c$ and $c-d$.

Example 4. Let us continue one more step and construct a complex X_4 from X_3 by attaching a 3-cell C along the 2-sphere formed by A and B. This creates a chain group C_3 generated by C, and we define a boundary homomorphism $\partial_3 : C_3 \to C_2$ sending C to $A-B$ since the cycle $A-B$ should become a boundary in the same way that the 1-dimensional cycle $a-b$ became the boundary of A. The result of this is that the quotient $H_2(X_4) = Ker(\partial_2)/Im(\partial_3)$ becomes trivial — all 2-dimensional cycles in X_4 are boundaries. Also $H_3(X_4) = Ker(\partial_3) = 0$, so $H_1(X_4) \approx \mathbb{Z} \times \mathbb{Z}$ is the only nontrivial homology group of X_4.

It is clear what the general pattern of the examples is. For a cell complex X one has chain groups $C_n(X)$ which are free abelian groups with basis the n-cells of X, and there are boundary homomorphisms $\partial_n : C_n(X) \to C_{n-1}(X)$, in terms of which one defines the homology group $H_n(X) = Ker(\partial_n)/Im(\partial_{n+1})$. The major difficulty is how to define ∂_n in general. For $n = 1$ this is easy: the boundary of an oriented edge is the vertex at its head minus the vertex at its tail. The next case $n = 2$ is also not hard, at least for cells attached along cycles which are simply loops of edges: then the boundary of the cell is this cycle of edges, with the appropriate signs taking orientations into account. But for larger n, matters become more complicated. Even if one restricts attention to cell complexes formed from polyhedral cells with nice attaching maps, there is still the matter of orientations to sort out.

The best solution to this problem seems to be to adopt an indirect approach. Arbitrary polyhedra can always be subdivided into special polyhedra called simplices (the triangle and the tetrahedron are the 2-and 3-dimensional instances) so there is no loss of generality, though initially there is some loss of efficiency, in restricting attention entirely to simplices. For simplices there is no difficulty in defining boundary maps or in handling orientations. So one obtains a homology theory, called simplicial homology, for cell complexes built from simplices. Still, this is a rather restricted class of spaces, and the theory itself has a certain rigidity which makes it awkward to work with.

The way around these obstacles is to step back from the geometry of spaces decomposed into simplices, and just to consider all possible continuous maps of simplices into a given space X. This produces tremendously large chain groups $C_n(X)$, but the quotients $H_n(X) = Ker(\partial_n)/Im(\partial_{n+1})$, called singular homology groups, turn out to be much smaller, at least for reasonably nice spaces X. In particular, for spaces like those in the four examples above, the singular homology groups coincide with the homology groups we computed from the cellular chains. And as we shall see later in this chapter, singular homology allows one to define these nice cellular homology groups for all cell complexes, and in particular to solve the problem of defining the boundary maps for cellular chains.

1. Simplicial and Singular Homology

The most important homology theory in algebraic topology, and the one we shall be studying almost exclusively, is so-called singular homology. But before starting this we will first introduce a more primitive version, called simplicial homology, to give some idea of how the technical apparatus works in a smaller-scale setting. The natural domain of definition for simplicial homology is a class of spaces we call Δ-complexes, which are a modest generalization of the more classical notion of a simplicial complex.

Δ-Complexes

We begin with three examples: the torus, the projective plane, and the Klein bottle. Each of these surfaces can be obtained from a square by identifying opposite edges, in the way indicated by the arrows in the figure below.

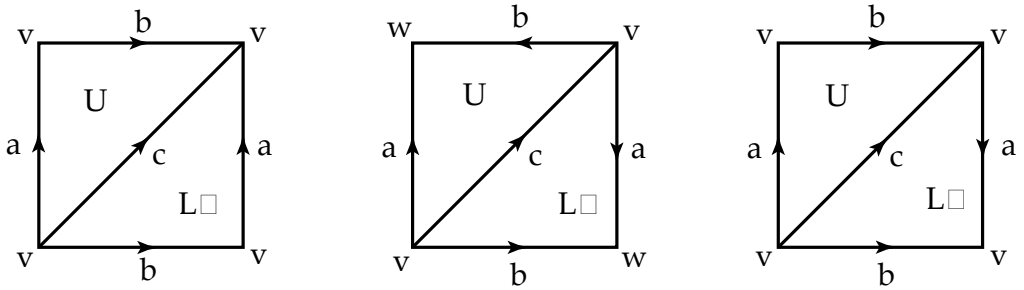

If we cut the square along a diagonal, we get two triangles, so each of these surfaces can also be constructed from two triangles by identifying certain pairs of edges.

In similar fashion, a polygon with any number of sides can be cut along diagonals into triangles, so in fact all closed surfaces can be constructed from triangles by identifying edges. Thus we have a single building block, the triangle, from which all surfaces can be constructed. Using only triangles we could also construct a large class of 2-dimensional spaces which are not surfaces by identifying more than two edges together at a time.

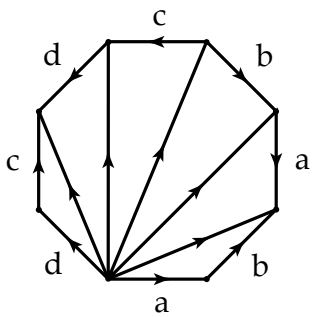

Δ-complexes are a generalization of this idea, using the n-dimensional analog of the triangle, the **n-simplex**. This is the smallest convex set in $\mathbb{R}^m$ containing $n+1$ points $v_0, \cdots, v_n$ which do not lie in a hyperplane of dimension less than n, where by a 'hyperplane' we mean the set of solutions of a system of inhomogeneous linear equations.

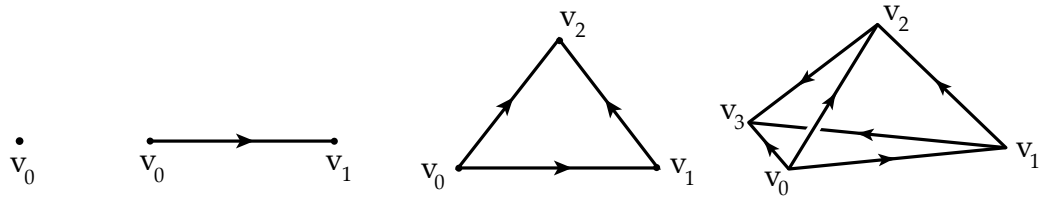

The points v_i are the **vertices** of the simplex, and the simplex itself will be denoted $[v_0, \cdots, v_n]$. For example, there is the standard n-simplex

$$\Delta^n = \left\{ (t_0, \cdots, t_n) \in \mathbb{R}^{n+1} \;\middle|\; \textstyle\sum_i t_i = 1 \text{ and } t_i \geq 0 \text{ for all } i \right\}$$

whose vertices are the unit vectors along the coordinate axes.

For purposes of homology it will be important to keep track of the order of the vertices of a simplex, so 'n-simplex' will really mean 'n-simplex with an ordering of its vertices.' A by-product of ordering the vertices of a simplex $[v_0, \cdots, v_n]$ is that this determines orientations of the edges $[v_i, v_j]$ according to increasing subscripts, as shown in the previous figure. Specifying the ordering of the vertices also determines a canonical linear homeomorphism from the standard n-simplex Δ^n onto any other n-simplex $[v_0, \cdots, v_n]$, preserving the order of vertices, namely, $(t_0, \cdots, t_n) \mapsto \sum_i t_i v_i$. The coefficients t_i are the **barycentric coordinates** of the point $\sum_i t_i v_i$ in $[v_0, \cdots, v_n]$.

A **face** of a simplex $[v_0, \cdots, v_n]$ is the subsimplex with vertices any non-empty subset of the v_i's. The subset need not be a proper subset, so $[v_0, \cdots, v_n]$ is regarded as a face of itself. The vertices of a face will always be ordered according to their order in the larger simplex.

A **Δ-complex** is a quotient space of a collection of disjoint simplices of various dimensions, obtained by identifying certain of their faces via the canonical linear homeomorphisms which preserve the ordering of vertices.

For example, the previously-described representations of the torus, projective plane, and Klein bottle as pairs of triangles with edges identified are in fact Δ-complex structures, because the indicated orientations of the three edges of each triangle are compatible with a unique ordering of the vertices of the triangle, and the identifications of edges preserve orientations.

In general, the edges in any Δ-complex X inherit well-defined orientations from the orderings of the vertices of the simplices from which X is built. These orientations are not completely arbitrary, since the orientations of the various edges in the boundary of each n-simplex of X must be related just as they are in a simplex $[v_0, \cdots, v_n]$, consistent with the ordering of the vertices. It is not hard to check that this compatibility condition on

orientations comes down to requiring that no 2-simplex of X has its three edges oriented cyclically, i.e., according to a clockwise or counterclockwise traversal of the three edges.

In the case of one-dimensional Δ-complexes this compatibility condition on orientations of edges is vacuous, so a one-dimensional Δ-complex is exactly the same thing as an oriented graph, that is, a graph with orientations specified for all its edges.

In the torus and Klein bottle examples, all three vertices of each triangle end up being identified to a single point, and in the projective plane example, two of the three vertices of each triangle are identified. Thus, certain identifications of points in the boundary of a single simplex are allowed in Δ-complexes. As a more extreme example, we could construct a Δ-complex from a 2-simplex by identifying all three of its edges together, preserving the orientations of these edges. With quite a bit of stretching and bending of the 2-simplex, these identifications can actually be performed in $\mathbb{R}^3$.

Since the identifications of faces of simplices which produce a Δ-complex X always preserve the orderings of vertices, these identifications never result in two distinct points in the interior of a face being identified in X. This means that X, as a set, is the disjoint union of a collection of *open simplices* — simplices with all their proper faces deleted. Moreover, each such open simplex e^n_α of dimension n comes equipped with a canonical map $\sigma_\alpha : \Delta^n \to X$ restricting to a homeomorphism from the interior of Δ^n onto e^n_α. Namely, the closure of e^n_α is the quotient of one of the simplices from which X was constructed, or a face of one of these simplices, and σ_α is the quotient map from this simplex or face to X. Note that the restriction of σ_α to any $(n-1)$-dimensional face of Δ^n is the canonical map σ_β for an open $(n-1)$-simplex e^{n-1}_β of X. Implicit here is the canonical identification of each $(n-1)$-dimensional face of Δ^n with the standard $(n-1)$-simplex Δ^{n-1}, preserving the order of vertices.

In the Appendix we show that the open simplices e^n_α are the cells of a CW structure with the σ_α's as characteristic maps. However, we won't need this fact in what follows.

Simplicial Homology

Our goal now is to define the simplicial homology groups of a Δ-complex X. Let $\Delta_n(X)$ be the free abelian group with basis the open n-simplices e^n_α of X. Elements of $\Delta_n(X)$, called **n-chains**, can be written as finite formal sums $\sum_\alpha n_\alpha e^n_\alpha$ or equivalently $\sum_\alpha n_\alpha \sigma_\alpha$ where $\sigma_\alpha : \Delta^n \to X$ is the canonical map with image the closure of e^n_α, described above. Such a sum $\sum_\alpha n_\alpha \sigma_\alpha$ can be thought of as a finite collection, or 'chain,' of n-simplices in X with integer multiplicities, the coefficients n_α.

As one can see in the next figure, the boundary of the n-simplex $[v_0, \cdots, v_n]$ consists of the various $(n-1)$-dimensional simplices $[v_0, \cdots, \widehat{v}_i, \cdots, v_n]$, where the 'hat' symbol

$\widehat{\ }$ over v_i indicates that this vertex is deleted from the sequence $v_0, \cdots, v_n$. We define a boundary homomorphism $\partial_n : \Delta_n(X) \to \Delta_{n-1}(X)$ by specifying its values on basis elements:

$$\partial_n(\sigma_\alpha) = \sum_i (-1)^i \sigma_\alpha \big| [v_0, \cdots, \widehat{v}_i, \cdots, v_n]$$

Thus $\partial_n(\sigma_\alpha)$ consists of the various $(n-1)$-dimensional faces of e^n_α, with certain signs. Heuristically, the signs are inserted to take orientations into account, so that all the faces of a simplex are coherently oriented, as in the following figure.

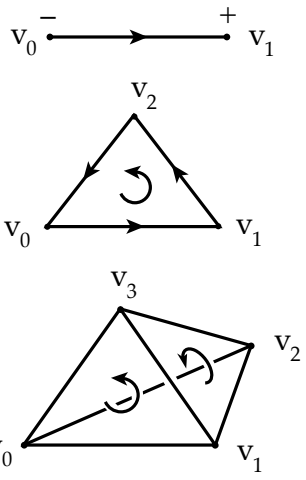

$$\partial[v_0, v_1] = [v_1] - [v_0]$$

$$\partial[v_0, v_1, v_2] = [v_1, v_2] - [v_0, v_2] + [v_0, v_1]$$

$$\partial[v_0, v_1, v_2, v_3] = [v_1, v_2, v_3] - [v_0, v_2, v_3] + [v_0, v_1, v_3] - [v_0, v_1, v_2]$$

In the last case, the orientations of the two hidden faces are also counterclockwise when viewed from outside the 3-simplex.

Lemma 2.1. *The composition* $\Delta_n(X) \xrightarrow{\partial_n} \Delta_{n-1}(X) \xrightarrow{\partial_{n-1}} \Delta_{n-2}(X)$ *is zero.*

Proof: $\partial_n(\sigma) = \sum_i (-1)^i \sigma \big| [v_0, \cdots, \widehat{v}_i, \cdots, v_n]$, hence

$$\partial_{n-1}\partial_n(\sigma) = \sum_{j<i} (-1)^i (-1)^j \sigma \big| [v_0, \cdots, \widehat{v}_j, \cdots, \widehat{v}_i, \cdots, v_n] + \sum_{j>i} (-1)^i (-1)^{j-1} \sigma \big| [v_0, \cdots, \widehat{v}_i, \cdots, \widehat{v}_j, \cdots, v_n].$$

The latter two summations cancel since after switching i and j in the second sum, it becomes the negative of the first. □

The algebraic situation we have now is a sequence of homomorphisms of abelian groups

$$\cdots \longrightarrow C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \longrightarrow \cdots \longrightarrow C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0$$

with $\partial_n\partial_{n+1} = 0$ for all n. Such a sequence is called a **chain complex**. Note that we have extended the sequence by a 0 at the right end, with $\partial_0 = 0$. From $\partial_n\partial_{n+1} = 0$ it follows that $Im(\partial_{n+1}) \subset Ker(\partial_n)$, where Im and Ker denote image and kernel. So we can define the n^{th} **homology group** of the chain complex to be the quotient group $H_n = Ker(\partial_n)/Im(\partial_{n+1})$. Elements of $Ker(\partial_n)$ are called **cycles** and elements of $Im(\partial_{n+1})$ are **boundaries**. Elements of H_n are cosets of $Im(\partial_{n+1})$, called **homology classes**. Cycles representing the same homology class are said to be **homologous**. This means their difference is a boundary.

Returning to the case that $C_n = \Delta_n(X)$, the homology group $Ker(\partial_n)/Im(\partial_{n+1})$ will be denoted $H_n^\Delta(X)$ and called the n^{th} **simplicial homology group** of X.

Example 1. $X = S^1$, with one vertex v and one edge e. Then $\Delta_0(S^1)$ and $\Delta_1(S^1)$ are both $\mathbb{Z}$ and the boundary map ∂_1 is zero since $\partial e = v - v$. The groups $\Delta_n(S^1)$ are 0 for $n \geq 2$ since there are no simplices in these dimensions. Hence

$$H_n^\Delta(S^1) \approx \begin{cases} \mathbb{Z} & \text{for } n = 0, 1 \\ 0 & \text{for } n \geq 2. \end{cases}$$

This is an illustration of the general fact that if the boundary maps in a chain complex are all zero, then the homology groups of the complex are isomorphic to the chain groups themselves.

Example 2. $X = T$, the torus with the Δ-complex structure pictured earlier, having one vertex, three edges a, b, and c, and two 2-simplices U and L. As in the previous example, $\partial_1 = 0$ so $H_0^\Delta(T) \approx \mathbb{Z}$. Since $\partial_2 U = a + b - c = \partial_2 L$ and $\{a, b, a + b - c\}$ is a basis for $\Delta_1(T)$, it follows that $H_1^\Delta(T) \approx \mathbb{Z}\oplus\mathbb{Z}$ with basis the homology classes $[a]$ and $[b]$. Since there are no 3-simplices, $H_2^\Delta(T)$ is equal to $Ker(\partial_2)$, which is infinite cyclic generated by $U - L$ since $\partial(pU + qL) = (p + q)(a + b - c) = 0$ only if $p = -q$. Thus

$$H_n^\Delta(T) \approx \begin{cases} \mathbb{Z}\oplus\mathbb{Z} & \text{for } n = 1 \\ \mathbb{Z} & \text{for } n = 0, 2 \\ 0 & \text{for } n \geq 3. \end{cases}$$

Example 3. $X = \mathbb{R}\mathrm{P}^2$, as pictured earlier, with two vertices v and w, three edges a, b, and c, and two 2-simplices U and L. Then $Im(\partial_1)$ is generated by $w - v$, so $H_0^\Delta(X) \approx \mathbb{Z}$ with either vertex as a generator. Since $\partial_2 U = -a + b + c$ and $\partial_2 L = a - b + c$, we see that ∂_2 is injective, so $H_2^\Delta(X) = 0$. Further, $Ker(\partial_1) \approx \mathbb{Z}\oplus\mathbb{Z}$ with basis $a - b$ and c, and $Im(\partial_2)$ is an index-two subgroup of $Ker(\partial_1)$ since we can choose c and $a - b + c$ as a basis for $Ker(\partial_1)$ and $a - b + c$ and $2c = (a - b + c) + (-a + b + c)$ as a basis for $Im(\partial_2)$. Thus $H_1^\Delta(X) \approx \mathbb{Z}_2$.

Example 4. We can obtain a Δ-complex structure on S^n by taking two copies of Δ^n and identifying their boundaries via the identity map. Labeling these two n-simplices U and L, then it is obvious that $Ker(\partial_n)$ is infinite cyclic generated by $U - L$. Thus $H_n^\Delta(S^n) \approx \mathbb{Z}$ in this Δ-complex structure on S^n. Computing the other homology groups would be more difficult.

Many similar examples could be worked out without much trouble, such as the other closed orientable and non-orientable surfaces. However, the calculations do tend to increase in complexity before long, particularly for higher dimensional complexes.

Some obvious general questions arise: Are the groups $H_n^\Delta(X)$ independent of the choice of Δ-complex structure on X, or in other words, do homeomorphic Δ-complexes have isomorphic homology groups? More generally, do homotopy equivalent Δ-complexes have isomorphic homology groups? To answer such questions and to develop a general theory it is best to leave the rather rigid simplicial realm and introduce the singular homology groups. These have the added advantage that they are defined for all spaces, not just Δ-complexes. At the end of this section, after some theory has been developed, we will show that simplicial and singular homology groups coincide for Δ-complexes.

Traditionally, simplicial homology is defined only for a special class of Δ-complexes called **simplicial complexes**, which are the Δ-complexes whose simplices are uniquely determined by their vertices. This amounts to saying that each n-simplex has n+1 distinct vertices, and no other n-simplex has this same set of vertices. Thus a simplicial complex can be described purely combinatorially as a set X_0 of vertices together with sets X_n of n-simplices, which are $(n+1)$-element subsets of X_0. The only requirement is that each $(k+1)$-element subset of the vertices of an n-simplex in X_n is a k-simplex, in X_k. From this combinatorial data a Δ-complex X can be constructed, once we choose a partial ordering of the vertices X_0 which restricts to a linear ordering on the vertices of each simplex in X_n. An exercise at the end of this section is to show that every Δ-complex can be subdivided to be a simplicial complex. In particular, every Δ-complex is then homeomorphic to a simplicial complex.

Compared with simplicial complexes, Δ-complexes have the advantage of simpler computations since fewer simplices are required. For example, to put a simplicial complex structure on the torus one needs at least 14 triangles, 21 edges, and 7 vertices, and for $\mathbb{R}\mathrm{P}^2$ one needs at least 10 triangles, 15 edges, and 6 vertices. This would slow down calculations considerably!

Singular Homology

A **singular n-simplex** in a space X is simply a map $\sigma : \Delta^n \to X$. The word 'singular' is used here to express the idea that σ need not be a nice embedding but can have 'singularities' where its image does not look at all like a simplex. All that is required is that σ be continuous. Let $C_n(X)$ be the free abelian group with basis the set of singular n-simplices in X. Elements of $C_n(X)$, called **n-chains**, or more precisely singular n-chains, are finite formal sums $\sum_i n_i\sigma_i$ for $n_i \in \mathbb{Z}$ and $\sigma_i : \Delta^n \to X$. A boundary map $\partial_n : C_n(X) \to C_{n-1}(X)$ is defined by the same formula as before:

$$\partial_n(\sigma) = \sum_i (-1)^i \sigma \big| [v_0, \cdots, \widehat{v}_i, \cdots, v_n]$$

Implicit in this formula is the canonical identification of $[v_0, \cdots, \widehat{v}_i, \cdots, v_n]$ with Δ^{n-1}, preserving the ordering of vertices, so that $\sigma \big| [v_0, \cdots, \widehat{v}_i, \cdots, v_n]$ is regarded as a map $\Delta^{n-1} \to X$, that is, a singular $(n-1)$-simplex.

Often we write the boundary map ∂_n from $C_n(X)$ to $C_{n-1}(X)$ simply as ∂ when this doesn't lead to serious ambiguities. The proof of Lemma 2.1 shows that $\partial_n\partial_{n+1} = 0$, or more concisely $\partial^2 = 0$, so we can define the **singular homology group** $H_n(X) = Ker(\partial_n)/Im(\partial_{n+1})$.

It is evident from the definition that homeomorphic spaces have isomorphic singular homology groups H_n, in contrast with the situation for H_n^Δ. On the other hand, since the groups $C_n(X)$ are so large, the number of singular n-simplices in X usually being uncountable, it is not at all clear that for a Δ-complex X with finitely many simplices, $H_n(X)$ should be finitely generated for all n, or that $H_n(X)$ should be zero for n larger than the dimension of X — two properties which are trivial for $H_n^\Delta(X)$.

Though singular homology looks so much more general than simplicial homology, it can actually be regarded as a special case of simplicial homology. Namely, for an arbitrary space X, we may define the **singular complex** $S(X)$ to be the Δ-complex with one n-simplex Δ^n_σ for each singular n-simplex $\sigma : \Delta^n \to X$, such an n-simplex Δ^n_σ being attached in the obvious way to the $(n-1)$-simplices $\Delta^{n-1}_{\sigma_i}$ where $\sigma_0, \cdots, \sigma_n$ are the restrictions of σ to the $(n-1)$-simplices in $\partial\Delta^n$. It is clear from the definitions that $H_n^\Delta(S(X))$ is identical with $H_n(X)$ for all n, and in this sense the singular homology group $H_n(X)$ is a special case of a simplicial homology group. One can regard $S(X)$ as a Δ-complex 'approximation' to X, an approximation which is very good from the viewpoint of homology theory. Geometrically, however, $S(X)$ doesn't look very much like X since it is so much larger than X.

Singular homology can be related to manifolds in the following way. A singular n-chain ξ can always be written in the form $\sum_i \varepsilon_i \sigma_i$ with $\varepsilon_i = \pm 1$, for singular n-simplices σ_i. When $\partial\xi$ is computed as a sum of singular $(n-1)$-simplices with signs ± 1, there may be some *cancelling pairs* consisting of two identical singular $(n-1)$-simplices with opposite signs. Choosing a maximal collection of such cancelling pairs, construct an n-dimensional Δ-complex K_ξ from a disjoint union of n-simplices Δ_i^n, one for each σ_i, by identifying the pairs of $(n-1)$-dimensional faces corresponding to the chosen cancelling pairs. The σ_i's then induce a map $K_\xi \to X$. If ξ is a cycle, all the $(n-1)$-simplices of K_ξ come from cancelling pairs, hence are faces of exactly two n-simplices of K_ξ. Thus K_ξ is a manifold, locally homeomorphic to $\mathbb{R}^n$, except at a subcomplex of dimension at most $n-2$. (A closer inspection shows that K_ξ is also a manifold near points in the interiors of $(n-2)$-simplices, so the nonmanifold points of K_ξ have dimension at most $n-3$ in fact.) Further, all the n-simplices of K_ξ can be coherently oriented by taking the signs of the σ_i's into account, so K_ξ is an oriented manifold away from its nonmanifold points; see §3.3 for a detailed development of the idea of orientability for manifolds. In particular, elements of $H_1(X)$ are represented by collections of oriented loops in X, and elements of $H_2(X)$ are represented by maps of closed oriented surfaces into X. In a similar fashion, one can see that an oriented 1-cycle $\coprod_\alpha S^1_\alpha \to X$ is zero in $H_1(X)$ iff it extends to a map of an oriented surface into X.

After these preliminary remarks let us begin to see what can be proved about singular homology.

Proposition 2.2. *If the path-components of X are denoted X_α, then $H_n(X)$ splits as the direct sum $\bigoplus_\alpha H_n(X_\alpha)$.*

Proof: Since a singular simplex always has path-connected image, $C_n(X)$ splits as the direct sum of its subgroups $C_n(X_\alpha)$. The boundary maps ∂_n preserve this direct sum decomposition, taking $C_n(X_\alpha)$ to $C_{n-1}(X_\alpha)$, so $Ker(\partial_n)$ and $Im(\partial_{n+1})$ split similarly as direct sums, hence the homology groups also split, $H_n(X) \approx \bigoplus_\alpha H_n(X_\alpha)$. □

Proposition 2.3. *If X is non-empty and path-connected then $H_0(X) \approx \mathbb{Z}$. Hence for any space X, $H_0(X)$ is a direct sum of $\mathbb{Z}$'s, one for each path-component of X.*

Proof: By definition, $H_0(X) = C_0(X)/Im(\partial_1)$ since $\partial_0 = 0$. Define a homomorphism $\varepsilon : C_0(X) \to \mathbb{Z}$ by $\varepsilon(\sum_i n_i\sigma_i) = \sum_i n_i$. This is obviously surjective if X is non-empty. The claim is that $Ker(\varepsilon) = Im(\partial_1)$ if X is path-connected, and hence ε induces an isomorphism $H_0(X) \approx \mathbb{Z}$.

To verify the claim, observe first that $Im(\partial_1) \subset Ker(\varepsilon)$ since for a singular 1-simplex $\sigma : \Delta^1 \to X$, $\varepsilon\partial_1(\sigma) = \varepsilon(\sigma\big|[v_1] - \sigma\big|[v_0]) = 1 - 1 = 0$. For the reverse inclusion $Ker(\varepsilon) \subset Im(\partial_1)$, suppose $\varepsilon(\sum_i n_i\sigma_i) = 0$, i.e., $\sum_i n_i = 0$. Here the σ_i's are singular 0-simplices, i.e., points of X. Choose a path τ_i from a basepoint x_0 to $\sigma_i(v_0)$ and let σ_0 be the singular 0-simplex with image x_0. Regarding the τ_i's as singular 1-simplices, we have $\partial(\sum_i n_i\tau_i) = \sum_i n_i\sigma_i - \sum_i n_i\sigma_0 = \sum_i n_i\sigma_i$ since $\sum_i n_i = 0$. Thus $\sum_i n_i\sigma_i$ is a boundary, which shows that $Ker(\varepsilon) \subset Im(\partial_1)$. □

Proposition 2.4. *If X is a point, then $H_n(X) = 0$ for $n > 0$ and $H_0(X) \approx \mathbb{Z}$.*

Proof: In this case there is a unique singular n-simplex σ_n for each n, and $\partial(\sigma_n) = \sum_i (-1)^i \sigma_{n-1}$, a sum of $n+1$ terms, which is therefore 0 for n odd and σ_{n-1} for n even, $n \neq 0$. Thus we have the chain complex

$$\cdots \longrightarrow \mathbb{Z} \xrightarrow{\approx} \mathbb{Z} \xrightarrow{0} \mathbb{Z} \xrightarrow{\approx} \mathbb{Z} \xrightarrow{0} \mathbb{Z} \longrightarrow 0$$

with boundary maps alternately isomorphisms and trivial maps, except at the last $\mathbb{Z}$. The homology groups of this complex are trivial except in dimension zero, where $H_0 \approx \mathbb{Z}$. □

It is often very convenient to have a slightly modified version of homology for which a point has trivial homology groups in all dimensions, including zero. This is done by defining the **reduced homology groups** $\widetilde{H}_n(X)$ to be the homology groups of the augmented chain complex

$$\cdots \longrightarrow C_2(X) \xrightarrow{\partial_2} C_1(X) \xrightarrow{\partial_1} C_0(X) \xrightarrow{\varepsilon} \mathbb{Z} \longrightarrow 0$$

where $\varepsilon(\sum_i n_i\sigma_i) = \sum_i n_i$ as in the proof of Proposition 2.3. Here we had better require X to be non-empty, to avoid having a nontrivial homology group in dimension -1. Since $\varepsilon\partial_1 = 0$, ε vanishes on $Im(\partial_1)$ and hence induces a map $H_0(X) \to \mathbb{Z}$ with kernel $\widetilde{H}_0(X)$, so $H_0(X) \approx \widetilde{H}_0(X) \oplus \mathbb{Z}$. Obviously $H_n(X) \approx \widetilde{H}_n(X)$ for $n > 0$.

Homotopy Invariance

The first substantial result we will prove about singular homology is that homotopy equivalent spaces have isomorphic homology groups. This will be done by showing that a map $f : X \to Y$ induces a homomorphism $f_* : H_n(X) \to H_n(Y)$ for each n, and that f_* is an isomorphism if f is a homotopy equivalence.

For a map $f : X \to Y$, an induced homomorphism $f_\sharp : C_n(X) \to C_n(Y)$ is defined by composing each singular n-simplex $\sigma : \Delta^n \to X$ with f to get a singular n-simplex

$f_\sharp(\sigma) = f\sigma : \Delta^n \to Y$, then extending $f_\sharp$ linearly via $f_\sharp(\sum_i n_i\sigma_i) = \sum_i n_i f_\sharp(\sigma_i) = \sum_i n_i f\sigma_i$. The maps $f_\sharp : C_n(X) \to C_n(Y)$ satisfy $f_\sharp\partial = \partial f_\sharp$ since

$$f_\sharp\partial(\sigma) = f_\sharp\big(\textstyle\sum_i (-1)^i \sigma \big| [v_0, \cdots, \widehat{v}_i, \cdots, v_n]\big) = \sum_i (-1)^i f\sigma \big| [v_0, \cdots, \widehat{v}_i, \cdots, v_n] = \partial f_\sharp(\sigma).$$

Thus we have a diagram

$$\begin{array}{ccccccccc}
\cdots & \longrightarrow & C_{n+1}(X) & \xrightarrow{\partial} & C_n(X) & \xrightarrow{\partial} & C_{n-1}(X) & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle f_\sharp} & & \downarrow{\scriptstyle f_\sharp} & & \downarrow{\scriptstyle f_\sharp} & & \\
\cdots & \longrightarrow & C_{n+1}(Y) & \xrightarrow{\partial} & C_n(Y) & \xrightarrow{\partial} & C_{n-1}(Y) & \longrightarrow & \cdots
\end{array}$$

such that in each square the composition $f_\sharp\partial$ equals the composition $\partial f_\sharp$. A diagram of maps with the property that any two compositions of maps starting at one point in the diagram and ending at another are equal is called a **commutative diagram**. In the present case commutativity of the diagram is equivalent to the commutativity relation $f_\sharp\partial = \partial f_\sharp$, but commutative diagrams can contain commutative triangles, pentagons, etc., as well as commutative squares.

The fact that the maps $f_\sharp : C_n(X) \to C_n(Y)$ satisfy $f_\sharp\partial = \partial f_\sharp$ is also expressed by saying that the $f_\sharp$'s define a **chain map** from the singular chain complex of X to that of Y. The relation $f_\sharp\partial = \partial f_\sharp$ implies that $f_\sharp$ takes cycles to cycles since $\partial\alpha = 0$ implies $\partial(f_\sharp\alpha) = f_\sharp(\partial\alpha) = 0$. Also, $f_\sharp$ takes boundaries to boundaries since $f_\sharp(\partial\beta) = \partial(f_\sharp\beta)$. Hence $f_\sharp$ induces a homomorphism $f_* : H_n(X) \to H_n(Y)$. An algebraic statement of what we have just proved is:

Proposition 2.5. *A chain map between chain complexes induces homomorphisms between the homology groups of the two complexes.* □

Two basic properties of induced homomorphisms which are important in spite of being rather trivial are:

(*i*) $(fg)_* = f_* g_*$ for a composed mapping $X \xrightarrow{g} Y \xrightarrow{f} Z$. This follows from associativity of compositions $\Delta^n \xrightarrow{\sigma} X \xrightarrow{g} Y \xrightarrow{f} Z$.

(*ii*) $\mathbb{1}_* = \mathbb{1}$ where $\mathbb{1}$ denotes the identity map of a space or a group.

Less trivially, we have:

Theorem 2.6. *Homotopic maps $f, g : X \to Y$ induce the same homomorphism $f_* = g_* : H_n(X) \to H_n(Y)$.*

In view of the formal properties $(fg)_* = f_* g_*$ and $\mathbb{1}_* = \mathbb{1}$, this immediately implies:

Corollary 2.7. *A homotopy equivalence* $f : X \to Y$ *induces isomorphisms* $f_* : H_n(X) \to H_n(Y)$ *for all* n. □

For example, if X is contractible then $\widetilde{H}_n(X) = 0$ for all n.

Proof of 2.6: The essential ingredient is a procedure for subdividing the product $\Delta^n \times I$ into $(n+1)$-simplices, as illustrated below in the cases $n = 1, 2$.

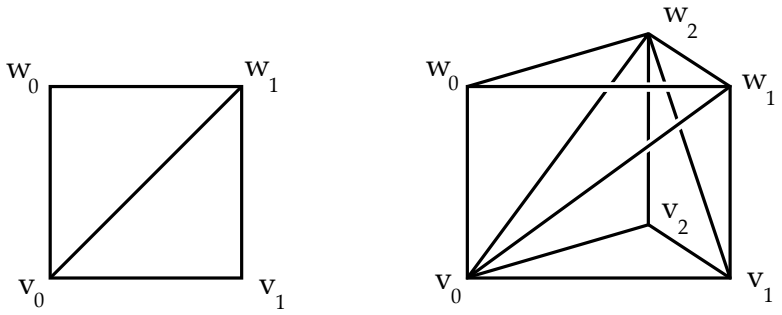

In $\Delta^n \times I$, let $\Delta^n \times \{0\} = [v_0, \cdots, v_n]$ and $\Delta^n \times \{1\} = [w_0, \cdots, w_n]$, where v_i and w_i have the same image under the projection $\Delta^n \times I \to \Delta^n$. Then the n-simplex $[v_0, \cdots, v_i, w_{i+1}, \cdots, w_n]$ is the graph of the function $\varphi_i : \Delta^n \to I$ given by $\varphi_i(t_0, \cdots, t_n) = t_{i+1} + \cdots + t_n$ since its vertices $v_0, \cdots, v_i, w_{i+1}, \cdots, w_n$ are on the graph of φ_i and it projects homeomorphically onto Δ^n under the projection $\Delta^n \times I \to \Delta^n$. The region between the graphs of φ_i and φ_{i-1} is $[v_0, \cdots, v_i, w_i, \cdots, w_n]$, which is a genuine $(n+1)$-simplex since w_i is not on the graph of φ_i, hence is not in the n-simplex $[v_0, \cdots, v_i, w_{i+1}, \cdots, w_n]$. The inequalities $0 = \varphi_n \le \varphi_{n-1} \le \cdots \le \varphi_0 \le \varphi_{-1} = 1$ show that $\Delta^n \times I$ is the union of the $(n+1)$-simplices $[v_0, \cdots, v_i, w_i, \cdots, w_n]$, each intersecting the next in an n-simplex face.

Given a homotopy $F : X \times I \to Y$ from f to g, define *prism* operators $P : C_n(X) \to C_{n+1}(Y)$ by

$$P(\sigma) = \sum_i (-1)^i F \circ (\sigma \times 1\!\!1) \,\big|\, [v_0, \cdots, v_i, w_i, \cdots, w_n]$$

for $\sigma : \Delta^n \to X$, where $F \circ (\sigma \times 1\!\!1)$ is the composition $\Delta^n \times I \to X \times I \to Y$. We shall show that these prism homomorphisms P satisfy the basic relation

$$\partial P = g_\sharp - f_\sharp - P\partial$$

This is an algebraic version of the statement that $\partial(\Delta^n \times I)$ consists of the top $\Delta^n \times \{1\}$,

the bottom $\Delta^n \times \{0\}$, and the sides $\partial\Delta^n \times I$. To prove the relation we calculate

$$\begin{aligned}\partial P(\sigma) = &\sum_{j\le i}(-1)^i(-1)^j F\circ(\sigma\times\mathbb{1})\big|[v_0,\cdots,\widehat{v}_j,\cdots,v_i,w_i,\cdots,w_n]\\ &+\sum_{j\ge i}(-1)^i(-1)^{j+1}F\circ(\sigma\times\mathbb{1})\big|[v_0,\cdots,v_i,w_i,\cdots,\widehat{w}_j,\cdots,w_n]\end{aligned}$$

The terms with $i=j$ in the two sums cancel except for $F\circ(\sigma\times\mathbb{1})\big|[\widehat{v}_0,w_0,\cdots,w_n]$ which is $g\circ\sigma = g_\sharp(\sigma)$, and $-F\circ(\sigma\times\mathbb{1})\big|[v_0,\cdots,v_n,\widehat{w}_n]$ which is $-f\circ\sigma = -f_\sharp(\sigma)$. The terms with $i\neq j$ are exactly $-P\partial(\sigma)$ since

$$\begin{aligned}P\partial(\sigma) = &\sum_{i<j}(-1)^i(-1)^j F\circ(\sigma\times\mathbb{1})\big|[v_0,\cdots,v_i,w_i,\cdots,\widehat{w}_j,\cdots,w_n]\\ &+\sum_{i>j}(-1)^{i-1}(-1)^j F\circ(\sigma\times\mathbb{1})\big|[v_0,\cdots,\widehat{v}_j,\cdots,v_i,w_i,\cdots,w_n]\end{aligned}$$

Now we can finish the proof of the theorem. If $\alpha\in C_n(X)$ is a cycle, then $g_\sharp(\alpha)-f_\sharp(\alpha) = \partial P(\alpha)+P\partial(\alpha) = \partial P(\alpha)$ since $\partial\alpha = 0$. Thus $g_\sharp(\alpha)$ and $f_\sharp(\alpha)$ are homologous, so $f_*(\alpha) = g_*(\alpha)$. □

The relationship $\partial P + P\partial = g_\sharp - f_\sharp$ is expressed by saying P is a **chain homotopy** between the chain maps $f_\sharp$ and $g_\sharp$. We have just shown:

Proposition 2.8. *Chain-homotopic chain maps induce the same homomorphism on homology.* □

For reduced homology groups there are also induced homomorphisms $f_*:\widetilde{H}_n(X)\to\widetilde{H}_n(Y)$ since $f_\sharp\varepsilon = \varepsilon f_\sharp$. These induced homomorphisms satisfy the same general properties.

Exact Sequences and Excision

It is time to introduce the most important algebraic tool in homology theory, exact sequences. A sequence of homomorphisms

$$\cdots\longrightarrow A_{n+1}\xrightarrow{\alpha_{n+1}}A_n\xrightarrow{\alpha_n}A_{n-1}\longrightarrow\cdots$$

is said to be **exact** if $Ker(\alpha_n) = Im(\alpha_{n+1})$ for each n. The inclusions $Im(\alpha_{n+1})\subset Ker(\alpha_n)$ are equivalent to $\alpha_n\alpha_{n+1}=0$, so the sequence is a chain complex, and the opposite inclusions $Ker(\alpha_n)\subset Im(\alpha_{n+1})$ say that the homology groups of this chain complex are trivial.

A number of basic algebraic concepts can be expressed in terms of exact sequences, for example:

(*i*) $0 \longrightarrow A \xrightarrow{\alpha} B$ is exact iff $Ker(\alpha) = 0$ i.e., α is injective.

(*ii*) $A \xrightarrow{\alpha} B \longrightarrow 0$ is exact iff $Im(\alpha) = B$ i.e., α is surjective.

(*iii*) $0 \longrightarrow A \xrightarrow{\alpha} B \longrightarrow 0$ is exact iff α is an isomorphism, by (i) and (ii).

(*iv*) $0 \longrightarrow A \xrightarrow{\alpha} B \xrightarrow{\beta} C \longrightarrow 0$ is exact iff α is injective, β is surjective, and $Ker(\beta) = Im(\alpha)$, so β induces an isomorphism $C \approx B/Im(\alpha)$. This can be written $C \approx B/A$ if we think of α as an inclusion of A as a subgroup of B.

An exact sequence $0 \to A \to B \to C \to 0$ as in (iv) is called a **short exact sequence**.

The second major tool for calculating homology, after homotopy invariance, is a long exact sequence of homology groups associated to certain pairs of spaces (X, A), pairs we call **good pairs** consisting of a space X and a non-empty closed subspace A which is a deformation retract of some open neighborhood in X. For example, if X is a CW complex and A is a non-empty subcomplex, then (X, A) is a good pair by Proposition A.4 in the Appendix.

Theorem 2.9. *If (X, A) is a good pair, then there is an exact sequence*

$$\cdots \longrightarrow \widetilde{H}_n(A) \xrightarrow{i_*} \widetilde{H}_n(X) \xrightarrow{j_*} \widetilde{H}_n(X/A) \xrightarrow{\partial} \widetilde{H}_{n-1}(A) \xrightarrow{i_*} \widetilde{H}_{n-1}(X) \longrightarrow \cdots$$
$$\cdots \longrightarrow \widetilde{H}_0(X/A) \longrightarrow 0$$

where i is the inclusion $A \hookrightarrow X$ and j is the quotient map $X \to X/A$.

The map ∂ will be constructed in the course of the proof. The idea is that an element $x \in \widetilde{H}_n(X/A)$ can be represented by a chain α in X with $\partial\alpha$ a cycle in A whose homology class is $\partial x \in \widetilde{H}_{n-1}(A)$.

Corollary 2.10. *$\widetilde{H}_n(S^n) \approx \mathbb{Z}$ and $\widetilde{H}_i(S^n) = 0$ for $i \neq n$.*

Proof: Take $(X, A) = (D^n, S^{n-1})$ so $X/A = S^n$. The long exact sequence of homology groups for the pair (X, A) has every third term $\widetilde{H}_i(D^n)$ zero since D^n is contractible. Exactness of the sequence then implies that the maps $\widetilde{H}_i(S^n) \xrightarrow{\partial} \widetilde{H}_{i-1}(S^{n-1})$ are isomorphims for all i and n. The result now follows by induction on n, starting with the case of S^0, where the result holds by Propositions 2.2 and 2.4. □

As a corollary of this calculation we have the following classical theorem of Brouwer, the 2-dimensional case of which was proved in §1.1.

Corollary 2.11. *∂D^n is not a retract of D^n. Hence every map $f: D^n \to D^n$ has a fixed point.*

Proof: If $r: D^n \to \partial D^n$ were a retraction, we would have $ri = \mathbb{1}$ for $i: \partial D^n \to D^n$ the inclusion map, hence the composition

$$\widetilde{H}_{n-1}(\partial D^n) \xrightarrow{i_*} \widetilde{H}_{n-1}(D^n) \xrightarrow{r_*} \widetilde{H}_{n-1}(\partial D^n)$$

would be the identity map on $\widetilde{H}_{n-1}(\partial D^n) \approx \mathbb{Z}$. But r_* and i_* are both 0 since $\widetilde{H}_{n-1}(D^n) = 0$, and we have a contradiction. The fixed-point result follows as in Theorem 1.10. □

Now we turn to the derivation of the long exact sequence of homology groups in the preceding theorem. There are four main steps:

(1) Define *relative* homology groups $H_n(X, A)$ for any pair (X, A) and show they fit into a long exact sequence

$$\cdots \to H_n(A) \to H_n(X) \to H_n(X, A) \to H_{n-1}(A) \to \cdots \to H_0(X, A) \to 0$$

There is also an analogous long exact sequence of reduced homology groups, though it turns out that $\widetilde{H}_n(X, A) = H_n(X, A)$ if $A \neq \varnothing$. Relative homology groups are quite useful tools in themselves, enhancing the flexibility of singular homology considerably.

(2) Develop a technical tool called barycentric subdivision which allows homology groups to be computed using arbitrarily small singular simplices.

(3) Apply this tool to prove a fundamental *excision* property of relative homology groups: $H_n(X, A) \approx H_n(X - Z, A - Z)$ for suitable subsets $Z \subset A$.

(4) Apply the excision property to show that $H_n(X, A) \approx \widetilde{H}_n(X/A)$ for good pairs (X, A).

Plugging the isomorphism in (4) into the reduced version of the long exact sequence in (1) then give the exact sequence we are looking for.

(1) Relative Homology Groups

Given a space X and a subspace $A \subset X$, define $C_n(X, A)$ to be the quotient group $C_n(X)/C_n(A)$. This quotient could also be viewed as the subgroup of $C_n(X)$ with basis the singular n-simplices $\sigma: \Delta^n \to X$ whose image is not contained in A, but it is usually better to regard $C_n(X, A)$ as a quotient of $C_n(X)$. Since the boundary map $\partial: C_n(X) \to C_{n-1}(X)$ takes $C_n(A)$ to $C_{n-1}(A)$, it induces a quotient boundary map $\partial: C_n(X, A) \to$

$C_{n-1}(X, A)$. The relation $\partial^2 = 0$ holds for this since it holds for the usual boundary map on $C_n(X)$. So we have a chain complex

$$\cdots \longrightarrow C_n(X, A) \xrightarrow{\partial} C_{n-1}(X, A) \longrightarrow \cdots$$

and we can define the **relative homology group** $H_n(X, A)$ to be $Ker(\partial)/Im(\partial)$ at $C_n(X, A)$ in this chain complex.

By considering the definition of the relative boundary map we see:

— Elements of $H_n(X, A)$ are represented by **relative cycles**: n-chains $\alpha \in C_n(X)$ such that $\partial\alpha \in C_{n-1}(A)$.

— A relative cycle α is trivial in $H_n(X, A)$ iff it is a **relative boundary**: $\alpha = \partial\beta + \gamma$ for some $\beta \in C_{n+1}(X)$ and $\gamma \in C_n(A)$.

These properties make precise the intuitive idea that $H_n(X, A)$ is 'homology of X modulo A.'

Consider now the diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & C_n(A) & \xrightarrow{i} & C_n(X) & \xrightarrow{j} & C_n(X, A) & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial} & & \downarrow{\scriptstyle \partial} & & \downarrow{\scriptstyle \partial} & & \\
0 & \longrightarrow & C_{n-1}(A) & \xrightarrow{i} & C_{n-1}(X) & \xrightarrow{j} & C_{n-1}(X, A) & \longrightarrow & 0
\end{array}$$

where i is inclusion and j is the quotient map. The diagram is commutative by the definition of the boundary maps. Putting these diagrams together for varying n, we have a **short exact sequence of chain complexes**, i.e., a commutative diagram

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
\cdots & \longrightarrow & A_{n+1} & \xrightarrow{\partial} & A_n & \xrightarrow{\partial} & A_{n-1} & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle i} & & \downarrow{\scriptstyle i} & & \downarrow{\scriptstyle i} & & \\
\cdots & \longrightarrow & B_{n+1} & \xrightarrow{\partial} & B_n & \xrightarrow{\partial} & B_{n-1} & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle j} & & \downarrow{\scriptstyle j} & & \downarrow{\scriptstyle j} & & \\
\cdots & \longrightarrow & C_{n+1} & \xrightarrow{\partial} & C_n & \xrightarrow{\partial} & C_{n-1} & \longrightarrow & \cdots \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

where the rows are chain complexes and the columns are exact. We shall show that when we pass to homology groups, this short exact sequence of chain complexes stretches out into a long exact sequence of homology groups

$$\cdots \longrightarrow H_n(A) \xrightarrow{i_*} H_n(B) \xrightarrow{j_*} H_n(C) \xrightarrow{\partial} H_{n-1}(A) \xrightarrow{i_*} H_{n-1}(B) \longrightarrow \cdots$$

where $H_n(A)$ denotes $Ker(\partial)/Im(\partial)$ at A_n, etc.

The commutativity of the squares in the preceding large diagram means that i and j are chain maps, which therefore induce maps i_* and j_* on homology. To define the boundary map $\partial : H_n(C) \to H_{n-1}(A)$, let $c \in C_n$ be a cycle. Since j is onto, $c = j(b)$ for some $b \in B_n$. The element $\partial b \in B_{n-1}$ is in $Ker(j)$ since $j(\partial b) = \partial j(b) = \partial c = 0$. So $\partial b = i(a)$ for some $a \in A_{n-1}$ since $Ker(j) = Im(i)$. Note that $\partial a = 0$ since $i(\partial a) = \partial i(a) = \partial\partial b = 0$ and i is injective. We define $\partial : H_n(C) \to H_{n-1}(A)$ by sending the homology class of c to the homology class of a, $\partial[c] = [a]$. This is well-defined since a different choice of b would have the form $b + i(a')$, changing a to $a + \partial a'$ which is homologous to a ; and a different choice of c within its homology class would have the form $c + \partial c' = c + \partial j(b') = c + j(\partial b') = j(b + \partial b')$, changing b to $b + \partial b'$ and hence leaving ∂b and therefore also a unchanged. The map $\partial : H_n(C) \to H_{n-1}(A)$ is a homomorphism since if $\partial[c_1] = [a_1]$ and $\partial[c_2] = [a_2]$ via elements b_1 and b_2 as above, then $j(b_1 + b_2) = j(b_1) + j(b_2) = c_1 + c_2$ and $i(a_1 + a_2) = i(a_1) + i(a_2) = \partial b_1 + \partial b_2 = \partial(b_1 + b_2)$ so $\partial([c_1] + [c_2]) = [a_1] + [a_2]$.

Theorem 2.12. *The sequence of homology groups*

$$\cdots \longrightarrow H_n(A) \xrightarrow{i_*} H_n(B) \xrightarrow{j_*} H_n(C) \xrightarrow{\partial} H_{n-1}(A) \xrightarrow{i_*} H_{n-1}(B) \longrightarrow \cdots$$

is exact.

Proof: There are six things to verify:

$Im(i_*) \subset Ker(j_*)$. This is immediate since $ji = 0$ implies $j_* i_* = 0$.

$Im(j_*) \subset Ker(\partial)$. We have $\partial j_* = 0$ since in this case $\partial b = 0$ in the definition of ∂.

$Im(\partial) \subset Ker(i_*)$. Here $i_* \partial = 0$ since $i_* \partial$ takes $[c]$ to $[\partial b] = 0$.

$Ker(j_*) \subset Im(i_*)$. A homology class in $Ker(j_*)$ is represented by a cycle $b \in B_n$ with $j(b)$ a boundary, so $j(b) = \partial c'$ for some $c' \in C_{n+1}$. Since j is surjective, $c' = j(b')$ for some $b' \in B_{n+1}$. We have $j(b - \partial b') = j(b) - j(\partial b') = j(b) - \partial j(b') = 0$ since $\partial j(b') = \partial c' = j(b)$. So $b - \partial b' = i(a)$ for some $a \in A_n$. This a is a cycle since $i(\partial a) = \partial i(a) = \partial(b - \partial b') = \partial b = 0$ and i is injective. Thus $i_*[a] = [b - \partial b'] = [b]$, showing that i_* maps onto $Ker(j_*)$.

$Ker(\partial) \subset Im(j_*)$. In the notation used in the definition of ∂, if c represents a homology class in $Ker(\partial)$, then $a = \partial a'$ for some $a' \in A_n$. The element $b - i(a')$ is a cycle since $\partial(b - i(a')) = \partial b - \partial i(a') = \partial b - i(\partial a') = \partial b - i(a) = 0$. And $j(b - i(a')) = j(b) - ji(a') = j(b) = c$, so j_* maps $[b - i(a')]$ to $[c]$.

$Ker(i_*) \subset Im(\partial)$. Given a cycle $a \in A_{n-1}$ such that $i(a) = \partial b$ for some $b \in B_n$, then $j(b)$ is a cycle since $\partial j(b) = j(\partial b) = ji(a) = 0$, and ∂ takes $[j(b)]$ to $[a]$. □

This theorem represents the beginnings of the subject of homological algebra. The method of proof is aptly called *diagram chasing*.

Returning to topology, the preceding algebraic theorem yields a long exact sequence of homology groups:

$$\cdots \longrightarrow H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\partial} H_{n-1}(A) \xrightarrow{i_*} H_{n-1}(X) \longrightarrow \cdots$$
$$\cdots \longrightarrow H_0(X, A) \longrightarrow 0$$

The boundary map $\partial : H_n(X, A) \to H_{n-1}(A)$ has a very simple description: If a class $[\alpha] \in H_n(X, A)$ is represented by a relative cycle α, then $\partial[\alpha]$ is the class of the cycle $\partial\alpha$ in $H_{n-1}(A)$. This is immediate from the algebraic definition of the boundary homomorphism in the long exact sequence of homology groups associated to a short exact sequence of chain complexes.

This long exact sequence makes precise the idea that the groups $H_n(X, A)$ measure the difference between the groups $H_n(X)$ and $H_n(A)$. In particular, the exact sequence implies that $H_n(X, A) = 0$ for all n iff the inclusion $A \hookrightarrow X$ induces isomorphisms $H_n(A) \approx H_n(X)$ for all n.

There is a completely analogous long exact sequence of reduced homology groups. This comes from applying the preceding algebraic machinery to the short exact sequence of chain complexes formed by the short exact sequences $0 \to C_n(A) \to C_n(X) \to C_n(X, A) \to 0$ in non-negative dimensions, augmented by the short exact sequence $0 \to \mathbb{Z} \xrightarrow{\mathbb{1}} \mathbb{Z} \to 0 \to 0$ in dimension -1. Thus $\widetilde{H}_n(X, A) = H_n(X, A)$ for all n.

Example 1. When $(X, A) = (D^n, \partial D^n)$ the long exact sequence of reduced homology groups implies that $H_k(D^n, \partial D^n) \approx \widetilde{H}_{k-1}(S^{n-1})$ for all k since $\widetilde{H}_k(D^n) = 0$ for all k. Thus

$$H_k(D^n, \partial D^n) \approx \begin{cases} \mathbb{Z} & \text{for } k = n \\ 0 & \text{otherwise.} \end{cases}$$

Example 2. Applying the long exact sequence of reduced homology groups to a pair (X, x_0) with $x_0 \in X$ yields isomorphisms $H_n(X, x_0) \approx \widetilde{H}_n(X)$ for all n since $\widetilde{H}_n(x_0) = 0$ for all n.

There are induced homomorphisms for relative homology just like in the nonrelative, or 'absolute,' case. A map $f : X \to Y$ with $f(A) \subset B$, or more concisely $f : (X, A) \to (Y, B)$, induces homomorphisms $f_\sharp : C_n(X, A) \to C_n(Y, B)$ since $f_\sharp : C_n(X) \to C_n(Y)$

takes $C_n(A)$ to $C_n(B)$, so we get a well-defined map on quotients. The relation $f_\sharp\partial = \partial f_\sharp$ holds for relative chains since it holds for absolute chains. By Proposition 2.6 we then have induced homomorphisms $f_*: H_n(X,A) \to H_n(Y,B)$.

Proposition 2.13. *If $f, g: (X,A) \to (Y,B)$ are homotopic through maps $(X,A) \to (Y,B)$, then $f_* = g_*: H_n(X,A) \to H_n(Y,B)$.*

Proof: The prism operator P from the proof of Theorem 2.6 takes $C_n(A)$ to $C_{n+1}(B)$, hence induces $P: C_n(X,A) \to C_{n+1}(Y,B)$. Since we are just passing to quotient groups, the formula $\partial P + P\partial = g_\sharp - f_\sharp$ remains valid. Thus the maps $f_\sharp$ and $g_\sharp$ on relative chain groups are chain homotopic, and therefore induce the same homomorphism on relative homology groups. □

An easy generalization of the long exact sequence of a pair (X,A) is the long exact sequence of a triple (X,A,B), where $B \subset A \subset X$:

$$\cdots \longrightarrow H_n(A,B) \longrightarrow H_n(X,B) \longrightarrow H_n(X,A) \longrightarrow H_{n-1}(A,B) \longrightarrow \cdots$$

This is the long exact sequence of homology groups associated to the short exact sequence of chain complexes which in dimension n is the short exact sequence

$$0 \longrightarrow C_n(A,B) \longrightarrow C_n(X,B) \longrightarrow C_n(X,A) \longrightarrow 0$$

For example, taking B to be a point, the long exact sequence of the triple (X,A,B) becomes the long exact sequence of reduced homology for the pair (X,A).

(2) Barycentric Subdivision

Let $\mathcal{U} = \{U_j\}$ be a collection of subspaces of X whose interiors form an open cover of X, and let $C_n^{\mathcal{U}}(X)$ be the subgroup of $C_n(X)$ generated by singular simplices with image in some U_j. The boundary map $\partial: C_n(X) \to C_{n-1}(X)$ takes $C_n^{\mathcal{U}}(X)$ to $C_{n-1}^{\mathcal{U}}(X)$, so the groups $C_n^{\mathcal{U}}(X)$ form a chain complex. We denote the homology groups of this chain complex by $H_n^{\mathcal{U}}(X)$.

Proposition 2.14. *The inclusion $\iota: C_n^{\mathcal{U}}(X) \hookrightarrow C_n(X)$ is a chain homotopy equivalence, i.e., there is a chain map $\rho: C_n(X) \to C_n^{\mathcal{U}}(X)$ such that $\iota\rho$ and $\rho\iota$ are chain homotopic to the identity. Hence ι induces isomorphisms $H_n^{\mathcal{U}}(X) \approx H_n(X)$ for all n.*

Proof: First we define the notion of barycentric subdivision. The points of a simplex $[v_0, \cdots, v_n]$ are the linear combinations $\sum_i t_i v_i$ with $t_i \geq 0$ and $\sum_i t_i = 1$. The t_i's are the barycentric coordinates of the point $\sum_i t_i v_i$, and the **barycenter** or 'center of gravity'

of the simplex is the point b whose barycentric coordinates are all equal, namely $t_i = 1/(n+1)$ for each i. The **barycentric subdivision** of $[v_0, \cdots, v_n]$ is the decomposition of $[v_0, \cdots, v_n]$ into the n-simplices $[b, w_0, \cdots, w_{n-1}]$ where, inductively, $[w_0, \cdots, w_{n-1}]$ is an $(n-1)$-simplex in the barycentric subdivision of a face $[v_0, \cdots, \widehat{v}_i, \cdots, v_n]$. The induction starts with $n = 0$ when the barycentric subdivision of $[v_0]$ is defined to be just $[v_0]$ itself. The next two cases $n = 1, 2$ and part of the case $n = 3$ are shown in the figure below.

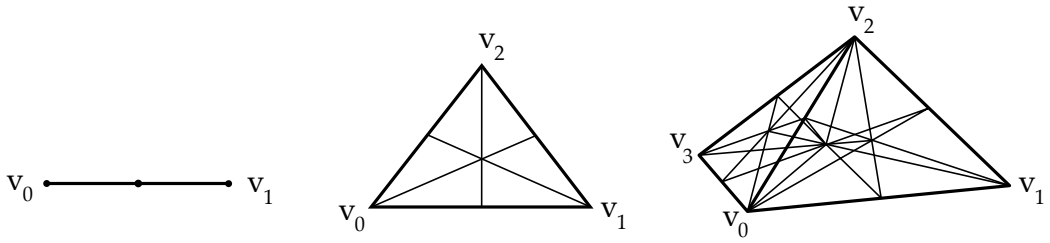

It follows from the inductive definition that the vertices of simplices in the barycentric subdivision of $[v_0, \cdots, v_n]$ are the barycenters of the various subsimplices $[v_{i_0}, \cdots, v_{i_k}]$ of $[v_0, \cdots, v_n]$. Note that the barycenter of $[v_{i_0}, \cdots, v_{i_k}]$ has barycentric coordinates $t_i = 1/(k+1)$ for $i = i_0, \cdots, i_k$ and $t_i = 0$ otherwise.

A fact we will need is that the diameter of each simplex of the barycentric subdivision of $[v_0, \cdots, v_n]$ is at most $n/(n+1)$ times the diameter of $[v_0, \cdots, v_n]$. Here the diameter of a simplex is by definition the maximum distance between any two of its points, and we are using the metric from the ambient Euclidean space $\mathbb{R}^m$ containing $[v_0, \cdots, v_n]$. The diameter of a simplex equals the maximum distance between any two of its vertices because the distance between two points v and $\sum_i t_i v_i$ of $[v_0, \cdots, v_n]$ is:

$$|v - \textstyle\sum_i t_i v_i| = |\sum_i t_i (v - v_i)| \le \sum_i t_i |v - v_i| \le \sum_i t_i \max |v - v_i| = \max |v - v_i|$$

To obtain the bound $n/(n+1)$ on the ratio of diameters it suffices by induction on n to consider the distance from the barycenter b of $[v_0, \cdots, v_n]$ to the barycenter of a subsimplex $[v_{i_0}, \cdots, v_{i_k}]$. Taking $v = b$ in the previous displayed inequalities, we see that it in fact suffices to consider the distance from b to a vertex v_i. If b_i is the barycenter of $[v_0, \cdots, \widehat{v}_i, \cdots, v_n]$, with all barycentric coordinates equal to $1/n$ except for $t_i = 0$, then

$$b = \frac{1}{n+1}\, v_i \;+\; \frac{n}{n+1}\, b_i$$

The sum of the two coefficients is 1, so b lies on the segment from v_i to b_i and the distance from b to v_i is $n/(n+1)$ times the length of this segment. Hence the distance from b to v_i is at most $n/(n+1)$ times the diameter of $[v_0, \cdots, v_n]$.

The significance of the factor $n/(n+1)$ is that by repeated barycentric subdivision we can produce simplices of arbitrarily small diameter since $(n/(n+1))^r$ approaches 0 as r goes to infinity.

We define inductively a chain $S([v_0,\cdots,v_n]) \in C_n([v_0,\cdots,v_n])$ which is a signed sum of the n-simplices in the barycentric subdivision of $[v_0,\cdots,v_n]$, namely $S([v_0]) = [v_0]$ and $S([v_0,\cdots,v_n]) = b * S(\partial[v_0,\cdots,v_n]) = \sum_i (-1)^i b * S([v_0,\cdots,\widehat{v}_i,\cdots,v_n])$ where b is the barycenter of $[v_0,\cdots,v_n]$ and the operation $b*$ is linear and sends a simplex $[w_0,\cdots,w_k]$ to $b*[w_0,\cdots,w_k] = [b,w_0,\cdots,w_k]$. More generally, define a homomorphism $S\colon C_n(X) \to C_n(X)$ by $S(\sigma) = \sigma_\sharp(S(\Delta^n))$ for $\sigma\colon \Delta^n \to X$, so that $S(\sigma)$ is a signed sum of the restrictions of σ to the n-simplices of the barycentric subdivision of Δ^n.

The homomorphism S is a chain map, $S\partial = \partial S$, since first of all

$$\begin{aligned}
\partial S(\Delta^n) &= \partial(b * S(\partial\Delta^n)) \\
&= S(\partial\Delta^n) - b * \partial S(\partial\Delta^n) \\
&\qquad \text{because } \partial \text{ acts by deleting vertices one at a time, and deleting initial} \\
&\qquad \text{vertices of the terms of } b * S(\partial\Delta^n) \text{ yields } S(\partial\Delta^n) \text{ while deleting} \\
&\qquad \text{the later vertices yields } -b * \partial S(\partial\Delta^n), \\
&= S(\partial\Delta^n) - b * S\partial(\partial\Delta^n) \\
&\qquad \text{by induction on } n, \\
&= S(\partial\Delta^n) \\
&\qquad \text{since } \partial(\partial\Delta^n) = 0
\end{aligned}$$

Hence for a singular n-simplex σ, $\partial S(\sigma) = \partial\sigma_\sharp S(\Delta^n) = \sigma_\sharp \partial S(\Delta^n) = \sigma_\sharp S(\partial\Delta^n) = S(\partial\sigma)$.

Next we construct a chain homotopy $T\colon C_n(X) \to C_{n+1}(X)$ between S and $\mathbb{1}$, satisfying $\partial T + T\partial = \mathbb{1} - S$. This comes from a subdivision of $\Delta^n \times [0,1]$ which restricts to $S(\Delta^n)$ on $\Delta^n \times \{1\}$ and to Δ^n on $\Delta^n \times \{0\}$, defined algebraically by the inductive formula $T(\Delta^n) = b * (\Delta^n - T(\partial\Delta^n))$ where b is the barycenter of $\Delta^n \times \{1\}$ and Δ^n is identified with $\Delta^n \times \{0\}$.

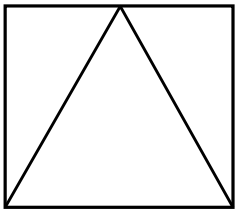

n = 1

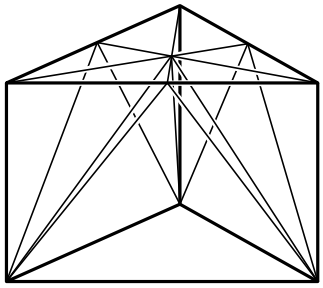

n = 2

For a singular n-simplex $\sigma\colon \Delta^n \to X$, $T(\sigma)$ is defined to be $\sigma_\sharp p_\sharp T(\Delta^n)$, where $p\colon \Delta^n \times [0,1] \to \Delta^n$ is the projection, so $T(\sigma)$ is a signed sum of the restrictions of σp to the

various $(n+1)$-simplices of $T(\Delta^n)$. To check that $\partial T + T\partial = \mathbb{1} - S$ we compute

$$\begin{aligned}
\partial T(\Delta^n) &= \partial(b * (\Delta^n - T(\partial\Delta^n))) \\
&= \Delta^n - T(\partial\Delta^n) - b * (\partial\Delta^n - \partial T(\partial\Delta^n)) \\
&\qquad \text{as before,} \\
&= \Delta^n - T(\partial\Delta^n) - b * (S(\partial\Delta^n) + T\partial(\partial\Delta^n)) \\
&\qquad \text{since } \mathbb{1} - \partial T = S + T\partial \text{ on } \partial\Delta^n \text{ by induction on } n, \\
&= \Delta^n - T(\partial\Delta^n) - b * S(\partial\Delta^n) \\
&\qquad \text{since } \partial\partial = 0, \\
&= \Delta^n - T(\partial\Delta^n) - S(\Delta^n)
\end{aligned}$$

Hence $\partial T(\sigma) = \partial\sigma_\sharp p_\sharp T(\Delta^n) = \sigma_\sharp p_\sharp \partial T(\Delta^n) = \sigma_\sharp p_\sharp(\Delta^n - T(\partial\Delta^n) - S(\Delta^n)) = \sigma - T(\partial\sigma) - S(\sigma)$.

Finally, we define a homomorphism $D : C_n(X) \to C_{n+1}(X)$ by specifying its value on a singular simplex $\sigma : \Delta^n \to X$ to be the sum $\sum_{0\le i<m} TS^i(\sigma)$ where $m = m(\sigma)$ is the smallest m such that the iterated subdivision $S^m(\sigma)$ lies in $C_n^{\mathcal{U}}(X)$. Such an m exists since the diameter of the simplices of $S^r(\Delta^n)$ will be less than a Lebesgue number of the cover of Δ^n by the open sets $\sigma^{-1}(int(U_j))$ if r is large enough. Geometrically, the formula $\sum_{0\le i<m} TS^i(\sigma)$ corresponds to stacking the m prisms $T(\Delta^n), TS(\Delta^n), \cdots, TS^{m-1}(\Delta^n)$ on top of each other to form one large prism with Δ^n on the bottom and $S^m(\Delta^n)$ on the top. Note that $D(\sigma) = 0$ if $\sigma \in C_n^{\mathcal{U}}(X)$ since $m = 0$ in this case.

We have

$$\begin{aligned}
\partial D(\sigma) = \sum_{0\le i<m} \partial TS^i(\sigma) &= \sum_{0\le i<m} [S^i(\sigma) - S^{i+1}(\sigma) - TS^i(\partial\sigma)] \\
&= \sigma - S^m(\sigma) - \sum_{0\le i<m} TS^i(\partial\sigma)
\end{aligned}$$

The last summation here contains every term in $D(\partial\sigma)$ since $m(\sigma\,|\,[v_0, \cdots, \widehat{v}_j, \cdots, v_n]) \le m(\sigma)$. Moreover, the terms of $\sum_{0\le i<m} TS^i(\partial\sigma)$ not in $D(\partial\sigma)$ all lie in $C_n^{\mathcal{U}}(X)$. Therefore the displayed formula can be written $\partial D(\sigma) + D(\partial\sigma) = \sigma - \rho(\sigma)$ where the chain $\rho(\sigma)$ lies in $C_n^{\mathcal{U}}(X)$. This formula defines $\rho : C_n(X) \to C_n^{\mathcal{U}}(X)$, and these ρ's form a chain map since $\partial\rho(\sigma) = \partial\sigma - \partial D\partial(\sigma) = \rho(\partial\sigma)$. The equation $\partial D(\sigma) + D(\partial\sigma) = \sigma - \rho(\sigma)$ says that $\partial D + D\partial = \mathbb{1} - \iota\rho$ for $\iota : C_n^{\mathcal{U}}(X) \hookrightarrow C_n(X)$ the inclusion. And $\rho\iota = \mathbb{1}$ since D is identically zero on $C_n^{\mathcal{U}}(X)$, as noted earlier. □

(3) Excision

A fundamental property of relative homology groups is given by the following Excision Theorem, describing when a subset $Z \subset A$ can be deleted without affecting the relative groups $H_n(X, A)$.

Theorem 2.15. *Given subspaces $Z \subset A \subset X$ such that the closure of Z is contained in the interior of A, then the inclusion $(X - Z, A - Z) \hookrightarrow (X, A)$ induces isomorphisms $H_n(X - Z, A - Z) \to H_n(X, A)$ for all n. Equivalently, for subspaces $A, B \subset X$ whose interiors cover X, then the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms $H_n(B, A \cap B) \to H_n(X, A)$ for all n.*

Proof: The first version implies the second by taking $Z = A - B$. In the other direction, take $B = X - Z$, so that $A \cap B = A - Z$. The condition $cl(Z) \subset int(A)$ is equivalent to $X = int(A) \cup int(B)$ since $X - int(B) = cl(Z)$.

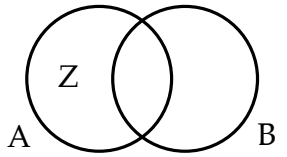

We prove the second version. For the cover $\mathcal{U} = \{A, B\}$ we introduce the suggestive notation $C_n(A + B)$ for $C_n^{\mathcal{U}}(X)$, the sums of chains in A and chains in B. The proof of Proposition 2.14 concluded with formulas $\partial D + D\partial = \mathbb{1} - \iota\rho$ and $\rho\iota = \mathbb{1}$. All the maps appearing in these formulas take chains in A to chains in A, hence induce quotient maps when we factor out chains in A. These quotient maps automatically satisfy the same two formulas, so the inclusion $C_n(A + B)/C_n(A) \hookrightarrow C_n(X)/C_n(A)$ induces an isomorphism on homology. The map $C_n(B)/C_n(A \cap B) \to C_n(A + B)/C_n(A)$ induced by inclusion is obviously an isomorphism since both quotient groups are free with basis the singular n-simplices in B which do not lie in A. Hence we obtain the desired isomorphism $H_n(B, A \cap B) \approx H_n(X, A)$, induced by inclusion. □

(4) Relative vs. Absolute Homology

All that remains in the proof of Theorem 2.9 is to replace relative homology groups with absolute homology groups.

Proposition 2.16. *For good pairs (X, A), the quotient map $q : (X, A) \to (X/A, A/A)$ induces isomorphisms $q_* : H_n(X, A) \to H_n(X/A, A/A) \approx \widetilde{H}_n(X/A)$ for all n.*

Proof: Let V be an open neighborhood of A in X which deformation retracts onto A. We have a commutative diagram

$$\begin{array}{ccccc} H_n(X, A) & \longrightarrow & H_n(X, V) & \longleftarrow & H_n(X - A, V - A) \\ \downarrow{\scriptstyle q_*} & & \downarrow{\scriptstyle q_*} & & \downarrow{\scriptstyle q_*} \\ H_n(X/A, A/A) & \longrightarrow & H_n(X/A, V/A) & \longleftarrow & H_n(X/A - A/A, V/A - A/A) \end{array}$$

The upper left horizontal map is an isomorphism since in the long exact sequence of the triple (X, V, A) the groups $H_n(V, A)$ are zero for all n, because in the long exact sequence of the pair (V, A) the maps $H_n(A) \to H_n(V)$ are isomorphisms by the assumption that

V deformation retracts to A. The same argument shows the lower left horizontal map is an isomorphism as well. The other two horizontal maps are isomorphisms directly from excision. The right-hand vertical map q_* is an isomorphism since q restricts to a homeomorphism on the complement of A. From the commutativity of the diagram it follows that the left-hand q_* is an isomorphism. □

This completes the proof of Theorem 2.9. □

The proposition implies that the excision property holds also for subcomplexes of CW complexes:

Corollary 2.17. *If the CW complex X is the union of subcomplexes A and B, then the inclusion $(B, A\cap B) \hookrightarrow (X, A)$ induces isomorphisms $H_n(B, A\cap B) \to H_n(X, A)$ for all n.*

Proof: We may assume we are not in the trivial case $A \cap B = \varnothing$. Since CW pairs are good, the proposition allows us to pass to the quotient spaces $B/(A\cap B)$ and X/A, which are homeomorphic. □

A second application of the preceding proposition is:

Corollary 2.18. *For a wedge sum $\bigvee_\alpha X_\alpha$, the inclusions $i_\alpha : X_\alpha \hookrightarrow \bigvee_\alpha X_\alpha$ induce an isomorphism $\bigoplus_\alpha i_{\alpha *} : \bigoplus_\alpha \widetilde{H}_n(X_\alpha) \to \widetilde{H}_n(\bigvee_\alpha X_\alpha)$, provided that the wedge sum is formed at basepoints $x_\alpha \in X_\alpha$ such that the pairs $(X_\alpha, \{x_\alpha\})$ are good.*

Proof: This follows from the proposition by taking $(X, A) = (\coprod_\alpha X_\alpha, \coprod_\alpha \{x_\alpha\})$. □

Here is another application of the machinery we have developed, a classical result of Brouwer from around 1910 known as 'invariance of dimension,' which says in particular that $\mathbb{R}^m$ is not homeomorphic to $\mathbb{R}^n$ if $m \neq n$.

Theorem 2.19. *If non-empty open sets $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ are homeomorphic, then $m = n$.*

Proof: For $x \in U$, excision gives $H_k(U, U - \{x\}) \approx H_k(\mathbb{R}^m, \mathbb{R}^m - \{x\})$. From the long exact sequence for the pair $(\mathbb{R}^m, \mathbb{R}^m - \{x\})$ we get $H_k(\mathbb{R}^m, \mathbb{R}^m - \{x\}) \approx \widetilde{H}_{k-1}(\mathbb{R}^m - \{x\})$. Since $\mathbb{R}^m - \{x\}$ deformation retracts onto S^{m-1}, we conclude that $H_k(U, U - \{x\})$ is $\mathbb{Z}$ for $k = m$ and 0 otherwise. By the same reasoning $H_k(V, V - \{y\})$ is $\mathbb{Z}$ for $k = n$ and 0 otherwise. Since a homeomorphism $h : U \to V$ induces isomorphisms $H_k(U, U - \{x\}) \to H_k(V, V - \{h(x)\})$ for all k, we must have $m = n$. □

Generalizing the idea of this proof, the **local homology groups** of a space X at a point $x \in X$ are defined to be the groups $H_n(X, X - \{x\})$. For any open neighborhood U of x, excision gives isomorphisms $H_n(X, X - \{x\}) \approx H_n(U, U - \{x\})$, so these groups depend only on the local topology of X near x. A homeomorphism $f : X \to Y$ must induce isomorphisms $H_n(X, X - \{x\}) \approx H_n(Y, Y - \{f(x)\})$ for all x and n, so these local homology groups can be used to tell when spaces are not locally homeomorphic at certain points, as in the preceding proof. The exercises give some more examples of this.

(5) Naturality

The exact sequences we have been constructing have an extra property which eventually becomes important at key points in various arguments, though initially this property may seem just an idle technicality, not very interesting. We shall discuss this property now rather than interrupting later arguments to check it when it is needed, but the reader may well want to postpone a careful reading of what follows.

The property is called **naturality**. For example, to say that the long exact sequence of a pair is natural means that whenever one has a map $f : (X, A) \to (Y, B)$, then the diagram

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & H_n(A) & \xrightarrow{i_*} & H_n(X) & \xrightarrow{j_*} & H_n(X, A) & \xrightarrow{\partial} & H_{n-1}(A) & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} & & \\
\cdots & \longrightarrow & H_n(B) & \xrightarrow{i_*} & H_n(Y) & \xrightarrow{j_*} & H_n(Y, B) & \xrightarrow{\partial} & H_{n-1}(B) & \longrightarrow & \cdots
\end{array}$$

is commutative. Commutativity of the squares involving i_* and j_* follows from the obvious commutativity of

$$\begin{array}{ccccccccc}
0 & \longrightarrow & C_n(A) & \xrightarrow{i} & C_n(X) & \xrightarrow{j} & C_n(X, A) & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle f_\sharp} & & \downarrow{\scriptstyle f_\sharp} & & \downarrow{\scriptstyle f_\sharp} & & \\
0 & \longrightarrow & C_n(B) & \xrightarrow{i} & C_n(Y) & \xrightarrow{j} & C_n(Y, B) & \longrightarrow & 0
\end{array}$$

Commutativity of the square containing $\partial : H_n(X, A) \to H_{n-1}(A)$ can be deduced from the relation $f_\sharp \partial = \partial f_\sharp$ on chains, which we noted when we defined induced homomorphisms. Namely, for a class $[\alpha] \in H_n(X, A)$ represented by a relative cycle α, we have $f_* \partial [\alpha] = f_* [\partial \alpha] = [f_\sharp \partial \alpha] = [\partial f_\sharp \alpha] = \partial [f_\sharp \alpha] = \partial f_* [\alpha]$.

Alternatively, we could appeal to the general algebraic fact that the long exact sequence of homology groups associated to a short exact sequence of chain complexes is natural: Given two short exact sequences of chain complexes and a map between them,

producing commutative diagrams

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A_n & \xrightarrow{i} & B_n & \xrightarrow{j} & C_n & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\
0 & \longrightarrow & A'_n & \xrightarrow{i'} & B'_n & \xrightarrow{j'} & C'_n & \longrightarrow & 0
\end{array}$$

for all n, then the diagram

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & H_n(A) & \xrightarrow{i_*} & H_n(B) & \xrightarrow{j_*} & H_n(C) & \xrightarrow{\partial} & H_{n-1}(A) & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle\alpha_*} & & \downarrow{\scriptstyle\beta_*} & & \downarrow{\scriptstyle\gamma_*} & & \downarrow{\scriptstyle\alpha_*} & & \\
\cdots & \longrightarrow & H_n(A') & \xrightarrow{i'_*} & H_n(B') & \xrightarrow{j'_*} & H_n(C') & \xrightarrow{\partial} & H_{n-1}(A') & \longrightarrow & \cdots
\end{array}$$

is commutative. Commutativity of the first two squares is obvious since $\beta i = i'\alpha$ implies $\beta_* i_* = i'_* \alpha_*$ and $\gamma j = j'\beta$ implies $\gamma_* j_* = j'_* \beta_*$. For the third square, recall the definition of the map $\partial : H_n(C) \to H_{n-1}(A)$, $\partial[c] = [a]$ where $c = j(b)$ and $i(a) = \partial b$. Then $\partial[\gamma(c)] = [\alpha(a)]$ since $\gamma(c) = \gamma j(b) = j'(\beta(b))$ and $i'(\alpha(a)) = \beta i(a) = \beta\partial(b) = \partial\beta(b)$. Hence $\partial\gamma_*[c] = \alpha_*[a] = \alpha_*\partial[c]$.

This algebraic fact also implies naturality of the long exact sequence of a triple and the long exact sequence of reduced homology of a pair.

Finally, there is the naturality of the long exact sequence in Theorem 2.9, i.e., commutativity of the diagram

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & \widetilde{H}_n(A) & \xrightarrow{i_*} & \widetilde{H}_n(X) & \xrightarrow{q_*} & \widetilde{H}_n(X/A) & \xrightarrow{\partial} & \widetilde{H}_{n-1}(A) & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle \overline{f}_*} & & \downarrow{\scriptstyle f_*} & & \\
\cdots & \longrightarrow & \widetilde{H}_n(B) & \xrightarrow{i_*} & \widetilde{H}_n(Y) & \xrightarrow{q_*} & \widetilde{H}_n(Y/B) & \xrightarrow{\partial} & \widetilde{H}_{n-1}(B) & \longrightarrow & \cdots
\end{array}$$

where i and q denote inclusions and quotient maps, and $\overline{f} : X/A \to Y/B$ is induced by f. The first two squares commute since $fi = if$ and $\overline{f}q = qf$. The third square expands into

$$\begin{array}{ccccccc}
\widetilde{H}_n(X/A) & \xrightarrow[\approx]{j_*} & H_n(X/A, A/A) & \xleftarrow[\approx]{q_*} & H_n(X, A) & \xrightarrow{\partial} & \widetilde{H}_{n-1}(X/A) \\
\downarrow{\scriptstyle \overline{f}_*} & & \downarrow{\scriptstyle \overline{f}_*} & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} \\
\widetilde{H}_n(Y/B) & \xrightarrow[\approx]{j_*} & H_n(Y/B, B/B) & \xleftarrow[\approx]{q_*} & H_n(Y, B) & \xrightarrow{\partial} & \widetilde{H}_{n-1}(Y/B)
\end{array}$$

We have already shown commutativity of the first and third squares, and the second square commutes since $\overline{f}q = qf$.

The Equivalence of Simplicial and Singular Homology

We can use the preceding results to show that the simplicial and singular homology groups of Δ-complexes are always isomorphic. For the proof of this it will be convenient to consider the relative case as well, so let X be a Δ-complex with $A \subset X$ a subcomplex, i.e., A is a union of simplices of X such that if a simplex of X is in A then so are all the faces of this simplex. This implies that A is a Δ-complex itself. Relative groups $H_n^\Delta(X, A)$ can be defined in the same way as for singular homology, via relative chains $\Delta_n(X, A) = \Delta_n(X)/\Delta_n(A)$, and this yields a long exact sequence of simplicial homology groups for the pair (X, A) by the same algebraic argument as for singular homology. There is a canonical homomorphism $H_n^\Delta(X, A) \to H_n(X, A)$ induced by the chain map $\Delta_n(X, A) \to C_n(X, A)$ sending each n-cell e_α^n of $X - A$ to its characteristic map $\sigma_\alpha : \Delta^n \to X$. The possibility $A = \varnothing$ is not excluded, in which case the relative groups reduce to absolute groups.

Theorem 2.20. *The homomorphisms $H_n^\Delta(X, A) \to H_n(X, A)$ are isomorphisms for all n and all Δ-complex pairs (X, A).*

Proof: First we do the case $(X, A) = (\Delta^k, \partial\Delta^k)$. We know that the groups $H_n^\Delta(\Delta^k, \partial\Delta^k)$ and $H_n(\Delta^k, \partial\Delta^k)$ are abstractly isomorphic for all n, being $\mathbb{Z}$ for $n = k$ and zero otherwise; for simplicial homology this is evident since $\Delta_n(\Delta^k, \partial\Delta^k)$ is zero except for $n = k$ when it is $\mathbb{Z}$, and for singular homology this calculation was an example earlier in this section. So it must only be shown that the canonical map $H_k^\Delta(\Delta^k, \partial\Delta^k) \to H_k(\Delta^k, \partial\Delta^k)$ is an isomorphism. Consider the triple $(\Delta^k, \partial\Delta^k, \Lambda)$ where Λ is the union of all but one of the $(k-1)$-simplices of $\partial\Delta^k$. We have a commutative diagram

$$\begin{array}{ccccc}
H_k^\Delta(\Delta^k, \partial\Delta^k) & \xrightarrow{\partial}_{\approx} & H_{k-1}^\Delta(\partial\Delta^k, \Lambda) & \xleftarrow[\approx]{} & H_{k-1}^\Delta(\Delta^{k-1}, \partial\Delta^{k-1}) \\
\downarrow & & \downarrow & & \downarrow \\
H_k(\Delta^k, \partial\Delta^k) & \xrightarrow{\partial}_{\approx} & H_{k-1}(\partial\Delta^k, \Lambda) & \xleftarrow[\approx]{} & H_{k-1}(\Delta^{k-1}, \partial\Delta^{k-1})
\end{array}$$

where the ∂ maps are part of the long exact sequences of the triple and the other horizontal maps are induced by the inclusion of the face Δ^{k-1} which is the closure of $\partial\Delta^k - \Lambda$. The upper two horizontal maps are clearly isomorphisms directly from the definitions. The lower ∂ map is an isomorphism since Δ^k deformation retracts onto Λ, making $H_n(\Delta^k, \Lambda) = 0$ for all n. The remaining horizontal map is an isomorphism by Corollary 2.17. The result now follows by induction on k, via the commutativity of the diagram.

Next consider the general absolute case. We have a commutative diagram of exact

sequences:

$$\begin{array}{ccccccccc} H_{n+1}^{\Delta}(X^k, X^{k-1}) & \longrightarrow & H_n^{\Delta}(X^{k-1}) & \longrightarrow & H_n^{\Delta}(X^k) & \longrightarrow & H_n^{\Delta}(X^k, X^{k-1}) & \longrightarrow & H_{n-1}^{\Delta}(X^{k-1}) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ H_{n+1}(X^k, X^{k-1}) & \longrightarrow & H_n(X^{k-1}) & \longrightarrow & H_n(X^k) & \longrightarrow & H_n(X^k, X^{k-1}) & \longrightarrow & H_{n-1}(X^{k-1}) \end{array}$$

The characteristic maps of the k-simplices give a map $\coprod_\alpha (\Delta_\alpha^k, \partial\Delta_\alpha^k) \to (X^k, X^{k-1})$ inducing isomorphisms on both simplicial and singular homology, using Proposition 2.16 in the case of singular homology. Hence the first and fourth vertical maps in the diagram are isomorphisms by the preceding special case. By induction on k the second and fifth vertical maps are isomorphisms. The following basic algebraic lemma, which will be used many times in the sequel, then implies that the middle vertical map is an isomorphism, which finishes the proof when X is finite-dimensional and $A = \varnothing$.

The Five-Lemma. *Suppose one has a commutative diagram of abelian groups*

$$\begin{array}{ccccccccc} A & \xrightarrow{i} & B & \xrightarrow{j} & C & \xrightarrow{k} & D & \xrightarrow{\ell} & E \\ \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \downarrow{\scriptstyle\delta} & & \downarrow{\scriptstyle\varepsilon} \\ A' & \xrightarrow{i'} & B' & \xrightarrow{j'} & C' & \xrightarrow{k'} & D' & \xrightarrow{\ell} & E' \end{array}$$

in which the two rows are exact and α, β, δ, and ε are isomorphisms. Then γ is an isomorphism also.

Proof: It suffices to show:

(a) γ is surjective if β and δ are surjective and ε is injective.

(b) γ is injective if β and δ are injective and α is surjective.

These are proved by straightforward diagram chasing. There is really no choice about how the argument can proceed, and it would be a good exercise for the reader to close the book now and reconstruct the proof without looking.

To prove (a), suppose $c' \in C'$. Then $k'(c) = \delta(d)$ for some $d \in D$ since δ is surjective. Since ε is injective and $\varepsilon\ell(d) = \ell'\delta(d) = \ell'k'(c) = 0$ we deduce that $\ell(d) = 0$, hence $d = k(c)$ for some $c \in C$ by exactness of the upper row. The difference $c' - \gamma(c)$ maps to 0 under k' since $k'(c') - k'\gamma(c) = k'(c) - \delta k(c) = k'(c') - \delta(d) = 0$. Therefore $c' - \gamma(c) = j'(b')$ for some $b' \in B'$ by exactness. Since β is surjective, $b' = \beta(b)$ for some $b \in B$, and then $\gamma(c + j(b)) = \gamma(c) + \gamma j(b) = \gamma(c) + j'\beta(b) = \gamma(c) + j'(b') = c'$, showing that γ is surjective.

To prove (b), suppose $\gamma(c) = 0$. Since δ is injective, $\delta k(c) = k'\gamma(c) = 0$ implies $k(c) = 0$, so $c = j(b)$ for some $b \in B$. The element $\beta(b)$ satisfies $j'\beta(b) = \gamma j(b) = \gamma(c) =$

0, so $\beta(b) = i'(a')$ for some $a' \in A'$. Since α is surjective, $a' = \alpha(a)$ for some $a \in A$. Since β is injective, $\beta(i(a) - b) = \beta i(a) - \beta(b) = i'\alpha(a) - \beta(b) = i'(a') - \beta(b) = 0$ implies $i(a) - b = 0$. Thus $b = i(a)$, and hence $c = j(b) = ji(a) = 0$ since $ji = 0$. This shows γ has trivial kernel. □

Returning to the proof of the theorem, we next consider the case that X is infinite-dimensional, where we will use the following fact: a compact set in X can meet only finitely many open simplices e^n_α. This is a general fact about CW complexes proved in the Appendix, but here is a direct proof for Δ-complexes. If a compact set C intersected infinitely many open simplices, it would contain an infinite sequence of points x_i each lying in a different open simplex. Then the sets $U_i = X - \bigcup_{j\neq i}\{x_j\}$, which are open since their pre-images under the characteristic maps of simplices are clearly open, form an open cover of C with no finite subcover.

This can be applied to show the map $H^\Delta_n(X) \to H_n(X)$ is surjective, as follows. Represent a given element of $H_n(X)$ by a singular n-cycle z. This is a linear combination of finitely many singular simplices with compact images, meeting only finitely many cells of X, hence contained in X^k for some k. We have shown that $H^\Delta_n(X^k) \to H_n(X^k)$ is an isomorphism, in particular surjective, so z is homologous in X^k to a simplicial cycle, a fortiori in X. This gives surjectivity. Injectivity is similar: If a simplicial n-cycle z is the boundary of a singular chain in X, this chain has compact image and hence must lie in some X^k, so z represents an element of the kernel of $H^\Delta_n(X^k) \to H_n(X^k)$. But we know this map is injective, so z is a simplicial boundary in X^k, a fortiori in X.

It remains to do the case of arbitrary X with $A \neq \varnothing$, but this follows from the absolute case by applying the five-lemma to the canonical map from the long exact sequence of simplicial homology groups for the pair (X, A) to the corresponding long exact sequence of singular homology groups. □

The theorem says in particular that $H_n(\Delta^n, \partial\Delta^n)$ is generated by the identity map $\Delta^n \to \Delta^n$, viewed as a singular n-simplex which is a relative cycle. Similarly, if we view S^n as the Δ-complex obtained from two copies of Δ^n by identifying their boundaries, then the difference between these two singular n-simplices in S^n is a cycle representing a generator of $H_n(S^n)$, since this is true for simplicial homology, as we saw in an example early in this section.

We can also deduce that $H_n(X)$ is finitely generated if X is a Δ-complex with finitely many n-simplices, since in this case the simplicial chain group $\Delta_n(X)$ is finitely generated, hence also its subgroup of cycles and therefore also the latter group's quotient

$H_n^\Delta(X)$. If we write $H_n(X)$ as the direct sum of cyclic groups, then the number of $\mathbb{Z}$ summands is known traditionally as the n^{th} **Betti number** of X, and integers specifying the orders of the finite cyclic summands are called **torsion coefficients**.

It is a curious historical fact that homology was not thought of originally as a sequence of groups, but rather as Betti numbers and torsion coefficients. One can after all compute Betti numbers and torsion coefficients from the simplicial boundary maps without actually mentioning the homology groups. This computational viewpoint, with homology being numbers rather than groups, prevailed from when Poincaré first started serious work on homology around 1900, up until the 1920's when the more abstract viewpoint of groups entered the picture. During this period 'homology' meant primarily 'simplicial homology,' and it was another 20 years before the shift to singular homology was complete, with the final definition of singular homology emerging only in a 1944 paper of Eilenberg, after contributions from quite a few others, particularly Alexander and Lefschetz. Within the next few years the rest of the basic structure of homology theory as we have presented it fell into place, and the first definitive treatment appeared in the classic book [Eilenberg-Steenrod] of 1952.

Exercises

1. What familiar space is the quotient Δ-complex of a 2-simplex $[v_0, v_1, v_2]$ obtained by identifying the edges $[v_0, v_1]$ and $[v_1, v_2]$, preserving the ordering of vertices?

2. Compute the simplicial homology groups of the triangular parachute obtained from Δ^2 by identifying its three vertices to a single point.

3. Compute the simplicial homology groups of the Klein bottle using the Δ-complex structure described in §2.1.

4. Find a way of identifying pairs of faces of Δ^3 to produce a Δ-complex structure on S^3 having a single 3-simplex, and compute the simplicial homology groups of this Δ-complex.

5. Compute the homology groups of the Δ-complex X obtained from Δ^n by identifying all pairs of faces having the same dimension. Thus X has a single k-simplex for each $k \leq n$.

6. Construct a Δ-complex structure on $\mathbb{R}\mathrm{P}^n$ as a quotient of a Δ-complex structure on S^n having vertices the two vectors of length 1 along each coordinate axis in $\mathbb{R}^{n+1}$.

7. Show that the quotient space obtained from a finite collection of disjoint 2-simplices by identifying pairs of edges is a surface, i.e., locally homeomorphic to $\mathbb{R}^2$. [Optional refinement: show the edges can always be oriented so as to define a Δ-complex structure on the surface.]

8. Show that if A is a retract of X then the map $H_n(A) \to H_n(X)$ induced by the inclusion $A \subset X$ is injective.

9. Show that a map $f : X \to Y$ induces homomorphisms $f_* : \widetilde{H}_n(X) \to \widetilde{H}_n(Y)$.

10. Show that chain homotopy of chain maps is an equivalence relation.

11. Determine whether there exists a short exact sequence $0 \to \mathbb{Z}_4 \to \mathbb{Z}_8 \oplus \mathbb{Z}_2 \to \mathbb{Z}_4 \to 0$. More generally, determine which abelian groups A fit into a short exact sequence $0 \to \mathbb{Z}_{p^m} \to A \to \mathbb{Z}_{p^n} \to 0$ with p prime. What about short exact sequences $0 \to \mathbb{Z} \to A \to \mathbb{Z}_n \to 0$?

12. Show that for the subspace $\mathbb{Q} \subset \mathbb{R}$, the relative homology group $H_1(\mathbb{R}, \mathbb{Q})$ is free abelian and find a basis.

13. Show that $S^1 \times S^1$ and $S^1 \vee S^1 \vee S^2$ have isomorphic homology groups in all dimensions, but their universal covering spaces do not.

14. Show that $\widetilde{H}_n(X) \approx \widetilde{H}_{n+1}(SX)$ for all n, where SX is the suspension of X.

15. Refining the preceding problem, construct chain maps $s : C_n(X) \to C_{n+1}(SX)$ which induce isomorphisms $\widetilde{H}_n(X) \to \widetilde{H}_{n+1}(SX)$.

16. Show that the second barycentric subdivision of a Δ-complex is a simplicial complex. Namely, show that the first barycentric subdivision produces a Δ-complex with the property that each simplex has all its vertices distinct, then show that for a Δ-complex with this property, barycentric subdivision produces a simplicial complex.

17. Show that each n-simplex in the barycentric subdivision of Δ^n is defined by n inequalities $t_{i_0} \le t_{i_1} \le \cdots \le t_{i_n}$ in its barycentric coordinates, where $(i_0, \cdots, i_n)$ is a permutation of $(0, \cdots, n)$.

18. Find an explicit, non-inductive formula for the barycentric subdivision operator $S : C_n(X) \to C_n(X)$.

19. Show that $H_1(X, A)$ is not isomorphic to $\widetilde{H}_1(X/A)$ if $X = [0, 1]$ and A is the sequence $1, 1/2, 1/3, \cdots$ together with its limit 0. [See §1.2 where we showed that $\pi_1(X/A)$ is uncountable.]

20. For a space A, let CA denote the cone $(A \times I)/(A \times \{0\})$. If A is a subspace of X, let $X \cup CA$ be the space obtained from the disjoint union of X and CA by identifying the subspaces $A = A \times \{1\} \subset CA$ and $A \subset X$. Show that $H_n(X, A) \approx \widetilde{H}_n(X \cup CA)$. [Thus relative homology always reduces to absolute homology.]

21. Let $f : (X, A) \to (Y, B)$ be a map such that both $f : X \to Y$ and the restriction $f : A \to B$ are homotopy equivalences.

(a) Show that $f_*: H_n(X, A) \to H_n(Y, B)$ is an isomorphism for all n.

(b) For the case of the inclusion $f: (D^n, S^{n-1}) \hookrightarrow (D^n, D^n - \{0\})$, show that f is not a homotopy equivalence of pairs, i.e., there is no map $g: (D^n, D^n - \{0\}) \to (D^n, S^{n-1})$ such that fg and gf are homotopic to the identity through maps of pairs.

22. Let X be the cone on the 1-skeleton of Δ^3, i.e., the union of all line segments joining points in the six edges of Δ^3 to the barycenter of Δ^3. Compute the local homology groups $H_n(X, X - \{x\})$ for all $x \in X$. Define ∂X to be the subspace of points x such that $H_n(X, X - \{x\}) = 0$ for all n, and compute the local homology groups $H_n(\partial X, \partial X - \{x\})$. Use these calculations to determine which subsets $A \subset X$ have the property that $f(A) \subset A$ for all homeomorphisms $f: X \to X$.

2. Computations and Applications

Now that the basic properties of homology have been established, we can begin to move a little more freely. Our first topic, exploiting the calculation of $H_n(S^n)$, is Brouwer's notion of degree for maps $S^n \to S^n$. Historically, Brouwer's introduction of this concept about 1910 preceded the rigorous development of homology, so his definition was rather different, using the technique of simplicial approximation which we explain in §2.4c. The later definition in terms of homology is certainly more elegant, though perhaps with some loss of geometric intuition. More in the spirit of Brouwer's definition is a third approach using differential topology, presented very lucidly in [Milnor 65].

Degree

For a map $f: S^n \to S^n$, the induced $f_*: \widetilde{H}_n(S^n) \to \widetilde{H}_n(S^n)$ is a homomorphism from an infinite cyclic group to itself and so must be of the form $f_*(\alpha) = d\alpha$ for some integer d depending only on f. This integer is called the **degree** of f, $deg(f)$. Here are some basic properties of degree.

(a) $deg(\mathbb{1}) = 1$, since $\mathbb{1}_* = \mathbb{1}$.

(b) $deg(f) = 0$ if f is not surjective. For if we choose a point $x_0 \in S^n - f(S^n)$ then f can be factored as a composition $S^n \to S^n - \{x_0\} \hookrightarrow S^n$ and $H_n(S^n - \{x_0\}) = 0$ since $S^n - \{x_0\}$ is contractible. Hence $f_* = 0$.

(c) If $f \simeq g$ then $deg(f) = deg(g)$ since $f_* = g_*$. The converse statement, that $f \simeq g$ if $deg(f) = deg(g)$, is a fundamental theorem of Hopf from around 1925 which we will prove in §4.1 as Corollary 4.8.

(d) $deg(fg) = deg(f)deg(g)$, since $(fg)_* = f_*g_*$. A consequence of this is that $deg(f) = \pm 1$ if f is a homotopy equivalence since $fg \simeq \mathbb{1}$ implies $deg(f)deg(g) = deg(\mathbb{1}) = 1$.

(e) $deg(f) = -1$ if f is a reflection of S^n, fixing the points in a subsphere S^{n-1} and interchanging the two complementary hemispheres. For we can give S^n a Δ-complex structure with these two hemispheres as its two n-simplices Δ_1^n and Δ_2^n, and the simplicial n-chain $\Delta_1^n - \Delta_2^n$ represents a generator of $H_n(S^n)$ as we saw at the end of the last section, so the reflection interchanging Δ_1^n and Δ_2^n sends this generator to its negative.

(f) The antipodal map $-\mathbb{1}: S^n \to S^n$, $x \mapsto -x$, has degree $(-1)^{n+1}$ since it is the composition of $n+1$ reflections, each changing the sign of one coordinate in $\mathbb{R}^{n+1}$.

(g) If $f: S^n \to S^n$ has no fixed points then $deg(f) = (-1)^{n+1}$. For if $f(x) \neq x$ then the line segment from $f(x)$ to $-x$, defined by $t \mapsto (1-t)f(x) - tx$ for $0 \leq t \leq 1$, does not pass through the origin. Hence if f has no fixed points, the formula $f_t(x) =$

$[(1-t)f(x)-tx]/|(1-t)f(x)-tx|$ defines a homotopy from f to the antipodal map. Note that the antipodal map has no fixed points, so the fact that maps without fixed points are homotopic to the antipodal map is a sort of converse statement.

An interesting application of the fact that the antipodal map of S^n has degree $(-1)^{n+1}$ is:

Theorem 2.21. *S^n has a continuous field of nonzero tangent vectors iff n is odd.*

Proof: Suppose $x \mapsto v(x)$ is a tangent vector field on S^n, assigning to a vector $x \in S^n$ the vector $v(x)$ tangent to S^n at x. Regarding $v(x)$ as a vector at the origin instead of at x, tangency just means that x and $v(x)$ are orthogonal in $\mathbb{R}^{n+1}$. If $v(x) \neq 0$ for all x we may normalize $v(x)$ so that $|v(x)| = 1$ by replacing $v(x)$ by $v(x)/|v(x)|$. Assuming this has been done, consider the great circle $(\cos t)x + (\sin t)v(x)$ in the plane spanned by x and $v(x)$. Letting t go from 0 to π defines a homotopy $f_t(x) = (\cos t)x + (\sin t)v(x)$ from the identity map of S^n to the antipodal map $-\mathbb{1}$. This implies that $deg(-\mathbb{1}) = deg(\mathbb{1})$, hence $(-1)^{n+1} = 1$ and n must be odd.

Conversely, if n is odd, say $n = 2k-1$, we can define $v(x_1, x_2, \cdots, x_{2k-1}, x_{2k}) = (-x_2, x_1, \cdots, -x_{2k}, x_{2k-1})$. Then $v(x)$ is orthogonal to x, so v is a tangent vector field on S^n which is nonzero at all points. □

For the much more difficult problem of finding the maximum number of tangent vector fields on S^n which are linearly independent at each point, see [VBKT] or [Husemoller].

Here is another nice application of degree, giving a partial answer to a question raised at the end of Chapter 1:

Proposition 2.22. *$\mathbb{Z}_2$ is the only nontrivial group which can act freely on S^n if n is even.*

Recall that an action of a group G on a space X is a homomorphism from G to the group $Homeo(X)$ of homeomorphisms $X \to X$, and the action is free if the homeomorphism corresponding to each nontrivial element of the group has no fixed points. In the case of S^n, the antipodal map $x \mapsto -x$ generates a free action of $\mathbb{Z}_2$ for each n.

Proof: Since the degree of a homeomorphism must be ± 1, an action of a group G on S^n determines a degree function $d: G \to \{\pm 1\}$. This is a homomorphism since $deg(fg) = deg(f)deg(g)$. If the action is free, then d sends every nontrivial element of G to $(-1)^{n+1}$ by property (g) above. Thus when n is even, d has trivial kernel, so $G \subset \mathbb{Z}_2$. □

We shall next describe a technique for computing degrees which can be applied to most maps that arise in practice. Suppose $f: S^n \to S^n$ has the property that for some point $y \in S^n$, the preimage $f^{-1}(y)$ consists of only finitely many points, say $x_1, \cdots, x_m$. Let $U_1, \cdots, U_m$ be disjoint neighborhoods of these points, mapped by f into a neighborhood V of y. Then Then $f(U_i - x_i \subset V - y$ for each i, and we have a commutative diagram

$$\begin{array}{ccccc}
 & & H_n(U_i, U_i - x_i) & \xrightarrow{f_*} & H_n(V, V-y) \\
 & \overset{\approx}{\swarrow} & \downarrow{\scriptstyle k_i} & & \downarrow{\scriptstyle \approx} \\
H_n(S^n, S^n - x_i) & \xleftarrow{p_i} & H_n(S^n, S^n - f^{-1}(y)) & \xrightarrow{f_*} & H_n(S^n, S^n - y) \\
 & \underset{\approx}{\nwarrow} & \uparrow{\scriptstyle j} & & \uparrow{\scriptstyle \approx} \\
 & & H_n(S^n) & \xrightarrow{f_*} & H_n(S^n)
\end{array}$$

where all the maps are the obvious ones, in particular k_i and p_i are induced by inclusions. The two isomorphisms in the upper half of the diagram come from excision, while the lower two isomorphisms come from exact sequences of pairs. Via these four isomorphisms, the top two groups in the diagram can be identified with $H_n(S^n) \approx \mathbb{Z}$, and the top homomorphism f_* becomes multiplication by an integer called the **local degree** of f at x_i, written $deg(f \,|\, x_i)$.

For example, if f is a homeomorphism, then y can be any point and there is only one corresponding x_i, so all the maps in the diagram are isomorphisms and $deg(f \,|\, x_i) = deg(f) = \pm 1$. More generally, if f maps each U_i homeomorphically onto V, then $deg(f \,|\, x_i) = \pm 1$ for each i. This situation occurs quite often in applications, and it is usually not hard to determine the correct sign by keeping track of orientations — see §3.3.

Here is the formula which reduces degree calculations to computing local degrees:

Proposition 2.23. $deg(f) = \sum_i deg(f \,|\, x_i)$.

Proof: By excision, the central term $H_n(S^n, S^n - f^{-1}(y))$ in the preceding diagram is the direct sum of the groups $H_n(U_i, U_i - x_i) \approx \mathbb{Z}$, with the projections onto its summands given by the maps p_i and with k_i the inclusion of the i^{th} summand. Identifying the outer groups in the diagram with $\mathbb{Z}$ as before, then $p_i j(1) = 1$ implies $j(1) = (1, \cdots, 1) = \sum_i k_i(1)$. Commutativity of the upper square says that the middle f_* takes $k_i(1)$ to $deg(f \,|\, x_i)$. The result then follows from commutativity of the lower square. □

Example 1. We can use this result to construct a map $S^n \to S^n$ of any given degree, for each $n \geq 1$. Let $q: S^n \to \bigvee_k S^n$ be the quotient map obtained by collapsing the complement of k disjoint balls B_i in S^n to a point, and let $p: \bigvee_k S^n \to S^n$ identify

all the summands to a single sphere. Consider the composition $f = pq$. For almost all $y \in S^n$ we have $f^{-1}(y)$ consisting of k points x_i in the interiors of the balls B_i. The local degree of f at x_i is ± 1 since f is a homeomorphism near x_i. By inserting reflections of the summands of $\bigvee_k S^n$ if necessary, we can make each local degree either $+1$ or -1, whichever we wish. Thus we can produce a map $S^n \to S^n$ of degree $\pm k$.

Example 2. A more direct construction in the case of S^1 is to let $f(z) = z^k$ for $k \in \mathbb{Z}$, where we view S^1 as the unit circle in $\mathbb{C}$. To show that $deg(f) = k$, observe that for any $y \in S^1$, $f^{-1}(y)$ consists of $|k|$ points, near each of which f is a local homeomorphism, with local degree $+1$ if $k > 0$ and local degree -1 if $k < 0$. Namely, near each of these points f can be homotoped, stretching by a factor of $|k|$ without changing local degree, to become the restriction of either a rotation or a reflection of S^1 according to whether k is positive or negative. A rotation has degree $+1$ since it is homotopic to the identity, while a reflection has degree -1, as we have seen. Since rotations and reflections are homeomorphisms, their degree equals their local degree at any point, as we noted earlier.

Another way of obtaining a map $S^n \to S^n$ of degree k is to take a repeated suspension of the map $z \mapsto z^k$ in Example 2, since suspension preserves degree:

Proposition 2.24. *$deg(Sf) = deg(f)$ where $Sf : S^{n+1} \to S^{n+1}$ is the suspension of $f : S^n \to S^n$.*

Proof: Let CS^n denote the cone $(S^n \times I)/(S^n \times 1)$ with base $S^n = S^n \times 0 \subset CS^n$, so CS^n/S^n is the suspension of S^n. The map f induces $Cf : (CS^n, S^n) \to (CS^n, S^n)$ with quotient Sf. The naturality of the boundary maps in the long exact sequence of the pair (CS^n, S^n) then gives a commutative diagram:

$$\begin{array}{ccc} \widetilde{H}_{n+1}(S^{n+1}) & \xrightarrow[\approx]{\partial} & \widetilde{H}_n(S^n) \\ \downarrow{\scriptstyle Sf_*} & & \downarrow{\scriptstyle f_*} \\ \widetilde{H}_{n+1}(S^{n+1}) & \xrightarrow[\approx]{\partial} & \widetilde{H}_n(S^n) \end{array}$$

Hence if f_* is multiplication by d, so is Sf_*. □

Cellular Homology

Cellular homology is a very efficient tool for computing the homology groups of CW complexes, based on degree calculations. Before giving the definition of cellular homology, we first establish a few preliminary facts:

Lemma 2.25. *If X is a CW complex, then:*

(a) *$H_k(X^n, X^{n-1})$ is zero for $k \neq n$, and is free abelian for $k = n$, with a basis in one-to-one correspondence with the n-cells of X.*

(b) *$H_k(X^n) = 0$ for $k > n$. In particular, if X is finite-dimensional then $H_k(X) = 0$ for $k > dim(X)$.*

(c) *The inclusion $i : X^n \hookrightarrow X$ induces an isomorphism $i_* : H_k(X^n) \to H_k(X)$ if $k < n$.*

Proof: Characteristic maps Φ_α for the n-cells e^n_α of X give a map $(\coprod_\alpha D^n_\alpha, \coprod_\alpha \partial D^n_\alpha) \to (X^n, X^{n-1})$ inducing an isomorphism on all relative homology groups since by Proposition 2.16 we can pass to the quotient map $\coprod_\alpha D^n_\alpha / \coprod_\alpha \partial D^n_\alpha \to X^n/X^{n-1}$, which is a homeomorphism. Thus $H_k(X^n, X^{n-1}) \approx \bigoplus_\alpha H_k(D^n_\alpha, \partial D^n_\alpha)$ which, from the long exact sequence of reduced homology groups for the pair $(D^n, \partial D^n)$, is zero unless $k = n$ when it is a direct sum $\bigoplus_\alpha \mathbb{Z}$ with one $\mathbb{Z}$ for each n-cell e^n_α. This proves (a).

From the long exact sequence for the pair (X^n, X^{n-1}) we obtain $H_k(X^{n-1}) \approx H_k(X^n)$ for $k \neq n,\ n-1$. In particular, if $k > n$ then $H_k(X^n) \approx H_k(X^{n-1}) \approx H_k(X^{n-2}) \approx \cdots \approx H_k(X^0) = 0$, proving (b). Further, if $k < n$ then $H_k(X^n) \approx H_k(X^{n+1}) \approx \cdots \approx H_k(X^{n+m})$ for all $m \geq 0$, proving (c) if X is finite-dimensional.

The proof of (c) when X is infinite dimensional requires more work, and this can be done in two different ways. The more direct approach is to descend to the chain level and use the fact that a singular chain in X has compact image, hence lies in a finite subcomplex of X by Proposition A.1 in the Appendix, and in particular in some finite skeleton X^m. So a k-cycle in X is a cycle in some X^m, and then by the finite-dimensional case of (c), the cycle is homologous to a cycle in X^n if $n > k$, so $i_* : H_k(X^n) \to H_k(X)$ is surjective. Similarly for injectivity, if a k-cycle in X^n bounds a chain in X, this chain lies in some X^m with $m \geq n$, so by the finite-dimensional case the cycle bounds a chain in X^n if $n > k$.

The other approach is more general. From the long exact sequence of the pair (X, X^n) it suffices to show $H_k(X, X^n) = 0$ for $k \leq n$. Since $H_k(X, X^n) \approx \widetilde{H}_k(X/X^n)$, this reduces us to showing:

$(*)$ $\widetilde{H}_k(X) = 0$ for $k \leq n$ if the n-skeleton of X is a point.

When X is finite-dimensional, $(*)$ is immediate from the finite-dimensional case of (c) which we have already shown. It will suffice therefore to reduce the infinite-dimensional case to the finite-dimensional case. This reduction will be achieved by stretching X out to a complex which is at least locally finite-dimensional, using a special case of the 'mapping telescope' construction described more fully in §3.4f.

Consider $X \times [0, \infty)$ with its product cell structure, where we give $[0, \infty)$ the cell

structure with the integer points as vertices. Let $T = \bigcup_i X^i \times [i, \infty)$, a subcomplex of $X \times [0, \infty)$. We claim that the quotient $X \times [0, \infty)/T$ is contractible. To see this, consider the translations $\tau_t : X \times [0, \infty) \to X \times [0, \infty)$, $\tau_t(x, s) = (x, s + t)$ for $t \geq 0$. Since $\tau_t(T) \subset T$ for all t and τ_t maps each cell of $X \times [0, \infty)$ into T for t sufficiently large, τ_t induces a contraction of $X \times [0, \infty)/T$. Since this quotient is contractible, we have $H_k(X) \approx H_k(X \times [0, \infty)) \approx H_k(T)$ for all k, from the long exact sequence of the pair $(X \times [0, \infty), T)$.

Recalling that X^0 is a point, let $R \subset T$ be the ray $X^0 \times [0, \infty)$, and let $Z \subset T$ be the union of this ray with all the subcomplexes $X^i \times \{i\}$. Then Z/R is a wedge of finite-dimensional complexes with n-skeleton a point, so the finite-dimensional case of $(*)$ together with Corollary 2.18 implies that $\widetilde{H}_k(Z/R) = 0$ for $k \leq n$. The same is therefore true for Z, from the long exact sequence of the pair (Z, R), since R is contractible. Similarly, T/Z is a wedge of finite-dimensional complexes with $(n+1)$-skeleton a point, so $\widetilde{H}_k(T/Z) = 0$ for $k \leq n+1$. The long exact sequence of the pair (T, Z) then implies that $\widetilde{H}_k(T) = 0$ for $k \leq n$, and we have proved $(*)$. □

Let X be a CW complex. Using Lemma 2.25, portions of the long exact sequences for the pairs (X^{n+1}, X^n), (X^n, X^{n-1}), and (X^{n-1}, X^{n-2}) fit into a diagram

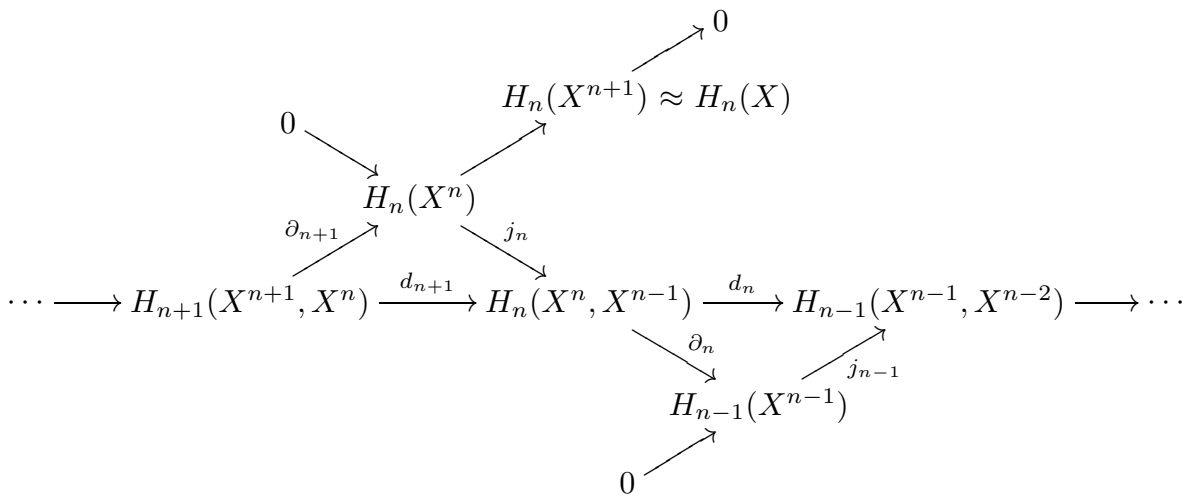

where d_{n+1} and d_n are defined as the compositions $j_n \partial_{n+1}$ and $j_{n-1} \partial_n$, i.e., the 'relativized' boundary maps ∂_{n+1} and ∂_n. The composition $d_n d_{n+1}$ includes two successive maps in one of the exact sequences, hence is zero. Thus the horizontal row in the diagram is a chain complex, called the **cellular chain complex** of X since $H_n(X^n, X^{n-1})$ is free with basis in one-to-one correspondence with the n-cells of X, so one can think of elements of $H_n(X^n, X^{n-1})$ as linear combinations of n-cells of X. The homology groups

of this cellular chain complex are called the **cellular homology groups** $H_n^{CW}(X)$ of X.

Theorem 2.26. $H_n^{CW}(X) \approx H_n(X)$.

Proof: From the diagram above, $H_n(X)$ can be identified with $H_n(X^n)/Im(\partial_{n+1})$. Since j_n is injective, it maps $Im(\partial_{n+1})$ isomorphically onto $Im(d_{n+1})$ and $H_n(X^n)$ isomorphically onto $Im(j_n) = Ker(\partial_n)$. Since j_{n-1} is injective, $Ker(\partial_n) = Ker(d_n)$. Thus j_n induces an isomorphism of the quotient $H_n(X^n)/Im(\partial_{n+1})$ onto $Ker(d_n)/Im(d_{n+1})$. □

Here are a few immediate applications:

(*i*) $H_n(X) = 0$ if X is a CW complex with no n-cells.

(*ii*) $H_n(X)$ is finitely generated if X is a CW complex with finitely many n-cells. For $Ker(d_n)$, as a subgroup of the finitely generated abelian group $H_n(X^n, X^{n-1})$, must be finitely generated, hence also the quotient $Ker(d_n)/Im(d_{n+1})$.

(*iii*) If X is a CW complex having no two of its cells in adjacent dimensions, then $H_n(X)$ is free abelian with basis in one-to-one correspondence with the n-cells of X. This is because the cellular boundary maps d_n are automatically zero in this case.

This last observation applies for example to $\mathbb{C}\mathrm{P}^n$ which has a CW structure with one cell of each even dimension $2k \leq 2n$. Thus:

$$H_i(\mathbb{C}\mathrm{P}^n) \approx \begin{cases} \mathbb{Z} & \text{for } i = 0, 2, 4, \cdots, 2n \\ 0 & \text{otherwise} \end{cases}$$

Another simple example would be $S^n \times S^n$, $n > 1$, with the product CW structure consisting of a 0-cell, two n-cells and a $2n$-cell.

It is possible to give more elementary proofs of the statements (*i*)-(*iii*) for finite-dimensional CW compexes by induction on the dimension, using only the basic results from the previous section. However, the viewpoint of cellular homology makes (*i*)-(*iii*) quite transparent.

Next we describe how the cellular boundary maps d_n can be computed in terms of degrees of maps of spheres. Consider the following commutative diagram, assuming $n > 1$ for the moment:

$$\begin{array}{ccccc}
H_n(D^n_\alpha, \partial D^n_\alpha) & \xrightarrow[\approx]{\partial} & \widetilde{H}_{n-1}(\partial D^n_\alpha) & \xrightarrow{\Delta_{\alpha\beta *}} & \widetilde{H}_{n-1}(S^{n-1}_\beta) \\
\downarrow{\scriptstyle \Phi_{\alpha *}} & & \downarrow{\scriptstyle \varphi_{\alpha *}} & & \uparrow{\scriptstyle q_{\beta *}} \\
H_n(X^n, X^{n-1}) & \xrightarrow{\partial_n} & \widetilde{H}_{n-1}(X^{n-1}) & \xrightarrow{q_*} & \widetilde{H}_{n-1}(X^{n-1}/X^{n-2}) \\
& \searrow^{d_n} & \downarrow{\scriptstyle j_{n-1}} & & \downarrow{\scriptstyle \approx} \\
& & H_{n-1}(X^{n-1}, X^{n-2}) & \xrightarrow{\approx} & H_{n-1}(X^{n-1}/X^{n-2}, X^{n-2}/X^{n-2})
\end{array}$$

Here Φ_α is the characteristic map of the cell e^n_α, φ_α is its attaching map, $q: X^{n-1} \to X^{n-1}/X^{n-2}$ is the quotient map, and $q_\beta: X^{n-1}/X^{n-2} \to S^{n-1}_\beta$ collapses the complement of the cell e^{n-1}_β to a point, the resulting quotient sphere being identified with $S^{n-1}_\beta = D^{n-1}_\beta/\partial D^{n-1}_\beta$ via the characteristic map Φ_β. The map $\Delta_{\alpha\beta}$ is the composition $q_\beta q \varphi_\alpha$.

The map $\Phi_{\alpha *}$ takes a chosen generator $[D^n_\alpha] \in H_n(D^n_\alpha, \partial D^n_\alpha)$ to a generator x_α of the $\mathbb{Z}$ summand of $H_n(X^n, X^{n-1})$ corresponding to e^n_α. Commutativity of the left half of the diagram then gives $d_n(x_\alpha) = j_{n-1}\varphi_{\alpha *}\partial[D^n_\alpha]$. In terms of the basis $\{y_\beta\}$ for $H_{n-1}(X^{n-1}, X^{n-2})$ corresponding to the cells e^{n-1}_β, the map $q_{\beta *}$ is the projection of $H_{n-1}(X^{n-1}/X^{n-2})$ onto its $\mathbb{Z}$ summand corresponding to y_β. Commutativity of the diagram then yields the formula we are looking for:

$$d_n(x_\alpha) = \sum_\beta d_{\alpha\beta} y_\beta \quad \text{where } d_{\alpha\beta} \text{ is the degree of the map } \Delta_{\alpha\beta}.$$

Note that $\sum_\beta d_{\alpha\beta} y_\beta$ has only finitely many nonzero terms since $H_{n-1}(X^{n-1}/X^{n-2}, x_0)$ is the direct sum of the groups $H_{n-1}(S^{n-1}_\beta)$.

When $n = 1$ we have $X^{n-2} = \varnothing$, and the preceding analysis applies if we define $Y/\varnothing$ to be the disjoint union of Y with a point. There is some logic to this definition if one considers that $Y/\varnothing$ should be a larger space than $Y/y_0 = Y$ for $y_0 \in Y$. For actual calculations it is not necessary to resort to this trick since the cellular boundary map $d_1: H_1(X^1, X^0) \to H_0(X^0)$ is the same as the simplicial boundary map $\Delta_1(X) \to \Delta_0(X)$, which is easy to compute. In many cases of interest X is connected and has only one 0-cell, and then d_1 must be 0, otherwise $H_0(X)$ would not be $\mathbb{Z}$.

Example 1. Let M_g be the closed orientable surface of genus g with its usual CW structure consisting of one 0-cell, $2g$ 1-cells, and one 2-cell attached by the product of commutators $[a_1, b_1] \cdots [a_g, b_g]$, the associated cellular chain complex is:

$$0 \longrightarrow \mathbb{Z} \xrightarrow{d_2} \mathbb{Z}^{2g} \xrightarrow{d_1} \mathbb{Z} \longrightarrow 0$$

As observed above, d_1 must be 0 since there is only one 0-cell. Also d_2 is 0 in this case because each a_i or b_i appears along with its inverse in $[a_1, b_1] \cdots [a_g, b_g]$, so the maps $\Delta_{\alpha\beta}$ are homotopic to constant maps. Since d_1 and d_2 are both zero, the homology groups of M_g are the same as the cellular chain groups, namely, $\mathbb{Z}$ in dimensions 0 and 2, and $\mathbb{Z}^{2g}$ in dimension 1.

Example 2. The closed non-orientable surface N_g of genus g has a cell structure with one 0-cell, g 1-cells, and one 2-cell attached by the word $a_1^2 a_2^2 \cdots a_g^2$. Again $d_1 = 0$,

and $d_2 : \mathbb{Z} \to \mathbb{Z}^g$ is specified by the equation $d_2(1) = (2, \cdots, 2)$ since each a_i appears in the attaching word of the 2-cell with total exponent 2, which means that each $\Delta_{\alpha\beta}$ is homotopic to the map $z \mapsto z^2$, of degree 2. Since $d_2(1) = (2, \cdots, 2)$, we have d_2 injective and hence $H_2(N_g) = 0$. If we change the basis for $\mathbb{Z}^g$ by replacing the last standard basis element $(0, \cdots, 0, 1)$ by $(1, \cdots, 1)$, we see that $H_1(N_g) \approx \mathbb{Z}^{g-1} \oplus \mathbb{Z}_2$.

Example 3. Let X be obtained from $S^1 \vee S^1$ by attaching two 2-cells by the words a^5b^{-3} and $b^3(ab)^{-2}$. Then $d_2 : \mathbb{Z}^2 \to \mathbb{Z}^2$ has matrix

$$\begin{pmatrix} 5 & -2 \\ -3 & 1 \end{pmatrix}$$

The two columns come from abelianizing a^5b^{-3} and $b^3(ab)^{-2}$ to $5a - 3b$ and $-2a + b$, in additive notation. The matrix has determinant -1, hence is an isomorphism. Therefore $\widetilde{H}_i(X) = 0$ for all i. Such a space X is called **acyclic**.

It is interesting to note that this acyclic complex X is not contractible, since its fundamental group is nontrivial. Namely, $\pi_1(X)$ has the presentation $\langle\, a, b \mid a^5b^{-3}, b^3(ab)^{-2} \,\rangle$, and there is a nontrivial homomorphism from this group to the group G of rotational symmetries of a regular dodecahedron, sending a to the rotation ρ_a through angle $2\pi/5$ about the axis through the center of a pentagonal face, and b to the rotation ρ_b through angle $2\pi/3$ about the axis through a vertex of this face. The composition $\rho_a\rho_b$ is a rotation through angle π about the axis through the midpoint of an edge abutting this vertex. Thus the relations $a^5 = b^3 = (ab)^2$ defining $\pi_1(X)$ become $\rho_a^5 = \rho_b^3 = (\rho_a\rho_b)^2 = 1$ in G, which means there is a well-defined homomorphism $\rho : \pi_1(X) \to G$ sending a to ρ_a and b to ρ_b. It is not hard to see that G is generated by ρ_a and ρ_b, so ρ is surjective. With more work one can compute that the kernel of ρ is $\mathbb{Z}_2$, generated by the element $a^5 = b^3 = (ab)^2$, and this $\mathbb{Z}_2$ is in fact the center of $\pi_1(X)$. In particular, $\pi_1(X)$ has order 120 since G has order 60.

Example 4: Moore Spaces. Given an abelian group G and an integer $n \geq 1$, we shall construct a CW complex X such that $H_n(X) \approx G$ and $\widetilde{H}_i(X) = 0$ for $i \neq n$. Such a space is called a **Moore space**, commonly written $M(G, n)$ to indicate the dependence on G and n. It is probably best for the definition of a Moore space to include the condition that $M(G, n)$ be simply-connected if $n > 1$. The ones we construct will have this property.

As an easy special case, when $G = \mathbb{Z}_m$ we can take X to be S^n with a cell e^{n+1} attached by a map $S^n \to S^n$ of degree m. More generally, any finitely generated G can be realized by taking wedge sums of examples of this type for finite cyclic summands of G, together with copies of S^n for infinite cyclic summands of G.

In the general nonfinitely generated case let $F \to G$ be a homomorphism of a free abelian group F onto G, sending a basis for F onto some set of generators of G. The kernel K of this homomorphism is a subgroup of a free abelian group, hence is itself free abelian. Choose bases $\{x_\alpha\}$ for F and $\{y_\beta\}$ for K, and write $y_\beta = \sum_\alpha d_{\beta\alpha} x_\alpha$. Let $X^n = \bigvee_\alpha S^n_\alpha$, so $H_n(X^n) \approx F$ via Corollary 2.18. We shall obtain X from X^n by attaching cells e^{n+1}_β via maps $f_\beta : S^n \to X^n$ such that the composition of f_β with the projection onto the summand S^n_α has degree $d_{\beta\alpha}$. Then the cellular boundary map d_{n+1} will be the inclusion $K \hookrightarrow F$, hence X will have the desired homology groups.

The construction of f_β generalizes the earlier construction of a map $S^n \to S^n$ of given degree. Let f_β map the complement of $\sum_\alpha |d_{\beta\alpha}|$ disjoint balls in S^n to the 0-cell of X^n, while mapping $|d_{\beta\alpha}|$ of the balls onto the summand S^n_α by maps of degree $+1$ if $d_{\beta\alpha} > 0$, or degree -1 in the opposite case.

Example 5. By taking a wedge sum of the Moore spaces constructed in Example 4 for varying n we obtain a connected CW complex with any prescribed sequence of homology groups in dimensions $1, 2, 3, \cdots$.

Example 6: $\mathbb{R}\mathrm{P}^n$. As we saw in Chapter 0, $\mathbb{R}\mathrm{P}^n$ has a CW structure with one cell in each dimension $\le n$, the attaching map for the cell e^{k+1} being the 2-sheeted covering projection $\varphi : S^k \to \mathbb{R}\mathrm{P}^k$. Computing the boundary map d_{k+1} is the same as computing the degree of the composition $S^k \xrightarrow{\varphi} \mathbb{R}\mathrm{P}^k \xrightarrow{q} \mathbb{R}\mathrm{P}^k/\mathbb{R}\mathrm{P}^{k-1} = S^k$, with q the quotient map. The map $q\varphi$ is a homeomorphism when restricted to each component of $S^k - S^{k-1}$, and these two homeomorphisms are obtained from each other by pre-composing with the antipodal map of S^k, which has degree $(-1)^{k+1}$. Hence $deg(q\varphi) = deg(\mathbb{1}) + deg(-\mathbb{1}) = 1 + (-1)^{k+1}$, and so d_{k+1} is either 0 or multiplication by 2 according to whether k is even or odd. Thus if n is even the cellular chain complex for $\mathbb{R}\mathrm{P}^n$ is

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

while for n odd it is

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

From this it follows that

$$H_k(\mathbb{R}\mathrm{P}^n) = \begin{cases} \mathbb{Z} & \text{for } k = 0 \text{ and for } k = n \text{ odd} \\ \mathbb{Z}_2 & \text{for } k \text{ odd, } 0 < k < n \\ 0 & \text{otherwise.} \end{cases}$$

Example 7: Lens Spaces. This example is somewhat more complicated. Given an integer $m > 1$ and integers $\ell_1, \cdots, \ell_n$ relatively prime to m, define the **lens space** $L =$

$L_m(\ell_1,\cdots,\ell_n)$ to be the quotient space $S^{2n-1}/\mathbb{Z}_m$ of the unit sphere $S^{2n-1}\subset\mathbb{C}^n$ under the action of $\mathbb{Z}_m$ generated by the rotation $\rho(z_1,\cdots,z_n)=(e^{2\pi i\ell_1/m}z_1,\cdots,e^{2\pi i\ell_n/m}z_n)$, rotating the j^{th} $\mathbb{C}$ factor of $\mathbb{C}^n$ by the angle $2\pi\ell_j/m$. In particular, when $m=2$, ρ is the antipodal map, so $L=\mathbb{R}\mathrm{P}^{2n-1}$ in this case. In the general case, the projection $S^{2n-1}\to L$ is a covering space since the action of $\mathbb{Z}_m$ on S^{2n-1} is free: Only the identity element fixes any point of S^{2n-1} since each point of S^{2n-1} has some coordinate z_j nonzero and then $e^{2\pi ik\ell_j/m}z_j\neq z_j$ for $0<k<m$, as a result of the assumption that ℓ_j is relatively prime to m.

We shall construct a CW structure on L with one cell e^k in each dimension $k\le 2n-1$ and show that the resulting cellular chain complex is

$$0\longrightarrow\mathbb{Z}\xrightarrow{0}\mathbb{Z}\xrightarrow{m}\mathbb{Z}\xrightarrow{0}\cdots\xrightarrow{0}\mathbb{Z}\xrightarrow{m}\mathbb{Z}\xrightarrow{0}\mathbb{Z}\longrightarrow 0$$

with boundary maps alternately 0 and multiplication by m. Hence

$$H_k\big(L_m(\ell_1,\cdots,\ell_n)\big)=\begin{cases}\mathbb{Z} & \text{for } k=0,2n-1\\ \mathbb{Z}_m & \text{for } k \text{ odd},\ 0<k<2n-1\\ 0 & \text{otherwise.}\end{cases}$$

To obtain the CW structure, first subdivide the unit circle C in the n^{th} $\mathbb{C}$ factor of $\mathbb{C}^n$ by taking the points $e^{2\pi ij/m}\in C$ as vertices, $j=1,\cdots,m$. Joining the j^{th} vertex of C to the unit sphere $S^{2n-3}\subset\mathbb{C}^{n-1}$ by arcs of great circles in S^{2n-1} yields a $(2n-2)$-dimensional ball B_j^{2n-2} bounded by S^{2n-3}. Specifically, B_j^{2n-2} consists of the points $\cos\theta\,(0,\cdots,0,e^{2\pi ij/m})+\sin\theta\,(z_1,\cdots,z_{n-1},0)$ for $0\le\theta\le\pi/2$. Similarly, joining the j^{th} edge of C to S^{2n-3} gives a ball B_j^{2n-1} bounded by B_j^{2n-2} and B_{j+1}^{2n-2}, subscripts being taken $\mod m$. The rotation ρ carries S^{2n-3} to itself and rotates C by the angle $2\pi\ell_n/m$, hence ρ permutes the B_j^{2n-2}'s and the B_j^{2n-1}'s. A suitable power of ρ, namely ρ^r where $r\ell_n\equiv 1 \mod m$, takes each B_j^{2n-2} and B_j^{2n-1} to the next one. Since ρ^r has order m, it is also a generator of the rotation group $\mathbb{Z}_m$, and hence we may obtain L as the quotient of one B_j^{2n-1} by identifying its two faces B_j^{2n-2} and B_{j+1}^{2n-2} together via ρ^r.

In particular, when $n=2$, B_j^{2n-1} is a lens-shaped 3-ball and L is obtained from this ball by identifying its two curved disk faces via ρ^r, which may be described as the composition of the reflection across the plane containing the rim of the lens, taking one face of the lens to the other, followed by a rotation of this face through the angle $2\pi\ell/m$ where $\ell=r\ell_1$. The figure illustrates the case $(m,\ell)=(7,2)$, with the two dots indicating a typical pair of identified points in the upper and

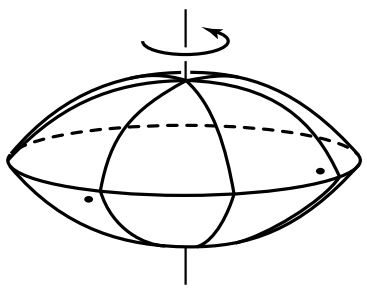

lower faces of the lens. Since the lens space L is determined by the rotation angle $2\pi\ell/m$, it is conveniently written $L_{\ell/m}$. Clearly only the $\mod m$ value of ℓ matters. It is a classical theorem of Reidemeister from the 1930's that $L_{\ell/m}$ is homeomorphic to $L_{\ell'/m'}$ iff $m' = m$ and $\ell' \equiv \pm\ell^{\pm 1} \mod m$. For example, when $m = 7$ there are only two distinct lens spaces $L_{1/7}$ and $L_{2/7}$. The 'if' part of this theorem is easy: Reflecting the lens through a mirror shows that $L_{\ell/m} \approx L_{-\ell/m}$, and by interchanging the roles of the two $\mathbb{C}$ factors of $\mathbb{C}^2$ one obtains $L_{\ell/m} \approx L_{\ell^{-1}/m}$. In the converse direction, $L_{\ell/m} \approx L_{\ell'/m'}$ clearly implies $m = m'$ since $\pi_1(L_{\ell/m}) \approx \mathbb{Z}_m$. The rest of the theorem takes considerably more work, however, involving either special 3-dimensional techniques or more algebraic methods which generalize to classify the higher-dimensional lens spaces as well. The latter approach is explained in [Cohen].

Returning to the construction of a CW structure on $L_m(\ell_1, \cdots, \ell_n)$, observe that the $(2n-3)$-dimensional lens space $L_m(\ell_1, \cdots, \ell_{n-1})$ sits in $L_m(\ell_1, \cdots, \ell_n)$ as the quotient of S^{2n-3}, and $L_m(\ell_1, \cdots, \ell_n)$ is obtained from this subspace by attaching two cells, of dimensions $2n-2$ and $2n-1$, coming from the interiors of B^{2n-1}_j and its two identified faces B^{2n-2}_j and B^{2n-2}_{j+1}. Inductively this gives a CW structure on $L_m(\ell_1, \cdots, \ell_n)$ with one cell e^k in each dimension $k \leq 2n-1$.

The boundary maps in the associated cellular chain complex are computed as follows. The first one, d_{2n-1}, is zero since the identification of the two faces of B^{2n-1}_j is via a reflection (degree -1) across B^{2n-1}_j fixing S^{2n-3}, followed by a rotation (degree $+1$), so $d_{2n-1}(e^{2n-1}) = e^{2n-2} - e^{2n-2} = 0$. The next boundary map d_{2n-2} takes e^{2n-2} to me^{2n-3} since the attaching map for e^{2n-2} is the quotient map $S^{2n-3} \to L_m(\ell_1, \cdots, \ell_{n-1})$ and the balls B^{2n-3}_j in S^{2n-3} which project down onto e^{2n-3} are permuted cyclically by the rotation ρ of degree $+1$. Inductively, the subsequent boundary maps d_k then alternate between 0 and multiplication by m.

Also of interest are the infinite-dimensional lens spaces $L_m(q_1, q_2, \cdots) = S^\infty/\mathbb{Z}_m$ defined in the same way as in the finite-dimensional case, starting from a sequence of integers $q_1, q_2, \cdots$ relatively prime to m. The space $L_m(q_1, q_2, \cdots)$ is the union of the increasing sequence of finite-dimensional lens spaces $L_m(q_1, \cdots, q_n)$ for $n = 1, 2, \ldots$, each of which is a subcomplex of the next in the cell structure we have just constructed, so $L_m(q_1, q_2, \cdots)$ is also a CW complex. Its cellular chain complex consists of a $\mathbb{Z}$ in each dimension with boundary maps alternately 0 and m, so its reduced homology consists of a $\mathbb{Z}_m$ in each odd dimension.

In the terminology of §1.4b, the infinite-dimensional lens space $L_m(q_1, q_2, \cdots)$ is a $K(\mathbb{Z}_m, 1)$ since its universal cover S^∞ is contractible, as we showed there. By Theo-

rem 1.32 the homotopy type of $L_m(q_1, q_2, \cdots)$ depends only on m, and not on the q_i's. This is not true in the finite-dimensional case, when two lens spaces $L_m(q_1, \cdots, q_n)$ and $L_m(q_1', \cdots, q_n')$ have the same homotopy type iff $q_1 \cdots q_n \equiv \pm k^n q_1' \cdots q_n' \mod m$ for some integer k. A proof of this is outlined in exercise 2 in §3.4e and exercise 45 in §4.1. For example, the three-dimensional lens spaces $L_{1/5}$ and $L_{2/5}$ are not homotopy equivalent, though they have the same fundamental group and the same homology groups. On the other hand, $L_{1/7}$ and $L_{2/7}$ are homotopy equivalent but not homeomorphic.

A consequence of our calculation of the homology of infinite-dimensional lens spaces is the following fact: *If for a given group G there exists a finite-dimensional CW complex $K(G,1)$, then G must be torsionfree.* For if G had torsion, it would have a subgroup $\mathbb{Z}_m$ for some $m > 1$, and the covering space of a $K(G,1)$ corresponding to this subgroup of $G = \pi_1(K(G,1))$ would be a $K(\mathbb{Z}_m, 1)$. If the $K(G,1)$ were a finite-dimensional CW complex, so would be the $K(\mathbb{Z}_m, 1)$, so the homology of the $K(\mathbb{Z}_m, 1)$ would be nonzero in only finitely many dimensions. But this contradicts our calculation of the homology of infinite-dimensional lens spaces, since the homotopy type of $K(\mathbb{Z}_m, 1)$ depends only on m by Theorem 1.32.

Euler Characteristic

This is an application of cellular homology to something quite classical indeed. For a finite CW complex X, the **Euler characteristic** $\chi(X)$ is defined to be the alternating sum $\sum_n (-1)^n c_n$ where c_n is the number of n-cells of X, generalizing the familiar formula *vertices − edges + faces* for 2-dimensional simplicial complexes. The following result shows that $\chi(X)$ depends only on the homotopy type of X, and in particular is independent of the choice of CW structure on X.

Theorem 2.27. $\chi(X) = \sum_n (-1)^n \operatorname{rank}(H_n(X))$.

Here the **rank** of a finitely generated abelian group is the number of $\mathbb{Z}$ summands when the group is expressed as a direct sum of cyclic groups. We shall need the following fact about rank, whose proof we leave as an exercise: If $0 \to A \to B \to C \to 0$ is a short exact sequence of finitely generated abelian groups, then $\operatorname{rank}(B) = \operatorname{rank}(A) + \operatorname{rank}(C)$.

Proof of 2.27: This is purely algebraic. Let

$$0 \longrightarrow C_k \xrightarrow{d_k} C_{k-1} \longrightarrow \cdots \longrightarrow C_1 \xrightarrow{d_1} C_0 \longrightarrow 0$$

be a chain complex of finitely generated abelian groups, with cycles $Z_n = Ker(d_n)$, boundaries $B_n = Im(d_{n+1})$, and homology $H_n = Z_n/B_n$. Thus we have short exact sequences

$0 \to Z_n \to C_n \to B_{n-1} \to 0$ and $0 \to B_n \to Z_n \to H_n \to 0$, hence

$$\operatorname{rank}(C_n) = \operatorname{rank}(Z_n) + \operatorname{rank}(B_{n-1})$$
$$\operatorname{rank}(Z_n) = \operatorname{rank}(B_n) + \operatorname{rank}(H_n)$$

Now substitute the second equation into the first, multiply the resulting equation by $(-1)^n$, and sum over n to get $\sum_n (-1)^n \operatorname{rank}(C_n) = \sum_n (-1)^n \operatorname{rank}(H_n)$. Applying this with $C_n = H_n(X^n, X^{n-1})$ then gives the theorem. □

For example, the surfaces M_g and N_g have Euler characteristics $\chi(M_g) = 2-2g$ and $\chi(N_g) = 2-g$. Thus all the orientable surfaces M_g are distinguished from each other by their Euler characteristics, as are the non-orientable surfaces N_g, and there are only the relations $\chi(M_g) = \chi(N_{2g})$.

Split Exact Sequences

Suppose one has a retraction $r : X \to A$, so $ri = \mathbb{1}$ where $i : A \to X$ is the inclusion. Then the induced map $i_* : H_n(A) \to H_n(X)$ is injective since $r_* i_* = \mathbb{1}$, and hence the boundary maps in the long exact sequence for (X, A) are zero. This means that the long exact sequence breaks up into short exact sequences:

$$0 \longrightarrow H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \longrightarrow 0$$

The relation $r_* i_* = \mathbb{1}$ actually gives more information than this, by the following piece of elementary algebra:

Splitting Lemma. *For a short exact sequence of abelian groups*

$$0 \longrightarrow A \xrightarrow{i} B \xrightarrow{j} C \longrightarrow 0$$

the following statements are equivalent:

(a) *There is a homomorphism $p : B \to A$ such that $pi = \mathbb{1} : A \to A$.*

(b) *There is a homomorphism $s : C \to B$ such that $js = \mathbb{1} : C \to C$.*

(c) *There is an isomorphism $B \approx A \oplus C$ making a commutative diagram*

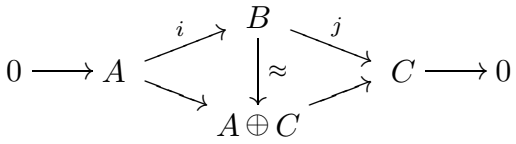

where the maps in the lower row are the obvious ones, $a \mapsto (a, 0)$, and $(a, c) \mapsto c$.

If these conditions are satisfied the exact sequence is said to **split**. Note that (c) is symmetric: There is no essential difference between the roles of A and C.

Sketch of Proof: For (a) $\Rightarrow$ (c) one checks that the map $B \to A \oplus C$, $b \mapsto (p(b), j(b))$, is an isomorphism with the desired properties. For (b) $\Rightarrow$ (c) one uses instead the map $A \oplus C \to B$, $(a, c) \mapsto i(a) + s(c)$. The opposite implications (c) $\Rightarrow$ (a) and (c) $\Rightarrow$ (b) are fairly obvious. If one wants to show (b) $\Rightarrow$ (a) directly, one can define $p(b) = i^{-1}(b - sj(b))$. Further details are left to the reader. □

Except for the implication (b) $\Rightarrow$ (c), the proof works equally well for nonabelian groups. In the nonabelian case, (b) is definitely weaker than (a) and (c), and short exact sequences satisfying (b) only determine B as a semi-direct product of A and C. The difficulty is that $s(C)$ might not be a normal subgroup of B. In the nonabelian case one defines 'splitting' to mean that (b) is satisfied.

In both the abelian and nonabelian contexts, if C is free then every exact sequence $0 \to A \xrightarrow{i} B \xrightarrow{j} C \to 0$ splits, since one can define $s : C \to B$ by choosing a basis $\{c_\alpha\}$ for C and letting $s(c_\alpha)$ be any element $b_\alpha \in B$ such that $j(b_\alpha) = c_\alpha$. The converse is also true: If every short exact sequence ending in C splits, then C is free. This is because for every C there is a short exact sequence $0 \to A \to B \to C \to 0$ with B free — choose generators for C and let B have a basis in one-to-one correspondence with these generators, then let $B \to C$ send each basis element to the corresponding generator — so if this sequence $0 \to A \to B \to C \to 0$ splits, C is isomorphic to a subgroup of a free group, hence is free.

From the lemma and the remarks preceding it we deduce that a retraction $r : X \to A$ gives a splitting $H_n(X) \approx H_n(A) \oplus H_n(X, A)$. This can be used to show the nonexistence of such a retraction in some cases, for example in the situation of the Brouwer fixed point theorem, where a retraction $D^n \to S^{n-1}$ would give a splitting $H_{n-1}(D^n) \approx H_{n-1}(S^{n-1}) \oplus H_{n-1}(D^n, S^{n-1})$. For a somewhat more subtle example, consider the mapping cylinder M_f of a degree m map $f : S^n \to S^n$. For the inclusion $i : S^n \hookrightarrow M_f$ of the domain S^n, the homology long exact sequence reduces to a short exact sequence

$$\begin{array}{ccccccccc} 0 & \longrightarrow & H_n(S^n) & \longrightarrow & H_n(M_f) & \longrightarrow & H_n(M_f, S^n) & \longrightarrow & 0 \\ & & \| & & \| & & \| & & \\ 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{m} & \mathbb{Z} & \longrightarrow & \mathbb{Z}_m & \longrightarrow & 0 \end{array}$$

If $m > 1$ this sequence does not split since $\mathbb{Z}$ is not isomorphic to $\mathbb{Z} \oplus \mathbb{Z}_m$, so S^n is not a retract of M_f. In the simplest case of the degree 2 map $S^1 \to S^1$, $z \mapsto z^2$, this says the Möbius band does not retract onto its boundary circle.

Mayer-Vietoris Sequences

Suppose a space X is decomposed as the union of two open sets A and B, or more generally A and B could be subspaces whose interiors cover X. Then the **Mayer-Vietoris sequence** is a long exact sequence

$$\cdots \longrightarrow H_n(A\cap B) \xrightarrow{\Phi} H_n(A)\oplus H_n(B) \xrightarrow{\Psi} H_n(X) \xrightarrow{\partial} H_{n-1}(A\cap B) \longrightarrow \cdots$$
$$\cdots \longrightarrow H_0(X) \longrightarrow 0$$

This is easy to derive from the machinery of §2.1. Let $C_n(A+B)$ be the subgroup of $C_n(X)$ consisting of chains which are sums of singular simplices in either A or B. The usual boundary map $\partial : C_n(X) \to C_{n-1}(X)$ takes $C_n(A+B)$ to $C_{n-1}(A+B)$, so the $C_n(A+B)$'s form a chain complex. According to Proposition 2.14, the inclusions $C_n(A+B) \hookrightarrow C_n(X)$ induce isomorphisms on homology groups. The Mayer-Vietoris sequence is then the long exact sequence of homology groups associated to the short exact sequence of chain complexes formed by the short exact sequences

$$0 \longrightarrow C_n(A\cap B) \xrightarrow{\varphi} C_n(A)\oplus C_n(B) \xrightarrow{\psi} C_n(A+B) \longrightarrow 0$$

where $\varphi(x) = (x,-x)$ and $\psi(x,y) = x+y$. The exactness of this short exact sequence can be checked as follows. First, $Ker(\varphi) = 0$ since a chain in $A\cap B$ which is zero as a chain in A (or in B) must be the zero chain. Next, $Im(\varphi) \subset Ker(\psi)$ since $\psi\varphi = 0$, and $Ker(\psi) \subset Im(\varphi)$ since for a pair $(x,y) \in C_n(A)\oplus C_n(B)$ the condition $x+y=0$ implies $x=-y$, so x is a chain in both A and B, i.e., $x \in C_n(A\cap B)$ and $(x,y) = (x,-x) \in Im(\varphi)$. Finally, exactness at $C_n(A+B)$ is immediate from the definition of $C_n(A+B)$.

There is also a formally identical Mayer-Vietoris sequence for reduced homology groups, obtained by augmenting the previous short exact sequence of chain complexes in the obvious way:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & C_0(A\cap B) & \xrightarrow{\varphi} & C_0(A)\oplus C_0(B) & \xrightarrow{\psi} & C_0(A+B) & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle\varepsilon} & & \downarrow{\scriptstyle\varepsilon\oplus\varepsilon} & & \downarrow{\scriptstyle\varepsilon} & & \\
0 & \longrightarrow & \mathbb{Z} & \xrightarrow{\varphi} & \mathbb{Z}\oplus\mathbb{Z} & \xrightarrow{\psi} & \mathbb{Z} & \longrightarrow & 0
\end{array}$$

The Mayer-Vietoris sequences also hold for decompositions $X = A\cup B$ such that A and B are deformation retracts of open neighborhoods U and V with $U\cap V$ deformation retracting onto $A\cap B$. This applies for example if X is a CW complex and A and B are subcomplexes, using the neighborhoods $N(A)$ and $N(B)$ constructed in the Appendix.

Example. Take $X = S^n$ with A and B the northern and southern hemispheres, so that $A \cap B = S^{n-1}$. Then in the reduced Mayer-Vietoris sequence the terms $\widetilde{H}_i(A) \oplus \widetilde{H}_i(B)$ are zero, so we obtain isomorphisms $\widetilde{H}_i(S^n) \approx \widetilde{H}_{i-1}(S^{n-1})$. This gives another way of calculating the homology groups of S^n by induction.

As this example illustrates, one often has a choice of using either Mayer-Vietoris sequences or long exact sequences of pairs, whichever is more convenient.

The Mayer-Vietoris sequence can be thought of as the homology version of the van Kampen theorem since if $A \cap B$ is path-connected, the H_1 terms of the reduced Mayer-Vietoris sequence yield an isomorphism $H_1(X) \approx H_1(A) \oplus H_1(B)/Im(\Phi)$. This is exactly the abelianized statement of the van Kampen theorem, and H_1 is the abelianization of π_1 for path-connected spaces, as we show in §2.4a.

A relative form of the Mayer-Vietoris sequence is sometimes useful. If one has a pair of spaces $(X, Y) = (A \cup B, C \cup D)$ such that X is the union of the interiors of A and B, and Y is the union of the interiors of C and D, with $C \subset A$ and $D \subset B$, then there is a relative Mayer-Vietoris sequence

$$\cdots \longrightarrow H_n(A \cap B, C \cap D) \xrightarrow{\Phi} H_n(A, C) \oplus H_n(B, D) \xrightarrow{\Psi} H_n(X, Y) \xrightarrow{\partial} H_{n-1}(A \cap B, C \cap D) \longrightarrow \cdots$$

To derive this, consider the commutative diagram

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & C_n(C \cap D) & \xrightarrow{\varphi} & C_n(C) \oplus C_n(D) & \xrightarrow{\psi} & C_n(C + D) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & C_n(A \cap B) & \xrightarrow{\varphi} & C_n(A) \oplus C_n(B) & \xrightarrow{\psi} & C_n(A + B) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & C_n(A \cap B, C \cap D) & \xrightarrow{\varphi} & C_n(A, C) \oplus C_n(B, D) & \xrightarrow{\psi} & C_n(A + B, C + D) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

where $C_n(C + D) \subset C_n(Y)$ consists of sums of chains in C and chains in D, $C_n(A + B)$ is defined similarly, and $C_n(A + B, C + D)$ is the quotient of these two groups. Thus the three columns of the diagram are exact. We have seen that the first two rows are exact, and we claim that the third row is exact also, with the maps φ and ψ induced from the φ and ψ in the second row. Since $\psi\varphi = 0$ in the second row, this holds also in the third row, so the third row is at least a chain complex. Viewing the three rows as chain complexes, the diagram then represents a short exact sequence of chain complexes. The associated long exact sequence of homology groups has two out of every three terms zero since the

first two rows of the diagram are exact. Hence the remaining homology groups are zero too and the third row is exact.

The third column maps to $0 \to C_n(Y) \to C_n(X) \to C_n(X,Y) \to 0$, inducing maps of homology groups which are isomorphisms for the X and Y terms as we have seen above. So by the five-lemma the maps $C_n(A+B, C+D) \to C_n(X,Y)$ also induce isomorphisms on homology. The relative Mayer-Vietoris sequence is then the long exact sequence of homology groups associated to the short exact sequence of chain complexes given by the third row of the diagram.

Homology with Coefficients

There is an easy generalization of the homology theory we have considered so far which behaves in a very similar fashion and sometimes offers technical advantages. The generalization is defined using chains of the form $\sum_i n_i\sigma_i$ where $\sigma_i : \Delta^n \to X$ is a singular n-simplex in X as before, but now the coefficients n_i are taken to lie in a fixed abelian group G rather than $\mathbb{Z}$. Such n-chains form an abelian group $C_n(X;G)$, and in the relative case one defines $C_n(X,A;G) = C_n(X;G)/C_n(A;G)$. The old formula for the boundary maps ∂ can still be used for arbitrary G, namely $\partial(\sum_i n_i\sigma_i) = \sum_{i,j}(-1)^j n_i\sigma_i\big|[v_0,\cdots,\widehat{v}_j,\cdots,v_n]$, and we again have $\partial^2 = 0$ by the same calculation as before. So the groups $C_n(X;G)$ and $C_n(X,A;G)$ form chain complexes, whose homology groups $H_n(X;G)$ and $H_n(X,A;G)$ are called **homology groups with coefficients in G**. Reduced groups $\widetilde{H}_n(X;G)$ are defined via the augmented chain complex $\cdots \to C_0(X;G) \xrightarrow{\varepsilon} G \to 0$ with ε defined as before, by summing coefficients.

All the theory we developed in §2.1 for $\mathbb{Z}$ coefficients carries over directly to general coefficient groups G, with no change in the proofs. The equivalence of simplicial and singular homology with G coefficients also carries over without complication. The same is true for Mayer-Vietoris sequences.

Differences between $H_n(X;G)$ and $H_n(X)$ begin to appear only when one starts making calculations. In the simplest case that X is a point, the method used for $G = \mathbb{Z}$ shows that $H_n(X;G)$ is G for $n = 0$ and 0 for $n > 0$. From this it follows just as for $G = \mathbb{Z}$ that $\widetilde{H}_n(S^k;G)$ is G for $n = k$ and 0 otherwise.

Cellular homology also generalizes to homology with coefficients, with the cellular chain group $H_n(X^n, X^{n-1})$ replaced by $H_n(X^n, X^{n-1};G)$ which is a direct sum of G's, one for each n-cell. The proof that the cellular homology groups $H_n^{CW}(X)$ agree with singular homology $H_n(X)$ extends immediately to give $H_n^{CW}(X;G) \approx H_n(X;G)$. The cellular boundary maps are given by the same formula as for $\mathbb{Z}$ coefficients, $d_n(\sum_\alpha n_\alpha e_\alpha^n) =$

$\sum_{\alpha,\beta} d_{\alpha\beta} n_\alpha e^{n-1}_\beta$ where $d_{\alpha\beta}$ is the degree of a certain map $S^{n-1} \to S^{n-1}$. This follows from:

Lemma 2.28. *If $f : S^k \to S^k$ has degree m, then $f_* : H_k(S^k; G) \to H_k(S^k; G)$ is multiplication by m.*

Proof: As a preliminary observation, note that a homomorphism $\varphi : G_1 \to G_2$ induces maps $\varphi_\sharp : C_n(X, A; G_1) \to C_n(X, A; G_2)$ commuting with the boundary maps, and hence there are induced homomorphisms $\varphi_* : H_n(X, A; G_1) \to H_n(X, A; G_2)$. These have various naturality properties. For example, they give a commutative diagram mapping the long exact sequence of homology for the pair (X, A) with G_1 coefficients to the corresponding sequence with G_2 coefficients. Also, the maps φ_* commute with homomorphisms f_* induced by maps $f : (X, A) \to (Y, B)$.

Now let $f : S^k \to S^k$ have degree m and let $\varphi : \mathbb{Z} \to G$ take 1 to a given element $g \in G$. Then we have a commutative diagram

$$\begin{array}{ccccccc}
\mathbb{Z} & \approx & \widetilde{H}_k(S^k; \mathbb{Z}) & \xrightarrow{f_*} & \widetilde{H}_k(S^k; \mathbb{Z}) & \approx & \mathbb{Z} \\
\downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle \varphi_*} & & \downarrow{\scriptstyle \varphi_*} & & \downarrow{\scriptstyle \varphi} \\
G & \approx & \widetilde{H}_k(S^k; G) & \xrightarrow{f_*} & \widetilde{H}_k(S^k; G) & \approx & G
\end{array}$$

where commutativity of the outer two squares comes from the inductive calculation of these homology groups, reducing to the case $k = 0$ when the commutativity is obvious.

Since the diagram commutes, the assumption that the map across the top takes 1 to m implies that the map across the bottom takes g to mg. □

An instructive example is $\mathbb{R}\mathrm{P}^n$ with the coefficient group G chosen to be a field F. The cellular chain complex is

$$\cdots \xrightarrow{0} F \xrightarrow{2} F \xrightarrow{0} F \xrightarrow{2} F \xrightarrow{0} F \longrightarrow 0$$

Hence if F has characteristic 2, for example if $F = \mathbb{Z}_2$, then $H_k(\mathbb{R}\mathrm{P}^n; F) \approx F$ for $0 \le k \le n$, a more uniform answer than with $\mathbb{Z}$ coefficients. On the other hand if F has characteristic different from 2, then the boundary maps $F \xrightarrow{2} F$ are isomorphisms, hence $H_k(\mathbb{R}\mathrm{P}^n; F)$ is F for $k = 0$ and for $k = n$ odd, and is zero otherwise.

This example shows that it is possible for $H_n(X; G)$ to be nonzero when $H_n(X; \mathbb{Z})$ is zero. This can be useful when considering maps between spaces. For example, let X be a Moore space $M(\mathbb{Z}_m, n)$, obtained from S^n by attaching a cell e^{n+1} by a map of degree m. The quotient map $q : X \to X/S^n = S^{n+1}$ induces trivial homomorphisms on reduced

homology with $\mathbb{Z}$ coefficients since the nonzero reduced homology groups of X and S^{n+1} occur in different dimensions. But with $\mathbb{Z}_m$ coefficients the story is different, as we can see by considering the long exact sequence of the pair (X, S^n), which contains the segment

$$0 = \widetilde{H}_{n+1}(S^n;\mathbb{Z}_m) \longrightarrow \widetilde{H}_{n+1}(X;\mathbb{Z}_m) \xrightarrow{q_*} \widetilde{H}_{n+1}(X/S^n;\mathbb{Z}_m)$$

Exactness says that q_* is injective, hence nonzero since $\widetilde{H}_{n+1}(X;\mathbb{Z}_m)$ is $\mathbb{Z}_m$, the cellular boundary map $H_{n+1}(X^{n+1}, X^n;\mathbb{Z}_m) \to H_n(X^n, X^{n-1};\mathbb{Z}_m)$ being $\mathbb{Z}_m \xrightarrow{m} \mathbb{Z}_m$. Thus homology with $\mathbb{Z}_m$ coefficients tells us that q is not homotopic to a constant map, which we would not know using only $\mathbb{Z}$ coefficients.

In §3.4a we will see that there is a general algebraic formula expressing homology with arbitrary coefficients in terms of homology with $\mathbb{Z}$ coefficients. Some easy special cases which give much of the flavor of the general result are included in the exercises.

Exercises

1. Prove the Brouwer fixed point theorem for maps $f : D^n \to D^n$ by applying degree theory to the map $S^n \to S^n$ which sends both hemispheres to the southern hemisphere via f. [This was Brouwer's original proof.]

2. Given a map $f : S^{2n} \to S^{2n}$, show that there is some point $x \in S^{2n}$ with either $f(x) = x$ or $f(x) = -x$. Deduce that every map $\mathbb{R}\mathrm{P}^{2n} \to \mathbb{R}\mathrm{P}^{2n}$ has a fixed point. Construct maps $\mathbb{R}\mathrm{P}^{2n-1} \to \mathbb{R}\mathrm{P}^{2n-1}$ without fixed points, e.g., from linear transformations $\mathbb{R}^{2n} \to \mathbb{R}^{2n}$ without eigenvectors.

3. Show that any two reflections of S^n across different hyperplanes are homotopic, in fact homotopic through reflections.

4. Show that any map $S^n \to S^n$ can be homotoped to have a fixed point if $n > 0$.

5. For an invertible linear transformation $f : \mathbb{R}^n \to \mathbb{R}^n$ show that the induced map on $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\}) \approx \widetilde{H}_{n-1}(\mathbb{R}^n - \{0\}) \approx \mathbb{Z}$ is $\mathbb{1}$ or $-\mathbb{1}$ according to whether the determinant of f is positive or negative. [Show that the matrix of f can be joined by a path of invertible matrices to a diagonal matrix with ± 1's on the diagonal.]

6. Use the chain maps in exercise 14 of §2.1 to give a direct proof that $deg(Sf) = deg(f)$ for $f : S^n \to S^n$.

7. Compute the homology groups of the 2-complexes described in exercises 4, 5, and 8 in §1.2.

8. In an exercise for §1.2 we described a 3-dimensional CW complex obtained from the cube I^3 by identifying opposite faces via a one-quarter twist. Compute the homology groups of this complex.

9. Show that the quotient map $S^1 \times S^1 \to S^2$ collapsing the subspace $S^1 \vee S^1$ to a point induces an isomorphism on H_2. Show also, using covering spaces, that any map $S^2 \to S^1 \times S^1$ is homotopic to a constant.

10. Show that if X is a CW complex then $H_n(X^n)$ is isomorphic to the kernel of the cellular boundary map $H_n(X^n, X^{n-1}) \to H_{n-1}(X^{n-1}, X^{n-2})$, hence is free.

11. Consider $\Delta^n = [v_0, \cdots, v_n]$ with its natural Δ-complex structure with k-simplices $[v_{i_0}, \cdots, v_{i_k}]$ for $i_0 < \cdots < i_k$. Compute the ranks of the simplicial (or cellular) chain groups $\Delta_i(\Delta^n)$ and the subgroups of cycles and boundaries. [Hint: Pascal's triangle.] Apply this, using also the previous problem, to show that the k-skeleton of Δ^n has homology groups $\widetilde{H}_i\big((\Delta^n)^k\big)$ equal to 0 for $i < k$, and free of rank $\binom{n}{k+1}$ for $i = k$.

12. Show the isomorphism $H_n^{CW}(X) \approx H_n(X)$ is natural in the following sense: A cellular map $f : X \to Y$, i.e., one satisfying $f(X^n) \subset Y^n$ for all n, induces chain maps f_* between the cellular chain complexes of X and Y, hence also $f_* : H_n^{CW}(X) \to H_n^{CW}(Y)$, and the latter map corresponds to $f_* : H_n(X) \to H_n(Y)$ under the isomorphism $H_n^{CW} \approx H_n$.

13. For a CW pair (X, A) show there is a relative cellular chain complex formed by the groups $H_i(X^i, X^{i-1} \cup A^i)$, having homology groups isomorphic to $H_n(X, A)$.

14. Compute $H_i(\mathbb{R}\mathrm{P}^n/\mathbb{R}\mathrm{P}^m)$ for $m < n$.

15. For finite CW complexes X and Y, show that $\chi(X \times Y) = \chi(X)\chi(Y)$.

16. If a finite CW complex X is the union of subcomplexes A and B, show that $\chi(X) = \chi(A) + \chi(B) - \chi(A \cap B)$.

17. For X a finite CW complex and $p : \widetilde{X} \to X$ an n-sheeted covering space, show that $\chi(\widetilde{X}) = n\chi(X)$.

18. Show that if M_g, the closed orientable surface of genus g, is a covering space of M_h then $g = n(h-1) + 1$ for some n, namely, n is the number of sheets in the covering. [In §1.3 we constructed an n-sheeted covering $M_g \to M_h$ if $g = n(h-1) + 1$.]

19. Show that for each $n \in \mathbb{Z}$ there is a unique function φ assigning an integer to each finite CW complex, such that (a) $\varphi(X) = \varphi(Y)$ if X and Y are homeomorphic, (b) $\varphi(X) = \varphi(A) + \varphi(X/A)$ if A is a subcomplex of X, and (c) $\varphi(S^0) = n$. For such a function φ show that $\varphi(X) = \varphi(Y)$ if $X \simeq Y$.

20. For a pair (X, A), consider the space $X \cup CA$ obtained from X by attaching a cone on A, as in problem 19 of §2.1.

(a) Show that X is a retract of $X \cup CA$ iff A is *contractible in* X, i.e., there is a homotopy $f_t : A \to X$ with f_0 the inclusion $A \hookrightarrow X$ and f_1 a constant map.

(b) Show that if A is contractible in X then $H_n(X,A) \approx \widetilde{H}_n(X) \oplus \widetilde{H}_{n-1}(A)$. [Hint: $(X \cup CA)/X = SA$, the suspension of A.]

21. Give an example of a 2-dimensional CW complex X which retracts onto a subcomplex A, but X is not homotopy equivalent to $A \vee (X/A)$, even though there is a map $X \to A \vee (X/A)$ inducing isomorphisms on all homology groups.

22. Use the Mayer-Vietoris sequence to show there are isomorphisms $\widetilde{H}_n(X \vee Y) \approx \widetilde{H}_n(X) \oplus \widetilde{H}_n(Y)$ if the basepoints of X and Y which are identified in $X \vee Y$ are deformation retracts of neighborhoods $U \subset X$ and $V \subset Y$.

23. Show that $H_n(X \times S^1) \approx H_n(X) \oplus H_{n-1}(X)$ using the Mayer-Vietoris sequence.

24. Show via Mayer-Vietoris sequences that there are isomorphisms $\widetilde{H}_n(SX) \approx \widetilde{H}_{n-1}(X)$ for all n, where SX is the suspension of X.

25. Suppose the space X is the union of open sets $A_1, \cdots, A_n$ such that each intersection $A_{i_1} \cap \cdots \cap A_{i_k}$ is either empty or has trivial reduced homology groups. Show that $\widetilde{H}_i(X) = 0$ for $i \geq n-1$, and give an example showing this inequality is best possible, for each n.

26. Derive a Mayer-Vietoris sequence in simplicial homology for a Δ-complex X decomposed as the union of subcomplexes A and B. [This is much easier than for singular homology since no barycentric subdivision is necessary.]

27. Show that a commutative diagram

$$\begin{array}{ccccccccc}
\cdots \longrightarrow & C_{n+1} & & & B_n & \longrightarrow C_n & & & B_{n-1} \longrightarrow \cdots \\
& \downarrow & \searrow & \nearrow & \downarrow & \downarrow & \searrow & \nearrow & \downarrow \\
\cdots \longrightarrow & E_{n+1} & & A_n & D_n & \longrightarrow E_n & & A_{n-1} & D_{n-1} \longrightarrow \cdots
\end{array}$$

with the two sequences across the top and bottom exact, gives rise to an exact sequence

$$\cdots \to E_{n+1} \to B_n \to C_n \oplus D_n \to E_n \to B_{n-1} \to \cdots$$

where the maps are obtained from those in the previous diagram in the obvious way, except that $B_n \to C_n \oplus D_n$ has a minus sign in one coordinate.

28. Use the preceding exercise to derive relative Mayer-Vietoris sequences in any homology theory for CW pairs $(X,Y) = (A \cup B, C \cup D)$ with $A = B$ or $C = D$.

29. From the long exact sequence of homology groups associated to the short exact sequence of chain complexes

$$0 \longrightarrow C_i(X) \xrightarrow{n} C_i(X) \longrightarrow C_i(X; \mathbb{Z}_n) \longrightarrow 0$$

deduce immediately that there are short exact sequences

$$0 \longrightarrow H_i(X)/nH_i(X) \longrightarrow H_i(X;\mathbb{Z}_n) \longrightarrow n\text{-}Torsion(H_{i-1}(X)) \longrightarrow 0$$

where $n\text{-}Torsion(G)$ denotes the kernel of $G \xrightarrow{n} G$. Use this to show that $\widetilde{H}_i(X;\mathbb{Z}_p) = 0$ for all i and all primes p iff $\widetilde{H}_i(X)$ is a vector space over $\mathbb{Q}$ for all i.

30. For X a finite CW complex and F a field, show that $\chi(X) = \sum_n (-1)^n \dim H_n(X;F)$.

31. Let X be a finite connected graph having no vertex which is the endpoint of just one edge, and suppose that $H_1(X;\mathbb{Z})$ is free abelian of rank $n > 1$, so the group of automorphisms of $H_1(X;\mathbb{Z})$ is $GL_n(\mathbb{Z})$, the $n \times n$ matrices with integer entries having an inverse matrix of the same sort. Show that if G is a finite group of homeomorphisms of X, then the homomorphism $G \to GL_n(\mathbb{Z})$ assigning to $g : X \to X$ the induced homomorphism $g_* : H_1(X;\mathbb{Z}) \to H_1(X;\mathbb{Z})$ is injective. Show the same result holds if the coefficient group $\mathbb{Z}$ is replaced by $\mathbb{Z}_m$ with $m > 2$. What goes wrong when $m = 2$?

32. Show that a chain complex of finitely generated free abelian groups splits as the direct sum of subcomplexes with only one or two nonzero terms, of the form

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{Z} \longrightarrow 0 \longrightarrow \cdots \qquad \text{or} \qquad \cdots \longrightarrow 0 \longrightarrow \mathbb{Z} \xrightarrow{m} \mathbb{Z} \longrightarrow 0 \longrightarrow \cdots$$

Use this result to deduce that if X is a CW complex with finitely many cells in each dimension, then $H_i(X;G)$ is the direct sum of the following groups:

(a) a copy of G for each $\mathbb{Z}$ summand of $H_i(X)$

(b) a copy of G/mG for each $\mathbb{Z}_m$ summand of $H_i(X)$

(c) a copy of the kernel of $G \xrightarrow{m} G$ for each $\mathbb{Z}_m$ summand of $H_{i-1}(X)$.

33. Use the Mayer-Vietoris sequence to show that a non-orientable closed surface, or more generally a finite simplicial complex X for which $H_1(X)$ contains torsion, cannot be embedded as a subspace of $\mathbb{R}^3$ in such a way as to have a neighborhood homeomorphic to the mapping cylinder of some map from a closed orientable surface to X. [This assumption on a neighborhood is in fact not needed if one deduces the result from Alexander duality in §3.3.]

3. The Formal Viewpoint

Sometimes it is good to step back from the forest of details and look for general patterns. In this rather brief section we will first describe the general pattern of homology by axioms, then we will look at some common formal features shared by many of the constructions we have made.

Axioms for Homology

For simplicity let us restrict attention to CW complexes and focus on reduced homology to avoid mentioning relative homology. Then a **(reduced) homology theory** assigns to each CW complex X a sequence of abelian groups $\widetilde{h}_n(X)$ and to each map $f:X\to Y$ between CW complexes a sequence of homomorphisms $f_*:\widetilde{h}_n(X)\to\widetilde{h}_n(Y)$ such that $(fg)_*=f_*g_*$ and $\mathbb{1}_*=\mathbb{1}$, and so that the following three axioms are satisfied.

(1) If $f\simeq g:X\to Y$, then $f_*=g_*:\widetilde{h}_n(X)\to\widetilde{h}_n(Y)$.

(2) There are boundary homomorphisms $\partial:\widetilde{h}_n(X/A)\to\widetilde{h}_{n-1}(A)$ defined for each CW pair (X,A), fitting into an exact sequence

$$\cdots\xrightarrow{\partial}\widetilde{h}_n(A)\xrightarrow{i_*}\widetilde{h}_n(X)\xrightarrow{q_*}\widetilde{h}_n(X/A)\xrightarrow{\partial}\widetilde{h}_{n-1}(A)\xrightarrow{i_*}\cdots$$

where i is the inclusion and q is the quotient map. Furthermore the boundary maps are natural: For $f:(X,A)\to(Y,B)$ inducing a quotient map $\overline{f}:X/A\to Y/B$, the diagrams

$$\begin{array}{ccc}\widetilde{h}_n(X/A) & \xrightarrow{\partial} & \widetilde{h}_{n-1}(A)\\ \downarrow{\scriptstyle \overline{f}_*} & & \downarrow{\scriptstyle f_*}\\ \widetilde{h}_n(Y/B) & \xrightarrow{\partial} & \widetilde{h}_{n-1}(B)\end{array}$$

commute.

(3) For a wedge sum $X=\bigvee_\alpha X_\alpha$ with inclusions $i_\alpha:X_\alpha\hookrightarrow X$, the direct sum map $\bigoplus_\alpha i_{\alpha*}:\bigoplus_\alpha\widetilde{h}_n(X_\alpha)\to\widetilde{h}_n(X)$ is an isomorphism for each n.

Negative values for the subscripts n are permitted. Ordinary singular homology is zero in negative dimensions by definition, but interesting homology theories with nontrivial groups in negative dimensions do exist.

The third axiom may seem less substantial than the first two, and indeed for finite wedge sums it can be deduced from the first two axioms, though not in general for infinite wedge sums, as an example in the exercises shows.

It is also possible, and not much more difficult, to give axioms for unreduced homology theories. One supposes one has relative groups $h_n(X,A)$ defined, specializing to absolute

groups by setting $h_n(X) = h_n(X, \varnothing)$. Axiom (1) is replaced by its obvious relative form, and the axiom (2) is broken into two parts, the first hypothesizing a long exact sequence involving these relative groups, with natural boundary maps, the second stating some version of excision, for example $h_n(X, A) \approx h_n(X/A, A/A)$ if one is dealing with CW pairs. In axiom (3) the wedge sum is replaced by disjoint union.

These axioms for unreduced homology are essentially the same as those originally laid out in the highly influential book [Eilenberg-Steenrod], except that axiom (3) was omitted since the focus there was on finite complexes, and there was another axiom specifying that the groups $h_n(\textit{point})$ are zero for $n \neq 0$, as is true for singular homology. This axiom was called the 'dimension axiom,' presumably because it specifies that a point has nontrivial homology only in dimension zero. It can be regarded as a normalization axiom, since one can trivially define a homology theory where it fails by setting $h_n(X, A) = H_{n+k}(X, A)$ for a fixed nonzero integer k. At the time there were no interesting homology theories known for which the dimension axiom did not hold, though only a short time afterwards a homology theory called 'bordism' was defined having the property that the bordism groups of a point are nonzero in infinitely many dimensions. Axiom (3) seems to have appeared first in [Milnor 62].

Reduced and unreduced homology theories are essentially equivalent. From an unreduced theory h one gets a reduced theory $\widetilde{h}$ by setting $\widetilde{h}_n(X)$ equal to the kernel of the canonical map $h_n(X) \to h_n(\textit{point})$. In the other direction, one sets $h_n(X) = \widetilde{h}_n(X_+)$ where X_+ is the disjoint union of X with a point. We leave it as an exercise to show that these two transformations between reduced and unreduced homology are inverses of each other. Just as with ordinary homology, one has $h_n(X) \approx \widetilde{h}_n(X) \oplus h_n(x_0)$ for any point $x_0 \in X$, since the long exact sequence of the pair (X, x_0) splits via a retraction $X \to x_0$. Note that $\widetilde{h}_n(x_0) = 0$ for all n, as can be seen by looking at the long exact sequence of reduced homology groups of the pair (x_0, x_0).

The groups $h_n(x_0) \approx \widetilde{h}_n(S^0)$ are called the **coefficients** of the homology theories h and $\widetilde{h}$, in analogy with the case of singular homology with coefficients. One can trivially realize any sequence of abelian groups G_i as the coefficient groups of a homology theory by setting $h_n(X, A) = \bigoplus_i H_{n-i}(X, A; G_i)$.

In general, homology theories are not uniquely determined by their coefficient groups, but this is true for singular homology:

Theorem 2.29. *If h is a homology theory defined for CW pairs, whose coefficient groups are zero except in dimension zero, then h is equivalent to ordinary singular homology with coefficients in the group $G = h_0(x_0)$: There are natural isomorphisms $h_n(X, A) \approx$*

$H_n(X, A; G)$ for all CW pairs (X, A) and all n.

This is a satisfying result, but not one with frequent applications, so we will only sketch the proof. The idea is that one can define cellular homology groups $h_n^{CW}(X)$ just as for ordinary homology, via the chain complex $\cdots \to h_n(X^n, X^{n-1}) \to h_{n-1}(X^{n-1}, X^{n-2}) \to \cdots$. Then one shows there are isomorphisms

$$h_n(X) \approx h_n^{CW}(X) \approx H_n^{CW}(X; G) \approx H_n(X; G)$$

The last of these isomorphisms is familiar to us, and the first is shown in the same way. For the middle isomorphism, the cellular chain groups in the two cases are isomorphic, with a copy of G for each cell, so one only has to check that the cellular boundary maps are the same. The calculation of cellular boundary maps in terms of degrees of certain maps between spheres works for the homology theory h, so what is needed is the fact that a map $S^n \to S^n$ which induces multiplication by m on ordinary homology also induces multiplication by m on $\widetilde{h}_n(S^n) \approx G$.

The easiest way to show this last fact is to use homotopy groups, in particular the fact that $\pi_n(S^n)$ is isomorphic to $\mathbb{Z}$ for $n > 0$, with the identity map of S^n a generator, as shown in §4.1. By assigning to each map $S^n \to S^n$ the induced homomorphism on $\widetilde{h}_n(S^n)$ we get a function from $\pi_n(S^n)$ to the group of homomorphisms $G \to G$, and this function is a homomorphism since $(f + g)_* = f_* + g_*$ as shown in Proposition 4.6. Since $\mathbb{1}_* = \mathbb{1}$, the result follows, at least for $n > 0$. For $n = 0$ there are only four maps $S^0 \to S^0$ to consider. Two of these are constant hence induce the trivial map on $\widetilde{h}_0$, and the identity map induces the identity, so one only has to show that the map $S^0 \to S^0$ interchanging the two points induces minus the identity. This can be done by a direct argument using $S^0 \vee S^0$, or by using suspensions, since this map suspends to a reflection of S^1, of degree -1, and suspension preserves degree for any homology theory, by the argument in Proposition 2.24.

It would be possible to avoid this fact about degrees by considering Δ-complexes rather than CW complexes, as is done for simplicial complexes in Chapter 3 of [Eilenberg-Steenrod]. One would only need the fact that reflections have degree -1 in any homology theory, which is easier to show.

There is a broader connection between suspension and homology: For any reduced homology theory $\widetilde{h}$ there are natural isomorphisms $\widetilde{h}_n(X) \approx \widetilde{h}_{n+1}(SX)$ for all n and X This comes from the long exact sequence of the pair (CX, X) where CX is the cone on X. Since CX is contractible, we have $\widetilde{h}_n(CX) = 0$ for all n, so the boundary maps in the long exact sequence are isomorphisms $\widetilde{h}_{n+1}(CX/X) \approx \widetilde{h}_n(X)$, and $CX/X = SX$.

Mayer-Vietoris sequences for any homology theory can be obtained in the following way. For a CW complex $X = A \cup B$ with A and B subcomplexes, consider the commutative diagram consisting of the exact sequences of the pairs (X, A) and $(B, A \cap B)$

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & h_{n+1}(B, A\cap B) & \longrightarrow & h_n(A\cap B) & \longrightarrow & h_n(B) & \longrightarrow & h_n(B, A\cap B) & \longrightarrow & \cdots \\
 & & \downarrow\approx & & \downarrow & & \downarrow & & \downarrow\approx & & \\
\cdots & \longrightarrow & h_{n+1}(X, A) & \longrightarrow & h_n(A) & \longrightarrow & h_n(X) & \longrightarrow & h_n(X, A) & \longrightarrow & \cdots
\end{array}$$

The vertical maps are induced by inclusions, and the vertical maps between relative groups are isomorphisms since $B/(A\cap B) = X/A$. Then it is a purely algebraic fact, whose proof is exercise 27 at the end of the previous section, that a diagram such as this with every third vertical map an isomorphism gives rise to an exact sequence

$$\cdots \longrightarrow h_n(A\cap B) \xrightarrow{\varphi} h_n(A) \oplus h_n(B) \xrightarrow{\psi} h_n(X) \xrightarrow{\partial} h_{n-1}(A\cap B) \longrightarrow \cdots$$

Categories and Functors

Formally, singular homology is a sequence of functions H_n which assign to each space X an abelian group $H_n(X)$ and to each map $f\colon X \to Y$ a homomorphism $H_n(f) = f_*\colon H_n(X) \to H_n(Y)$, and similarly for relative homology groups. This sort of situation arises quite often, and not just in algebraic topology, so it is useful to introduce some general terminology for it. Roughly speaking, 'functions' like H_n are called 'functors,' and the domains and ranges of these functors are called 'categories.' Thus for H_n, the domain category consists of topological spaces and continuous maps, or in the relative case, pairs of spaces and continuous maps of pairs, and the range category consists of abelian groups and homomorphisms. A key point is that one is interested not only in the objects in the category, e.g., spaces or groups, but also in the maps, or 'morphisms,' between these objects.

Now for the precise definitions. A **category** $\mathcal{C}$ consists of:

(1) a collection $\mathrm{Ob}(\mathcal{C})$ of **objects**.

(2) sets $\mathrm{Mor}(X, Y)$ of **morphisms** for each pair $X, Y \in \mathrm{Ob}(\mathcal{C})$, including a distinguished 'identity' morphism $\mathbb{1} = \mathbb{1}_X \in \mathrm{Mor}(X, X)$ for each X.

(3) a 'composition of morphisms' function $\circ\colon \mathrm{Mor}(X, Y) \times \mathrm{Mor}(Y, Z) \to \mathrm{Mor}(X, Z)$ for each triple $X, Y, Z \in \mathrm{Ob}(\mathcal{C})$, satisfying $f\circ\mathbb{1} = f$, $\mathbb{1}\circ f = f$, and $(f\circ g)\circ h = f\circ(g\circ h)$.

There are plenty of obvious examples, such as:

— The category of topological spaces, with continuous maps as the morphisms. Or we could restrict to special classes of spaces such as CW complexes, keeping continuous

maps as the morphisms. We could also restrict the morphisms, e.g., the category of topological spaces with homeomorphisms as the morphisms.

— The category of groups, with homomorphisms as morphisms. Or the subcategory of abelian groups, again with homomorphisms as the morphisms. Generalizing this is the category of modules over a fixed ring, with morphisms the module homomorphisms.

— The category of sets, with arbitrary functions as the morphisms. Or the morphisms could be restricted, e.g., to injections, or to surjections.

There are also many examples of categories where the morphisms are not simply functions, for example:

— Any group G can be viewed as a category with only one object and with G as the morphisms of this object, so that condition (3) reduces to two of the three axioms for a group. If we require only these two axioms, associativity and a left and right identity, we have a 'group without inverses,' usually called a *monoid* since it is the same thing as a category with one object.

— A partially ordered set $(X, \leq)$ can be considered a category where the objects are the elements of X and there is a unique morphism from x to y whenever $x \leq y$. The relation $x \leq x$ gives the morphism $\mathbb{1}$ and transitivity gives the composition $\text{Mor}(x, y) \times \text{Mor}(y, z) \to \text{Mor}(x, z)$. The condition that $x \leq y$ and $y \leq x \Rightarrow x = y$ then says that $\text{Mor}(x, y) \cup \text{Mor}(y, x)$ contains at most one element for all pairs $x, y \in X$, i.e., there is at most one morphism between any two objects.

— There is a 'homotopy category' whose objects are topological spaces and whose morphisms are homotopy classes of maps, rather than actual maps. This uses the fact that composition is well-defined on homotopy classes: $f_0 g_0 \simeq f_1 g_1$ if $f_0 \simeq f_1$ and $g_0 \simeq g_1$.

— Chain complexes are the objects of a category, with chain maps as morphisms. This category has various interesting subcategories, obtained by restricting the objects. For example, we could take chain complexes whose groups are zero in negative dimensions, or zero outside some finite range. Or we could restrict to exact sequences, or short exact sequences. In each case we take morphisms to be chain maps, i.e., commutative diagrams. Going a step further, there is a category whose objects are short exact sequences of chain complexes and whose morphisms are commutative diagrams of maps between such short exact sequences.

A **functor** F from a category $\mathcal{C}$ to a category $\mathcal{D}$ assigns to each object X in $\mathcal{C}$ an object $F(X)$ in $\mathcal{D}$ and to each morphism $f \in \text{Mor}(X, Y)$ a morphism $F(f) \in \text{Mor}(F(X), F(Y))$, such that $F(\mathbb{1}) = \mathbb{1}$ and $F(f \circ g) = F(f) \circ F(g)$. In the case of the

functor H_n the latter two conditions are the familiar properties $\mathbb{1}_* = \mathbb{1}$ and $(fg)_* = f_* g_*$ of induced maps. Strictly speaking, what we have just defined is a **covariant** functor. A **contravariant** functor would assign to $f \in \mathrm{Mor}(X, Y)$ a 'backwards' morphism $F(f) \in \mathrm{Mor}(F(Y), F(X))$ with $F(\mathbb{1}) = \mathbb{1}$ and $F(f \circ g) = F(g) \circ F(f)$. A classical example of this is the dual vector space functor, which assigns to a vector space V over a fixed scalar field K the dual vector space $F(V) = V^*$ of linear maps $V \to K$, and to each linear transformation $f : V \to W$ the dual map $F(f) = f^* : W^* \to V^*$, in the reverse direction. In the next chapter we shall study the contravariant version of homology, called cohomology.

A number of the constructions we have studied in this chapter are functors:

— The singular chain complex functor assigns to a space X the chain complex of singular chains in X and to a map $f : X \to Y$ the induced chain map. This is a functor from the category of spaces and continuous maps to the category of chain complexes and chain maps.

— The algebraic homology functor assigns to a chain complex its sequence of homology groups and to a chain map the induced homomorphisms on homology defines a functor from the category of chain complexes and chain maps to the category whose objects are sequences of abelian groups and whose morphisms are sequences of homomorphisms.

— The composition of the two preceding functors is the functor assigning to a space its singular homology groups.

— The first example above, the singular chain complex functor, can itself be regarded as the composition of two functors. The first functor assigns to a space X its singular complex $\Delta(X)$, a Δ-complex, and the second functor assigns to a Δ-complex its simplicial chain complex. This is what the two functors do on objects, and what they do on morphisms can be described in the following way. A map of spaces $f : X \to Y$ induces a map $f_* : \Delta(X) \to \Delta(Y)$ by composing singular simplices $\Delta^n \to X$ with f. The map f_* is a map between Δ-complexes taking the distinguished characteristic maps in the domain Δ-complex to the distinguished characteristic maps in the target Δ-complex. Call such maps **Δ-maps** and let them be the morphisms in the category of Δ-complexes. Note that a Δ-map induces a chain map between simplicial chain complexes, taking basis elements to basis elements, so we have a simplicial chain complex functor taking the category of Δ-complexes and Δ-maps to the category of chain complexes and chain maps.

— There is a functor assigning to a pair of spaces (X, A) the associated long exact sequence of homology groups. Morphisms in the domain category are maps of pairs, and in the target category morphisms are maps between exact sequences forming

commutative diagrams. This functor is the composition of two functors, the first assigning to (X, A) a short exact sequence of chain complexes, the second assigning to such a short exact sequence the associated long exact sequence of homology groups. Morphisms in the intermediate category are the evident commutative diagrams.

Another sort of process one encounters from time to time is the transformation of one functor into another, for example:

— Boundary maps $H_n(X, A) \to H_{n-1}(A)$ in singular homology, or indeed in any homology theory.

— Change-of-coefficient homomorphisms $H_n(X; G_1) \to H_n(X; G_2)$ induced by a homomorphism $G_1 \to G_2$, as in the proof of Lemma 2.28.

Formally, if one has two functors $F, G : \mathcal{C} \to \mathcal{D}$ then a **natural transformation** T from F to G assigns a morphism $T_X : F(X) \to G(X)$ to each object $X \in \mathcal{C}$, in such a way that for each morphism $f : X \to Y$ in $M\mathcal{C}$ the square

$$\begin{array}{ccc} F(X) & \xrightarrow{F(f)} & F(Y) \\ {\scriptstyle T_X}\downarrow & & \downarrow{\scriptstyle T_Y} \\ G(X) & \xrightarrow{G(f)} & G(Y) \end{array}$$

commutes. The case that F and G are contravariant rather than covariant would be similar.

We have been describing the passage from topology to the abstract world of categories and functors, but there is also a nice path in the opposite direction:

— To each category $\mathcal{C}$ there is associated a Δ-complex $B\mathcal{C}$ called the **classifying space** of $\mathcal{C}$, having as its n-simplices the strings of morphisms $X_0 \to X_1 \to \cdots \to X_n$ in $\mathcal{C}$. The $n+1$ faces of this simpex are obtained by composing two adjacent morphisms in the string (this yields $n-1$ faces) or by deleting the first or last morphism, to get the remaining two faces. In case $\mathcal{C}$ has a single object, whose morphisms form a group G, then $B\mathcal{C}$ is the same as the Δ-complex BG constructed in §1.4b, a $K(G, 1)$ space. In general, $B\mathcal{C}$ need not be a $K(G, 1)$. For example, if we start with a Δ-complex X and regard its set of simplices as a partially ordered set $\mathcal{C}(X)$ under the relation of inclusion of faces, then $B\mathcal{C}(X)$ is the barycentric subdivision of X.

— A functor $F : \mathcal{C} \to \mathcal{D}$ induces a map $B\mathcal{C} \to B\mathcal{D}$, in fact a Δ-map, sending the n-simplex $X_0 \to X_1 \to \cdots \to X_n$ to the n-simplex $F(X_0) \to F(X_1) \to \cdots \to F(X_n)$.

— A natural transformation from a functor F to a functor G induces a homotopy between the induced maps of classifying spaces. We leave this for the reader to make explicit, using the subdivision of $\Delta^n \times I$ into $(n+1)$-simplices described earlier in the chapter.

Exercises

1. If $T_n(X, A)$ denotes the torsion subgroup of $H_n(X, A; \mathbb{Z})$, show that the functors $(X, A) \mapsto T_n(X, A)$, with the obvious induced homomorphisms $T_n(X, A) \to T_n(Y, B)$ and boundary maps $T_n(X, A) \to T_{n-1}(A)$, do not define a homology theory. Do the same for the 'mod torsion' functor $MT_n(X, A) = H_n(X, A; \mathbb{Z})/T_n(X, A)$.

2. Define a candidate for a reduced homology theory on CW complexes by $\widetilde{h}_n(X) = \prod_i \widetilde{H}_i(X)/\bigoplus_i \widetilde{H}_i(X)$. Thus $\widetilde{h}_n(X)$ is independent of n and is zero if X is finite-dimensional, but is not identically zero, e.g., for $X = \bigvee_i S^i$. Show that the axioms for a homology theory are satisfied except that the wedge axiom fails.

4. Additional Topics

a. Homology and Fundamental Group

There is a close relation between $H_1(X)$ and $\pi_1(X)$. The idea is that a path $f : I \to X$ can be viewed as a singular 1-simplex in X, and if f is a loop, this singular 1-simplex is a cycle since $\partial f = f(1) - f(0)$.

Theorem 2.30. *By regarding loops as singular 1-cycles, we obtain a homomorphism $h : \pi_1(X, x_0) \to H_1(X)$ whose kernel is the commutator subgroup of $\pi_1(X)$. If X is path-connected, then h is surjective and induces an isomorphism from the abelianization of $\pi_1(X)$ onto $H_1(X)$.*

Proof: Recall the notation $f \simeq g$ for the relation of homotopy, fixing endpoints, between paths f and g. Regarding f and g as chains, the notation $f \sim g$ will mean that $f - g$ is the boundary of a 2-chain.

(i) If f is a constant path, then $f \sim 0$. Namely, f is the boundary of the constant singular 2-simplex σ since $\partial\sigma = \sigma\big|[v_1, v_2] - \sigma\big|[v_0, v_2] + \sigma\big|[v_0, v_1] = f - f + f = f$.

(ii) If $f \simeq g$ then $f \sim g$. To see this, consider a homotopy $F : I \times I \to X$ from f to g. This yields a pair of singular 2-simplices σ_1 and σ_2 in X by subdividing the square $I \times I$ into two triangles $[v_0, v_1, v_3]$ and $[v_0, v_2, v_3]$, as shown in the figure. When one computes $\partial(\sigma_1 - \sigma_2)$, the two restrictions of F to the diagonal of the square cancel, leaving $f - g$ together with two constant singular 1-simplices from the left and right edges of the square. By (i) these are boundaries, so $f - g$ is also a boundary.

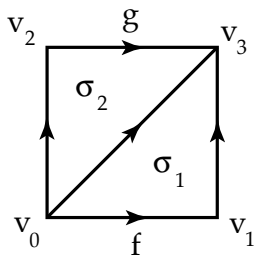

(iii) $f \cdot g \sim f + g$, where $f \cdot g$ denotes the product of the paths f and g. In fact we have the formula $\partial\sigma = g - f \cdot g + f$ where $\sigma : \Delta^2 \to X$ is the composition of the orthogonal projection of $\Delta^2 = [v_0, v_1, v_2]$ onto the edge $[v_0, v_2]$ followed by $f \cdot g : [v_0, v_2] \to X$.

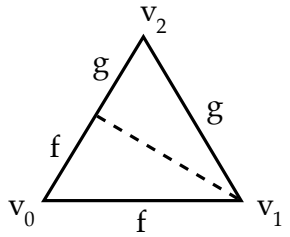

(iv) $\overline{f} \sim -f$, where $\overline{f}$ is the inverse path of f. This follows from the preceding three observations, which give $f + \overline{f} \sim f \cdot \overline{f} \sim 0$.

Applying (ii) and (iii) to loops, it follows that we have a well-defined homomorphism $h : \pi_1(X, x_0) \to H_1(X)$ sending the homotopy class of a loop f to the homology class of the 1-cycle f.

To show h is surjective when X is path-connected, let a 1-cycle $\sum n_i \sigma_i$ be given. For each point $x_j \in X$ which is an endpoint of some σ_i, choose a path γ_j from the basepoint

x_0 to x_j. Each σ_i then determines a loop $\widetilde{\sigma}_i$ going from x_0 to the starting point of σ_i via the appropriate γ_j, then along σ_i, and then back to x_0 along the inverse of a γ_j. Since $\sum n_i\sigma_i$ is a cycle, the coefficient of each x_j in $\partial(\sum n_i\sigma_i)$ is zero. This implies that the loop $\prod \widetilde{\sigma}_i^{n_i}$ traverses each γ_j the same number of times forward as backward. Hence $\sum n_i\widetilde{\sigma}_i \sim \sum n_i\sigma_i$ by (iii) and (iv). Since h takes the homotopy class of $\prod \widetilde{\sigma}_i^{n_i}$ to the homology class of $\sum n_i\widetilde{\sigma}_i$, which equals the homology class of $\sum n_i\sigma_i$, this shows h is surjective.

The commutator subgroup of $\pi_1(X)$ is contained in the kernel of h since $H_1(X)$ is abelian. To obtain the reverse inclusion we shall show that every class $[f]$ in $Ker(h)$ is trivial in the abelianization $\pi_1(X)_{ab}$ of $\pi_1(X)$.

If $h([f]) = 0$, then as a 1-cycle, f is the boundary of a 2-chain $\sum n_i\sigma_i$. The various restrictions $\sigma_i\,|\,[v_p, v_q]$ with $p < q$ define singular 1-simplices τ_j, and choosing paths from the basepoint to the endpoints of the τ_j's as before, each τ_j determines a loop $\widetilde{\tau}_j$. Corresponding to $\partial\sigma_i = \tau_j - \tau_k + \tau_\ell$ there is the loop $\widetilde{\tau}_j\overline{\widetilde{\tau}}_k\widetilde{\tau}_\ell$ which we denote $\widetilde{\partial\sigma_i}$. The map σ_i gives rise to a homotopy from $\widetilde{\partial\sigma_i}$ to the constant loop, as one can see from the figure. The formula $f = \partial(\sum n_i\sigma_i) = \sum n_i\partial\sigma_i$ implies that $[f] = \sum n_i[\widetilde{\partial\sigma_i}]$ in $\pi_1(X)_{ab}$ since all the τ_j's in $\sum n_i\partial\sigma_i$ cancel algebraically except for one, which equals f. We have already noted that $[\widetilde{\partial\sigma_i}] = 0$, so we conclude that $[f] = 0$ in $\pi_1(X)_{ab}$. □

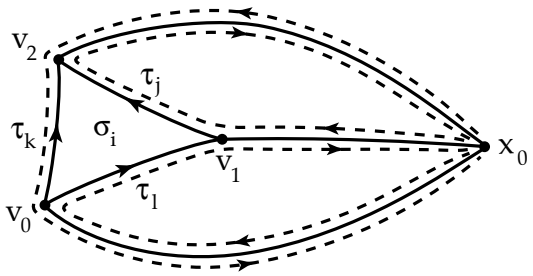

Exercises

1. Show that the map $h : \pi_1(X, x_0) \to H_1(X)$ can also be defined by $h([f]) = f_*(\alpha)$ where $f : S^1 \to X$ represents a given element of $\pi_1(X, x_0)$, f_* is the induced map on H_1, and α is a suitable generator of $H_1(S^1) \approx \mathbb{Z}$.

b. Theorems of Brouwer and Hopf

Algebraic topology can be used to give interesting point-set topology results using the following calculations:

Proposition 2.31.

(a) *If D is a subspace of S^n homeomorphic to D^k for some $k \geq 0$, then $\widetilde{H}_i(S^n - D) = 0$ for all i.*

(b) *If S is a subspace of S^n homeomorphic to S^k for some k with $0 \leq k < n$, then $\widetilde{H}_i(S^n - S)$ is $\mathbb{Z}$ for $i = n - k - 1$ and 0 otherwise.*

As a special case of (b) we have the Jordan curve theorem, that a subspace of S^2 homeomorphic to S^1 separates S^2 into two complementary components, or equivalently, path-components since open subsets of S^n are locally path-connected. One could just as well use $\mathbb{R}^2$ in place of S^2 here since deleting a point from an open set in S^2 does not affect its connectedness. More generally, (b) says that a subspace of S^n homeomorphic to S^{n-1} separates it into two components, and these components have the same homology groups as a point. However, there are examples where these complementary components are not simply-connected, for example the well-known Alexander horned sphere in S^3. These complications involving embedded S^{n-1}'s in S^n are all local in nature since it is known that any locally nice embedded S^{n-1} in S^n is equivalent to the standard $S^{n-1} \subset S^n$; see [M.Brown].

In contrast with this, when we come to embeddings of S^{n-2} in S^n, even locally nice embeddings need not be equivalent to the standard one. This is the subject of knot theory, including the classical case of knotted embeddings of S^1 in S^3 or $\mathbb{R}^3$. For embeddings of S^{n-2} in S^n the complement always has the same homology as S^1, according to the theorem, but the fundamental group can be quite different. In spite of the fact that the homology of a knot complement does not detect knottedness, it is still possible to use homology to distinguish different knots by looking at the homology of covering spaces of their complements. For more on this see [Rolfsen].

Proof: We prove (a) by induction on k. When $k = 0$, $S^n - D$ is homeomorphic to $\mathbb{R}^n$, so this case is trivial. For the induction step, let $h : I^k \to D$ be a homeomorphism, and consider the open sets $A = S^n - h(I^{k-1} \times [0, 1/2])$ and $B = S^n - h(I^{k-1} \times [1/2, 1])$, with $A \cap B = S^n - D$ and $A \cup B = S^n - h(I^{k-1} \times \{1/2\})$. By induction $\widetilde{H}_i(A \cup B) = 0$ for all i, so the Mayer-Vietoris sequence gives isomorphisms $\Phi : \widetilde{H}_i(S^n - D) \to \widetilde{H}_i(A) \oplus \widetilde{H}_i(B)$ for all i. Modulo signs, the two components of Φ are induced by the inclusions $S^n - D \hookrightarrow A$ and $S^n - D \hookrightarrow B$, so if there exists an i-dimensional cycle α in $S^n - D$ which is not a

boundary in $S^n - D$, then α is also not a boundary in at least one of A and B. When $i = 0$ the word 'cycle' here is to be interpreted in the sense of augmented chain complexes since we are dealing with reduced homology. Repeating this argument, we can further subdivide the last I factor of I^k into quarters, eighths, etc., producing a nested sequence of closed subintervals $I_1 \supset I_2 \supset \cdots$ with intersection a point $p \in I$, such that α is not a boundary in $S^n - h(I^{k-1} \times I_m)$ for any m. By induction on k, α is the boundary of a chain β in $S^n - h(I^{k-1} \times \{p\})$. This β is a finite linear combination of singular simplices with compact image in $S^n - h(I^{k-1} \times \{p\})$. The union of these images is covered by the nested sequence of open sets $S^n - h(I^{k-1} \times I_m)$, so by compactness β must actually be a chain in $S^n - h(I^{k-1} \times I_m)$ for some m. This contradiction shows that α must be a boundary in $S^n - D$, finishing the induction step.

Part (b) is also proved by induction on k, starting with the trivial case $k = 0$ when $S^n - S$ is homeomorphic to $S^{n-1} \times \mathbb{R}$. For the induction step, write $S = D_1 \cup D_2$ with D_1 and D_2 homeomorphic to D^k and $D_1 \cap D_2$ homeomorphic to S^{k-1}. The Mayer-Vietoris sequence for $A = S^n - D_1$ and $B = S^n - D_2$, which have trivial reduced homology by part (a), then gives isomorphisms $\widetilde{H}_i(S^n - S) \approx \widetilde{H}_{i+1}(S^n - (D_1 \cap D_2))$ for all i. □

The next theorem says that for subspaces of $\mathbb{R}^n$, the property of being open is a topological invariant. This accounts for the classical name 'Invariance of Domain' associated to this result, 'domain' being a synonym for 'open set' in $\mathbb{R}^n$.

Theorem 2.32. *If a subspace X of $\mathbb{R}^n$ is homeomorphic to an open set in $\mathbb{R}^n$, then X is itself open in $\mathbb{R}^n$.*

Proof: Regarding S^n as the one-point compactification of $\mathbb{R}^n$, an equivalent statement is that X is open in S^n, and this is what we shall prove. Each $x \in X$ has a neighborhood D in X homeomorphic to D^n, with x corresponding to the center of D^n under this homeomorphism. Let $S \subset D$ correspond to $S^{n-1} = \partial D^n$ under this homeomorphism. Then $S^n - D$ is open, and is connected by the case $k = n$, $i = 0$ in part (a) of the previous proposition, using the fact that connectedness is equivalent to path-connectedness for open sets in S^n. Also, $S^n - S$ is open and has exactly two components by the case $k = n - 1$, $i = 0$ of (b) of the proposition. Thus we have $S^n - S$ decomposed as the disjoint union of the connected sets $S^n - D$ and $D - S$, so these must be the two components of $S^n - S$. In particular, $D - S$ is open, being a component of the open set $S^n - S$. So $D - S$ is an open neighborhood of x in S^n, and since this neighborhood is contained in X we conclude that X is open. □

A corollary of this theorem is that if we have a connected space X and a compact

subspace $Y \subset X$, each of X and Y being locally homeomorphic to $\mathbb{R}^n$, then $X = Y$, i.e., $X - Y = \varnothing$. This follows since Y is open in X by the theorem, and Y is closed since it is compact, so $Y = X$ since X is connected. For example, this says that S^n cannot be embedded as a subspace of $\mathbb{R}^n$, since if it were we would have $S^n \approx \mathbb{R}^n$ but S^n is compact while $\mathbb{R}^n$ is not. A consequence of this is that $\mathbb{R}^n$ contains no subspace homeomorphic to $\mathbb{R}^m$ with $m > n$, for if it did it would also contain a subspace homeomorphic to S^n since $S^n \subset \mathbb{R}^m$ if $n < m$. More generally, there cannot be a continuous injection $\mathbb{R}^m \to \mathbb{R}^n$ when $m > n$ since this would give a continuous injection $S^n \to \mathbb{R}^n$, which would be a homeomorphism onto its image since S^n is compact.

The Invariance of Domain and the n-dimensional generalization of the Jordan curve theorem were first proved by Brouwer around 1910, at a very early stage in the development of algebraic topology.

Division Algebras

Here is an algebraic application of homology theory due to H. Hopf:

Theorem 2.33. *$\mathbb{R}$ and $\mathbb{C}$ are the only finite-dimensional division algebras over $\mathbb{R}$ which are commutative and have an identity.*

By definition, an **algebra** structure on $\mathbb{R}^n$ is a bilinear multiplication $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$, $(a, b) \mapsto ab$. Thus the product satisfies left and right distributivity, $a(b + c) = ab + ac$ and $(a + b)c = ac + bc$, and scalar associativity, $\alpha(ab) = (\alpha a)b = a(\alpha b)$ for $\alpha \in \mathbb{R}$. Commutativity, full associativity, and an identity element are not assumed. An algebra is a **division algebra** if the equations $ax = b$ and $xa = b$ are always solvable, if $a \neq 0$. In other words, the linear transformations $x \mapsto ax$ and $x \mapsto xa$ are surjective when $a \neq 0$. These are linear maps $\mathbb{R}^n \to \mathbb{R}^n$, so surjectivity is equivalent to having trivial kernel, i.e., the absence of zero-divisors.

The four classical examples are $\mathbb{R}$, $\mathbb{C}$, the quaternions $\mathbb{H}$, and the octonions $\mathbb{O}$. Frobenius proved in 1877 that $\mathbb{R}$, $\mathbb{C}$, and $\mathbb{H}$ are the only finite-dimensional associative division algebras over $\mathbb{R}$, and in 1898 Hurwitz proved that these three together with $\mathbb{O}$ are the only finite-dimensional division algebras over $\mathbb{R}$ with a product satisfying $|ab| = |a||b|$. See [Ebbinghaus]. We shall show in §3.2 that a finite-dimensional division algebra must have dimension a power of 2. By refinements of these topological arguments, Adams showed that in fact the only possible dimensions are 1, 2, 4, and 8. A very nice proof using K-theory is in [Adams-Atiyah], and an exposition of this proof can be found also in [VBKT].

Proof: Suppose first that $\mathbb{R}^n$ has a commutative division algebra structure. Define a map $f: S^{n-1} \to S^{n-1}$ by $f(x) = x^2/|x^2|$. This is well-defined since $x \neq 0$ implies $x^2 \neq 0$ in a division algebra, hence $|x^2| > 0$. The map f is continuous since the multiplication map $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ is bilinear, hence continuous. Since $f(-x) = f(x)$, f induces a quotient map $\overline{f}: \mathbb{R}\mathrm{P}^{n-1} \to S^{n-1}$. The following argument shows that $\overline{f}$ is injective. An equality $f(x) = f(y)$ implies $x^2 = \alpha^2 y^2$ for $\alpha = (|x^2|/|y^2|)^{1/2} > 0$. Thus we have $x^2 - \alpha^2 y^2 = 0$, which factors as $(x + \alpha y)(x - \alpha y) = 0$ using commutativity and the fact that α is a real scalar. Since there are no divisors of zero we deduce that $x = \pm\alpha y$. Since x and y are unit vectors and α is a positive real scalar, this yields $x = \pm y$, so x and y determine the same point of $\mathbb{R}\mathrm{P}^{n-1}$, completing the verification that $\overline{f}$ is injective.

Since $\overline{f}$ is an injective map of compact Hausdorff spaces, it must be a homeomorphism onto its image. By the remarks following the proof of Theorem 2.32, $\overline{f}$ must in fact be surjective if we are not in the trivial case $n = 1$. Thus we have a homeomorphism $\mathbb{R}\mathrm{P}^{n-1} \approx S^{n-1}$. This implies $n = 2$ since $\mathbb{R}\mathrm{P}^{n-1}$ and S^{n-1} have different fundamental groups (and different homology groups) if $n > 2$.

It remains to show that a 2-dimensional commutative division algebra A with identity is isomorphic to $\mathbb{C}$. This is elementary algebra: If $j \in A$ is not a real scalar multiple of the identity element $1 \in A$ and we write $j^2 = a + bj$ for $a, b \in \mathbb{R}$, then $(j - b/2)^2 = a + b^2/4$ so we may assume $b = 0$ by rechoosing j. If $a \geq 0$, say $a = c^2$, then $j^2 = c^2$ implies $(j + c)(j - c) = 0$ so $j = \pm c$, contradicting the choice of j. So $j^2 = -c^2$ and by rescaling j we may assume $j^2 = -1$, hence A is isomorphic to $\mathbb{C}$. □

Leaving out the last paragraph of the proof, we have shown that a finite-dimensional commutative division algebra, not necessarily with an identity, must have dimension at most 2. Oddly enough, there do exist 2-dimensional commutative division algebras without identity elements, for example $\mathbb{C}$ with the modified multiplication $z \cdot w = \overline{zw}$, the bar denoting complex conjugation.

Exercises

1. If $X \subset S^n$ is a subspace homeomorphic to $S^k \vee S^\ell$, compute $H_i(S^n - X)$. Do this also when X is homeomorphic to $S^k \amalg S^\ell$.

2. Show that $\widetilde{H}_i(S^n - X) \approx \widetilde{H}_{n-i-1}(X)$ when X is homeomorphic to a finite connected graph. [First do the case that the graph is a tree.]

3. In the unit sphere $S^{p+q-1} \subset \mathbb{R}^{p+q}$ let S^{p-1} and S^{q-1} be the subspheres consisting of points whose last q and first p coordinates are zero, respectively. Show that S^{p-1} and S^{q-1} are not the boundaries of any pair of disjointly embedded disks D^p and D^q in D^{p+q}.

4. Show that $\mathbb{R}^{2n+1}$ is not a division algebra over $\mathbb{R}$ if $n > 0$, by showing that, if it were, then the map $S^{2n} \to S^{2n}$, $x \mapsto ax/|ax|$ would be homotopic to $x \mapsto -ax/|ax|$, for nonzero $a \in \mathbb{R}^{2n+1}$, but these maps have different degrees.

c. Simplicial Approximation and the Lefschetz Fixed Point Theorem

Many spaces of interest in algebraic topology can be given the structure of simplicial complexes, and early in the history of the subject this structure was exploited as one of the main technical tools. Later, CW complexes largely superceded simplicial complexes in this role, but there are still some occasions when the extra structure of simplicial complexes can be quite handy. This will be illustrated nicely in our proof of the classical Lefschetz fixed point theorem, below.

One of the nicest features of simplicial complexes is that maps can be deformed to preserve some of the linear structure within simplices. This is the idea of 'simplicial approximation,' which goes back all the way to Brouwer and Alexander before 1920. To state this result in its most basic form we first need a definition. If K and L are simplicial complexes, then a map $f : K \to L$ is **simplicial** if it sends each simplex of K to a simplex of L by a linear map taking vertices to vertices. Since a linear map from a simplex to a simplex is uniquely determined by its values on vertices, this means that a simplicial map is uniquely determined by its values on vertices. Also, it is easy to see that a map from the vertices of K to the vertices of L extends to a simplicial map iff it sends the vertices of each simplex of K to the vertices of some simplex of L.

Here then is the **Simplicial Approximation Theorem**:

Theorem 2.34. *If K is a finite simplicial complex and L is an arbitrary simplicial complex, then any map $f : K \to L$ is homotopic to a map which is simplicial with respect to some iterated barycentric subdivision of K.*

Allowing subdivision of K is essential. For if we fix the simplicial structure on K and choose L also to be a finite simplicial complex, then there are only finitely many simplicial maps $K \to L$ since there are only finitely many ways to map the vertices of K to the vertices of L. Taking K and L to be simplicial structures on S^n, this means that simplicial maps $K \to L$ can have only finitely many degrees. This is still true if we allow the simplicial structure L to vary, since if L has more vertices than K then a simplicial map $K \to L$ cannot be surjective, hence must have degree zero, and there are only finitely many simplicial complexes L having no more vertices than K.

Before proving the simplicial approximation theorem we need some terminology and a lemma. The **star** $St(\sigma)$ of a simplex σ in a simplicial complex X is defined to be the subcomplex consisting of all the simplices of X which contain σ. Closely related to this is the **open star** $st(\sigma)$, which is the union of the interiors of all simplices containing σ,

where the **interior** of a simplex τ is by definition $\tau - \partial\tau$. Thus $st(\sigma)$ is an open set in X whose closure is $St(\sigma)$.

Lemma 2.35. *For vertices $v_1, \cdots, v_n$ of X, the intersection $st(v_1) \cap \cdots \cap st(v_n)$ is empty unless $v_1, \cdots, v_n$ are the vertices of a simplex σ of X, in which case $st(v_1) \cap \cdots \cap st(v_n) = st(\sigma)$.*

Proof: $st(v_1) \cap \cdots \cap st(v_n)$ consists of the interiors of all simplices τ whose vertex set contains $\{v_1, \cdots, v_n\}$. If $st(v_1) \cap \cdots \cap st(v_n) \neq \varnothing$, such a τ exists and contains the simplex $\sigma = [v_1, \cdots, v_n] \subset K$. Since the simplices τ containing $\{v_1, \cdots, v_n\}$ are just the simplices containing σ, we obtain the result that $st(v_1) \cap \cdots \cap st(v_n) = st(\sigma)$. □

Proof of 2.34: Choose a metric on K which restricts to the standard Euclidean metric on each simplex of K. For example, K can be viewed as a subcomplex of a simplex Δ^N whose vertices are all the vertices of K, and we can restrict a standard metric on Δ^N to give a metric on K. Let ε be a Lebesgue number for the open cover $\{ f^{-1}\big(st(w)\big) \mid w \text{ is a vertex of } L \}$ of K. After iterated barycentric subdivision of K we may assume that each simplex has diameter less than $\varepsilon/2$. The star of each vertex v of K then has diameter less than ε, hence this star maps by f to the open star of some vertex $g(v)$ of L. The resulting map $g\colon K^0 \to L^0$ thus satisfies $f\big(St(v)\big) \subset st\big(g(v)\big)$ for all vertices v of K.

To see that g extends to a simplicial map $g\colon K \to L$, consider a simplex $[v_1, \cdots, v_n]$ of K. An interior point x of this simplex lies in $st(v_i)$ for each i, so $f(x)$ lies in $st\big(g(v_i)\big)$ for each i, since $f\big(st(v_i)\big) \subset st\big(g(v_i)\big)$ by the definition of $g(v_i)$. Thus $st\big(g(v_1)\big) \cap \cdots \cap st\big(g(v_n)\big) \neq \varnothing$, so $[g(v_1), \cdots, g(v_n)]$ is a simplex of L by the lemma, and we can extend g linearly over $[v_1, \cdots, v_n]$. Both $f(x)$ and $g(x)$ lie in a single simplex of L since $g(x)$ lies in $[g(v_1), \cdots, g(v_n)]$ and $f(x)$ lies in the star of this simplex. So taking the linear path $(1-t)f(x)+tg(x)$, $0 \leq t \leq 1$, in the simplex containing $f(x)$ and $g(x)$ defines a homotopy from f to the simplicial map g. To check continuity of this homotopy it suffices to restrict to the simplex $[v_1, \cdots, v_n]$, where continuity is clear since $f(x)$ varies continuously in the star of $[g(v_1), \cdots, g(v_n)]$ and $g(x)$ varies continuously in $[g(v_1), \cdots, g(v_n)]$. □

The proof makes it clear that the simplicial approximation g can be chosen not just homotopic to f but also close to f if we allow subdivisions of L as well as K.

The Lefschetz Fixed Point Theorem

This very classical application of homology is a considerable generalization of the Brouwer fixed point theorem, and it is also related to the Euler characteristic formula.

For a homomorphism $\varphi : \mathbb{Z}^n \to \mathbb{Z}^n$ with matrix $[a_{ij}]$, the trace $tr(\varphi)$ is defined to be $\sum_i a_{ii}$, the sum of the diagonal elements of $[a_{ij}]$. Since $tr([a_{ij}][b_{ij}]) = tr([b_{ij}][a_{ij}])$, conjugate matrices have the same trace, and it follows that $tr(\varphi)$ is independent of the choice of basis for $\mathbb{Z}^n$. For a homomorphism $\varphi : A \to A$ of a finitely generated abelian group A we can then define $tr(\varphi)$ to be the trace of the induced homomorphism $\overline{\varphi} : A/Torsion \to A/Torsion$.

For a map $f : X \to X$ of a finite CW complex X, or more generally any space whose homology groups are finitely generated and vanish in high dimensions, the **Lefschetz number** $\tau(f)$ is defined to be $\sum_n (-1)^n tr(f_* : H_n(X) \to H_n(X))$. For example, if f is the identity, or is homotopic to the identity, then $\tau(f)$ is the Euler characteristic $\chi(X)$.

Here is the Lefschetz Fixed Point Theorem:

Theorem 2.36. *If X is a finite simplicial complex, or more generally a retract of a finite simplicial complex, and $f : X \to X$ is a map with $\tau(f) \neq 0$, then f has a fixed point.*

As we show in the Appendix, every compact, locally contractible space which can be embedded in a Euclidean space of some dimension is a retract of a finite simplicial complex. This includes compact manifolds and finite CW complexes, for example. The compactness hypothesis is essential, since a translation of $\mathbb{R}$ has $\tau = 1$ but no fixed points. For an example showing that local properties are also significant, let X be the compact subspace of $\mathbb{R}^2$ consisting of two concentric circles together with a copy of $\mathbb{R}$ between them whose two ends spiral into the two circles, wrapping around them infinitely often, and let $f : X \to X$ be a homeomorphism translating the $\mathbb{R}$ along itself and rotating the circles, with no fixed points. Since f is homotopic to the identity, we have $\tau(f) = \chi(X)$, which equals 1 since the three path components of X are two circles and a line.

If X has the same homology groups as a point, at least modulo torsion, the theorem says that every map has a fixed point. This holds for example for $\mathbb{R}\mathrm{P}^n$ if n is even. The case of projective spaces is interesting because of its connection with linear algebra. An invertible linear transformation $f : \mathbb{R}^n \to \mathbb{R}^n$ takes lines through 0 to lines through 0, hence induces a map $\overline{f} : \mathbb{R}\mathrm{P}^{n-1} \to \mathbb{R}\mathrm{P}^{n-1}$. Fixed points of $\overline{f}$ are equivalent to eigenvectors of f. The characteristic polynomial of f has odd degree if n is odd, hence has a real root, so an eigenvector exists in this case. This is in agreement with the observation above that every map $\mathbb{R}\mathrm{P}^{2k} \to \mathbb{R}\mathrm{P}^{2k}$ has a fixed point. On the other hand the rotation of $\mathbb{R}^{2k}$ defined by $f(x_1, \cdots, x_{2k}) = (x_2, -x_1, x_4, -x_3, \cdots, x_{2k}, -x_{2k-1})$, has no eigenvectors and its projectivization $\overline{f} : \mathbb{R}\mathrm{P}^{2k-1} \to \mathbb{R}\mathrm{P}^{2k-1}$ has no fixed points.

Similarly, in the complex case an invertible linear transformation $f : \mathbb{C}^n \to \mathbb{C}^n$ induces $\overline{f} : \mathbb{C}\mathrm{P}^{n-1} \to \mathbb{C}\mathrm{P}^{n-1}$, and this always has a fixed point since the characteristic

polynomial always has a complex root. Nevertheless, as in the real case there is a map $\overline{f}:\mathbb{C}\mathrm{P}^{2k-1}\to\mathbb{C}\mathrm{P}^{2k-1}$ without fixed points. Namely, consider the map $f:\mathbb{C}^{2k}\to\mathbb{C}^{2k}$ defined by $f(z_1,\cdots,z_{2k})=(\overline{z}_2,-\overline{z}_1,\overline{z}_4,-\overline{z}_3,\cdots,\overline{z}_{2k},-\overline{z}_{2k-1})$. This is only 'conjugate-linear' over $\mathbb{C}$, but this is still good enough to imply that f induces a well-defined map $\overline{f}$ on $\mathbb{C}\mathrm{P}^{2k-1}$. Then it is an easy algebra exercise to check that $\overline{f}$ has no fixed points. The similarity between the real and complex cases persists in the fact that every map $\mathbb{C}\mathrm{P}^{2k}\to\mathbb{C}\mathrm{P}^{2k}$ has a fixed point, though to deduce this one needs more structure than just homology, so this will be left as an exercise for §3.2, using cup products.

One could go further and consider the quaternionic case. Oddly enough, every map $\mathbb{H}\mathrm{P}^n\to\mathbb{H}\mathrm{P}^n$ has a fixed point if $n>1$, according to an exercise in §4.4k. When $n=1$ the antipodal map of $S^3=\mathbb{H}\mathrm{P}^1$ has no fixed points.

Proof of 2.36: The general case easily reduces to the case of finite simplicial complexes, for suppose $r:K\to X$ is a retraction of the finite simplicial complex K onto X. For $f:X\to X$, the map $fr:K\to X\subset K$ then has exactly the same fixed points as f. And since $r_*:H_n(K)\to H_n(X)$ is projection onto a direct summand, we clearly have $tr(f_*r_*)=tr(f_*)$, so $\tau(f_*r_*)=\tau(f_*)$.

For X a finite simplicial complex, suppose $f:X\to X$ has no fixed points. We claim there is a subdivision L of X, a further subdivision K of L, and a simplicial map $g:K\to L$ homotopic to f such that $g(\sigma)\cap\sigma=\varnothing$ for each simplex σ of K. To see this, first choose a metric d on X as in the proof of the preceding proposition. Since f has no fixed points, $d\big(x,f(x)\big)>0$ for all $x\in X$, so by the compactness of X there is an $\varepsilon>0$ such that $d\big(x,f(x)\big)>\varepsilon$ for all x. Choose a subdivision L of X so that the stars of all simplices have diameter less than $\varepsilon/2$. By the proposition, there is a subdivision K of L and a simplicial map $g:K\to L$ homotopic to f. By construction, this g has the property that for each simplex σ of K, $f(\sigma)$ is contained in the star of the simplex $g(\sigma)$. We may assume the subdivision K is chosen fine enough so that its simplices all have diameter less than $\varepsilon/2$. Then $g(\sigma)\cap\sigma=\varnothing$ for each simplex σ of K since for $x\in\sigma$, σ lies within distance $\varepsilon/2$ of x and $g(\sigma)$ lies within distance $\varepsilon/2$ of $f(x)$, while $d\big(x,f(x)\big)>\varepsilon$.

For such a $g:K\to L$, the Lefschetz numbers $\tau(f)$ and $\tau(g)$ are equal since f and g are homotopic. Since g is simplicial, it takes the n-skeleton K^n of K to the n-skeleton L^n of L, for each n. Since K is a subdivision of L, L^n is contained in K^n, and hence $g(K^n)\subset K^n$ for all n. Thus g induces a chain map of the cellular chain complex $\{H_n(K^n,K^{n-1})\}$ to itself.

The next thing to observe is that

$$\tau(g)=\sum_n(-1)^n tr\big(g_*:H_n(K^n,K^{n-1})\to H_n(K^n,K^{n-1})\big).$$

This is the analog of Theorem 2.27 for trace instead of rank, and is proved in precisely the same way, based on the elementary algebraic fact that trace is additive for endomorphisms of short exact sequences, i.e., given a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\ 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \end{array}$$

with exact rows, then $tr(\beta) = tr(\alpha) + tr(\gamma)$. This algebraic fact can be proved by reducing to the easy case that A, B, and C are free by first factoring out the torsion in B, hence also in A, then eliminating any remaining torsion in C by replacing A by a larger subgroup $A' \subset B$, with A having finite index in A'. The details of this argument are left to the reader.

Finally, note that $g_* : H_n(K^n, K^{n-1}) \to H_n(K^n, K^{n-1})$ has trace 0 since the matrix for g_* has zeros down the diagonal, in view of the fact that $g(\sigma) \cap \sigma = \varnothing$ for each n-simplex σ. So $\tau(f) = \tau(g) = 0$. □

Fixed point theory is a well-developed side branch of algebraic topology which will not be touched upon much in this book. An engaging introduction can be found in [R.Brown].

Simplicial Approximations to CW Complexes

The simplicial approximation theorem concerns maps, and there is a somewhat analogous statement about spaces:

Theorem 2.37. *Every CW complex X is homotopy equivalent to a simplicial complex, of the same dimension as X, finite if X is finite, and countable if X is countable.*

We will build a simplicial complex $Y \simeq X$ inductively as an increasing union of subcomplexes Y_n with $Y_n \simeq X^n$. For the inductive step, assuming we have already constructed $Y_n \simeq X^n$, let e^{n+1} be an $(n+1)$-cell of X attached by a map $\varphi : S^n \to X^n$. The map $S^n \to Y_n$ corresponding to φ under the homotopy equivalence $Y_n \simeq X^n$ is homotopic to a simplicial map $f : S^n \to Y_n$ by the simplicial approximation theorem, and the mapping cone C_f is homotopy equivalent to $C_\varphi = X^n \cup e^{n+1}$. If we knew that the mapping cone of a simplicial map was a simplicial complex, then by performing the same construction for all the $(n+1)$-cells of X we would have completed the induction step. Unfortunately, and somewhat surprisingly, mapping cones and mapping cylinders are rather awkward objects in the simplicial category. To avoid this awkwardness we will instead construct simplicial analogs of mapping cones and cylinders which have all the essential features of actual mapping cones and cylinders.

Let us first construct the simplicial analog of a mapping cylinder. For a simplicial map $f:K\to L$ this will be a simplicial complex $M(f)$ containing L and the barycentric subdivision K' of K as subcomplexes, such that there is a deformation retraction r_t of $M(f)$ onto L with $r_1\big|K'=f$. The construction proceeds one simplex of K at a time, by induction on dimension. To begin, the ordinary mapping cylinder of $f\big|K^0$ suffices for $M(f\big|K^0)$. Assume inductively that we have already constructed $M(f\big|K^{n-1})$. Let σ be an n-simplex of K and let $\tau=f(\sigma)$, a simplex of L of dimension n or less. In particular, we assume we have already constructed $M(f:\partial\sigma\to\tau)$ with the desired properties, and then we let $M(f:\sigma\to\tau)$ be the cone on $M(f:\partial\sigma\to\tau)$. The figure below shows the case of a surjective linear map $\Delta^2\to\Delta^1$.

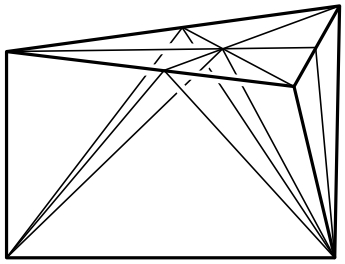

The space $M(f:\partial\sigma\to\tau)$ is contractible since by induction it deformation retracts onto τ which is contractible. The cone $M(f:\sigma\to\tau)$ is of course contractible, so the inclusion of $M(f:\partial\sigma\to\tau)$ into $M(f:\sigma\to\tau)$ is a homotopy equivalence. This implies that $M(f:\sigma\to\tau)$ deformation retracts onto $M(f:\partial\sigma\to\tau)$ by Corollary 0.5 (or one can give a direct argument using the fact that $M(f:\partial\sigma\to\tau)$ is contractible). Performing these constructions for all n-simplices σ, we obtain $M(f\big|K^n)$ with a deformation retraction onto $M(f\big|K^{n-1})$. Taking the union over all n yields $M(f)$ with a deformation retraction r_t onto L, namely the infinite concatenation of the previous deformation retractions, with the deformation retraction of $M(f\big|K^n)$ onto $M(f\big|K^{n-1})$ performed in the t-interval $[1/2^{n+2},1/2^{n+1}]$. The map $r_1\big|K$ produced by this process may not equal f, but it is homotopic to f via the linear homotopy $tf+(1-t)r_1$, which is defined since $r_1(\sigma)\subset f(\sigma)$ for all simplices σ of K. By applying the homotopy extension property to the homotopy of r_1 which equals $tf+(1-t)r_1$ on K and the identity map on L, we can improve our deformation retraction of $M(f)$ onto L so that its restriction to K at time 1 is f.

From the simplicial analog $M(f)$ of a mapping cylinder we construct the simplicial 'mapping cone' $C(f)$ by attaching the ordinary cone on K' to the subcomplex $K'\subset M(f)$.

Proof of 2.37: We will construct for each n a CW complex Z_n containing X^n as a deformation retract and also containing as a deformation retract a subcomplex Y_n which is a simplicial complex. Beginning with $Y_0=Z_0=X^0$, suppose inductively that we have

already constructed Y_n and Z_n. Let the cells e^{n+1}_α of X be attached by maps $\varphi_\alpha : S^n \to X^n$. Using the simplicial approximation theorem, there is a homotopy from φ_α to a simplicial map $f_\alpha : S^n \to Y_n$. The CW complex $W_n = Z_n \bigcup_\alpha M(f_\alpha)$ contains a simplicial subcomplex S^n_α homeomorphic to S^n at one end of $M(f_\alpha)$, and the homeomorphism $S^n \approx S^n_\alpha$ is homotopic in W_n to the map f_α, hence also to φ_α. Let Z_{n+1} be obtained from Z_n by attaching $D^{n+1}_\alpha \times I$'s via these homotopies between the φ_α's and the inclusions $S^n_\alpha \hookrightarrow W_n$. Thus Z_{n+1} contains X^{n+1} at one end, and at the other end we have a simplicial complex $Y_{n+1} = Y_n \bigcup_\alpha C(f_\alpha)$, where $C(f_\alpha)$ is obtained from $M(f_\alpha)$ by attaching a cone on the subcomplex S^n_α. Since $D^{n+1} \times I$ deformation retracts onto $\partial D^{n+1} \times I \cup D^{n+1} \times \{0\}$, we see that Z_{n+1} deformation retracts onto $Z_n \cup Y_{n+1}$, which in turn deformation retracts onto $Y_n \cup Y_{n+1} = Y_{n+1}$ by induction. Likewise, Z_{n+1} deformation retracts onto $X^{n+1} \cup W_n$ which deformation retracts onto $X^{n+1} \cup Z_n$ and hence onto $X^{n+1} \cup X^n = X^{n+1}$ by induction.

Let $Y = \bigcup_n Y_n$ and $Z = \bigcup_n Z_n$. The deformation retractions of Z_n onto X^n give deformation retractions of $X \cup Z_n$ onto X, and the infinite concatenation of the latter deformation retractions is a deformation retraction of Z onto X. Similarly, Z deformation retracts onto Y. □

Exercises

1. Use the Fixed Point Theorem to show that a map $S^n \to S^n$ has a fixed point unless its degree is equal to the degree of the antipodal map $x \mapsto -x$.

2. Verify that the map $f : \mathbb{C}^{2k} \to \mathbb{C}^{2k}$, $f(z_1, \cdots, z_{2k}) = (\overline{z}_2, -\overline{z}_1, \overline{z}_4, -\overline{z}_3, \cdots, \overline{z}_{2k}, -\overline{z}_{2k-1})$, induces a well-defined map $\mathbb{C}\mathrm{P}^{2k-1} \to \mathbb{C}\mathrm{P}^{2k-1}$ without fixed points.

3. If X is a finite simplicial complex and $f : X \to X$ is a simplicial homeomorphism, show that the Lefschetz number $\tau(f)$ equals the Euler characteristic of the set of fixed points of f. In particular, $\tau(f)$ is the number of fixed points if the fixed points are isolated. [Hint: Barycentrically subdivide X to make the fixed point set a subcomplex.]

4. Let M be a closed orientable surface embedded in $\mathbb{R}^3$ in such a way that reflection across a plane P defines a homeomorphism $r : M \to M$ fixing $M \cap P$, a collection of circles. Is it possible to homotope r to have no fixed points?

5. Let X be homotopy equivalent to a finite simplicial complex and let Y be homotopy equivalent to a finite or countably infinite simplicial complex. Using the simplicial approximation theorem, show that there are at most countably many homotopy classes of maps $X \to Y$.

6. Show that there are only countably many homotopy types of finite CW complexes.

4. Additional Topics

a. The Universal Coefficient Theorem for Homology

In this section we derive the general algebraic formula which computes homology with arbitrary coefficients in terms of homology with $\mathbb{Z}$ coefficients. The theory parallels rather closely the universal coefficient theorem for cohomology in §3.1.

The definition of homology with coefficients in §2.2 can be reformulated in terms of tensor products, so let us first review the basic facts about tensor products. For abelian groups A and B the tensor product $A\otimes B$ is defined to be the abelian group with generators $a\otimes b$ for $a\in A$, $b\in B$, and relations $(a+a')\otimes b = a\otimes b + a'\otimes b$ and $a\otimes(b+b') = a\otimes b + a\otimes b'$. So the zero element of $A\otimes B$ is $0\otimes 0 = 0\otimes b = a\otimes 0$, and $-(a\otimes b) = -a\otimes b = a\otimes(-b)$. Some readily verified elementary properties are:

(1) $A\otimes B \approx B\otimes A$
(2) $(\bigoplus_i A_i)\otimes B \approx \bigoplus_i (A_i\otimes B)$
(3) $(A\otimes B)\otimes C \approx A\otimes(B\otimes C)$
(4) $\mathbb{Z}\otimes A \approx A$
(5) $\mathbb{Z}_n\otimes A \approx A/nA$
(6) Homomorphisms $f: A\to A'$ and $g: B\to B'$ induce a homomorphism $f\otimes g: A\otimes B\to A'\otimes B'$ via $(f\otimes g)(a\otimes b) = f(a)\otimes g(b)$.
(7) A map $\varphi: A\times B\to C$ which is a homomorphism in each variable separately induces a homomorphism $A\otimes B\to C$ sending $a\otimes b$ to $\varphi(a,b)$.

In (1)-(5) the isomorphisms are the obvious ones, e.g., $a\otimes b\mapsto b\otimes a$ in (1) and $n\otimes a\mapsto na$ in (4). Properties (1), (2), (4), and (5) allow the calculation of tensor products of finitely generated abelian groups.

The chain group $C_n(X,A;G)$ defined in §2.2 is a direct sum of copies of G, one for each singular simplex in X not contained in A. From (2) and (4) it follows that $C_n(X,A;G)$ is naturally isomorphic to $C_n(X,A)\otimes G$. Under this isomorphism the boundary map $C_n(X,A;G)\to C_{n-1}(X,A;G)$ corresponds to $\partial\otimes 1\!\!1: C_n(X,A)\otimes G\to C_{n-1}(X,A)\otimes G$ where $\partial: C_n(X,A)\to C_{n-1}(X,A)$ is the usual boundary map for $\mathbb{Z}$ coefficients. Thus we have the following algebraic problem:

> Given a chain complex $\cdots\longrightarrow C_n \xrightarrow{\partial_n} C_{n-1}\longrightarrow\cdots$ of free abelian groups C_n, is it possible to compute the homology groups $H_n(C;G)$ of the associated chain complex $\cdots\longrightarrow C_n\otimes G \xrightarrow{\partial_n\otimes 1\!\!1} C_{n-1}\otimes G\longrightarrow\cdots$ just in terms of G and the homology groups $H_n(C)$ of the original complex?

To approach this problem, the idea will be to compare the chain complex C with two simpler subcomplexes, consisting of the cycles and the boundaries in C, and see what happens upon tensoring all three complexes with G.

Let $Z_n = Ker(\partial_n) \subset C_n$ and $B_n = Im(\partial_{n+1}) \subset C_n$. The restrictions of ∂_n to these two subgroups are zero, so they can be regarded as subcomplexes Z and B of C with trivial boundary maps. Thus we have a short exact sequence of chain complexes, part of which is:

$$(i) \qquad \begin{array}{ccccccccc} 0 & \longrightarrow & Z_n & \longrightarrow & C_n & \xrightarrow{\partial_n} & B_{n-1} & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \partial_n} & & \downarrow{\scriptstyle \partial_n} & & \downarrow{\scriptstyle \partial_{n-1}} & & \\ 0 & \longrightarrow & Z_{n-1} & \longrightarrow & C_{n-1} & \xrightarrow{\partial_{n-1}} & B_{n-2} & \longrightarrow & 0 \end{array}$$

The rows in this diagram split since each B_n is free, being a subgroup of the free group C_n. Thus $C_n \approx Z_n \oplus B_{n-1}$, but the chain complex C is not the direct sum of the chain complexes Z and B since the latter have trivial boundary maps but the boundary maps in C may be nontrivial. Now tensor with G to get a commutative diagram

$$(ii) \qquad \begin{array}{ccccccccc} 0 & \longrightarrow & Z_n \otimes G & \longrightarrow & C_n \otimes G & \xrightarrow{\partial_n \otimes \mathbb{1}} & B_{n-1} \otimes G & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \partial_n \otimes \mathbb{1}} & & \downarrow{\scriptstyle \partial_n \otimes \mathbb{1}} & & \downarrow{\scriptstyle \partial_{n-1} \otimes \mathbb{1}} & & \\ 0 & \longrightarrow & Z_{n-1} \otimes G & \longrightarrow & C_{n-1} \otimes G & \xrightarrow{\partial_{n-1} \otimes \mathbb{1}} & B_{n-2} \otimes G & \longrightarrow & 0 \end{array}$$

The rows are exact since the rows in (i) split and tensor products satisfy $(A \oplus B) \otimes G \approx A \otimes G \oplus B \otimes G$, so the rows in (ii) are split exact sequences too. Thus we have a short exact sequence of chain complexes $0 \to Z \otimes G \to C \otimes G \to B \otimes G \to 0$. Since the boundary maps are trivial in $Z \otimes G$ and $B \otimes G$, the associated long exact sequence of homology groups has the form

$$(iii) \qquad \cdots \longrightarrow B_n \otimes G \longrightarrow Z_n \otimes G \longrightarrow H_n(C;G) \longrightarrow B_{n-1} \otimes G \longrightarrow Z_{n-1} \otimes G \longrightarrow \cdots$$

The 'boundary' maps $B_n \otimes G \to Z_n \otimes G$ in this sequence are the obvious ones, $i_n \otimes \mathbb{1}$, where $i_n : B_n \to Z_n$ is the inclusion. This is evident from the definition of the boundary map in a long exact sequence of homology groups: In the diagram (ii) one takes an element of $B_{n-1} \otimes G$, pulls it back via $(\partial_n \otimes \mathbb{1})^{-1}$ to $C_n \otimes G$, then applies $\partial_n \otimes \mathbb{1}$ to get into $C_{n-1} \otimes G$, then pulls back to $Z_{n-1} \otimes G$.

The long exact sequence (iii) can be broken up into short exact sequences

$$(iv) \qquad 0 \longrightarrow Coker(i_n \otimes \mathbb{1}) \longrightarrow H_n(C;G) \longrightarrow Ker(i_{n-1} \otimes \mathbb{1}) \longrightarrow 0$$

where $Coker(i_n \otimes \mathbb{1}) = (Z_n \otimes G)/Im(i_n \otimes \mathbb{1})$. The next lemma shows that $Coker(i_n \otimes \mathbb{1})$ is just $H_n(C) \otimes G$.

Lemma 3.28. *If the sequence of abelian groups* $A \xrightarrow{i} B \xrightarrow{j} C \longrightarrow 0$ *is exact, then so is* $A\otimes G \xrightarrow{i\otimes \mathbb{1}} B\otimes G \xrightarrow{j\otimes \mathbb{1}} C\otimes G \longrightarrow 0$.

Proof: Certainly the compositions of two successive maps in the latter sequence are zero. Also, $j\otimes \mathbb{1}$ is clearly surjective since j is. To check exactness at $B\otimes G$, let us define a map $\varphi : C\times G \to B\otimes G/Im(i\otimes \mathbb{1})$ by $\varphi(c,g) = b\otimes g$ where $j(b) = c$. Note that φ is well-defined since if $j(b) = j(b') = c$ then $b\otimes g - b'\otimes g = (b-b')\otimes g \in Im(i\otimes \mathbb{1})$. Since φ is a homomorphism in each variable separately, it induces a homomorphism $C\otimes G \to B\otimes G/Im(i\otimes \mathbb{1})$ which is clearly an inverse to the natural quotient map $B\otimes G/Im(i\otimes \mathbb{1}) \to C\otimes G$. Thus the kernel of $j\otimes \mathbb{1}$ is exactly the image of $i\otimes \mathbb{1}$. □

It remains to understand $Ker(i_{n-1}\otimes \mathbb{1})$, or equivalently $Ker(i_n\otimes \mathbb{1})$. The situation is that tensoring the short exact sequence

$$(v) \qquad 0 \longrightarrow B_n \xrightarrow{i_n} Z_n \longrightarrow H_n(C) \longrightarrow 0$$

with G produces a sequence which becomes exact only by insertion of the extra term $Ker(i_n\otimes \mathbb{1})$:

$$(vi) \qquad 0 \longrightarrow Ker(i_n\otimes \mathbb{1}) \longrightarrow B_n\otimes G \xrightarrow{i_n\otimes \mathbb{1}} Z_n\otimes G \longrightarrow H_n(C)\otimes G \longrightarrow 0$$

What we will show is that $Ker(i_n\otimes \mathbb{1})$ does not really depend on B_n and Z_n but only on their quotient $H_n(C)$, and of course G.

In the terminology of §3.1, the sequence (v) is a free resolution of H, an exact sequence

$$\cdots \longrightarrow F_2 \xrightarrow{f_2} F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} H \longrightarrow 0$$

with each F_n free. Tensoring this free resolution with a fixed group G produces a chain complex

$$\cdots \longrightarrow F_1\otimes G \xrightarrow{f_1\otimes \mathbb{1}} F_0\otimes G \xrightarrow{f_0\otimes \mathbb{1}} H\otimes G \longrightarrow 0$$

By Lemma 3.28 this is exact at $F_0\otimes G$ and $H\otimes G$, but to the left of these two terms it may not be exact. For the moment let us write $H_n(F\otimes G)$ for the homology group $Ker(f_n\otimes \mathbb{1})/Im(f_{n+1}\otimes \mathbb{1})$.

Lemma 3.29. *For any two free resolutions* F *and* F' *of* H *there are canonical isomorphisms* $H_n(F\otimes G) \approx H_n(F'\otimes G)$ *for all* n.

Proof: We use Lemma 3.1(a). In the situation described there, tensoring the two free resolutions with G gives two chain complexes with chain maps $\alpha_n\otimes \mathbb{1}$ between them.

These induce homomorphisms on homology, $\alpha_*: H_n(F\otimes G) \to H_n(F'\otimes G)$, which do not depend on the choice of α_n's since if α_n and α_n' are chain homotopic via a chain homotopy λ_n then $\alpha_n \otimes \mathbb{1}$ and $\alpha_n' \otimes \mathbb{1}$ are chain homotopic via $\lambda_n \otimes \mathbb{1}$. These induced homomorphisms $\alpha_*: H_n(F\otimes G) \to H_n(F'\otimes G)$ obviously satisfy $(\beta\alpha)_* = \beta_*\alpha_*$ for a composition $H \xrightarrow{\alpha} H' \xrightarrow{\beta} H''$ with a free resolution F'' of H'' also given. In particular, if we take α to be an isomorphism with β its inverse, and $F'' = F$, then $\beta_*\alpha_* = (\beta\alpha)_* = \mathbb{1}_* = \mathbb{1}$, and similarly with β and α reversed, so α_* is an isomorphism. Specializing further, taking α to be the identity but with two different free resolutions F and F', we get a canonical isomorphism $\mathbb{1}_*: H_n(F\otimes G) \to H_n(F'\otimes G)$. □

The group $H_n(F\otimes G)$, which depends only on H and G, is denoted $\mathrm{Tor}_n(H,G)$. Since a free resolution $0 \to F_1 \to F_0 \to H \to 0$ always exists, as noted in §3.1, it follows that $\mathrm{Tor}_n(H,G) = 0$ for $n > 1$. Usually $\mathrm{Tor}_1(H,G)$ is written simply as $\mathrm{Tor}(H,G)$. As we shall see later, $\mathrm{Tor}(H,G)$ provides a measure of the common torsion of H and G, hence the name 'Tor.'

What is $\mathrm{Tor}_0(H,G)$? With the definition given above it would be zero since Lemma 3.28 implies that $F_1\otimes G \to F_0\otimes G \to H\otimes G \to 0$ is exact. It is probably better to modify the definition of $H_n(F\otimes G)$ to be the homology groups of the sequence $\cdots \to F_1\otimes G \to F_0\otimes G \to 0$, omitting the term $H\otimes G$ which can be regarded as a kind of augmentation. With this revised definition, Lemma 3.28 then gives an isomorphism $\mathrm{Tor}_0(H,G) \approx H\otimes G$.

We should remark that $\mathrm{Tor}(H,G)$ is a functor of both G and H: Homomorphisms $\alpha: H \to H'$ and $\beta: G \to G'$ induce homomorphisms $\alpha_*: \mathrm{Tor}(H,G) \to \mathrm{Tor}(H',G)$ and $\beta_*: \mathrm{Tor}(H,G) \to \mathrm{Tor}(H,G')$, satisfying $(\alpha\alpha')_* = \alpha_*\alpha_*'$, $(\beta\beta')_* = \beta_*\beta_*'$, and $\mathbb{1}_* = \mathbb{1}$. The induced map α_* was constructed in the proof of Lemma 3.29, while for β the construction of β_* is obvious.

Before going into calculations of $\mathrm{Tor}(H,G)$ let us finish analyzing the earlier exact sequence (iv). Recall that we have a chain complex C of free abelian groups, with homology groups denoted $H_n(C)$, and tensoring C with G gives another complex $C\otimes G$ whose homology groups are denoted $H_n(C;G)$. The following result is known as the **Universal Coefficient Theorem** since it describes homology with arbitrary coefficients in terms of homology with $\mathbb{Z}$ coefficients, which are therefore 'universal.'

Theorem 3.30. *If C is a chain complex of free abelian groups, then there are natural short exact sequences*

$$0 \longrightarrow H_n(C)\otimes G \longrightarrow H_n(C;G) \longrightarrow \mathrm{Tor}(H_{n-1}(C),G) \longrightarrow 0$$

for all n and all G, and these sequences split, though not naturally.

Naturality means that a chain map $C \to C'$ induces a map between the corresponding short exact sequences, with commuting squares.

Proof: The exact sequence in question is (iv) since we have shown that we can identify $Coker(i_n \otimes \mathbb{1})$ with $H_n(C) \otimes G$ and $Ker(i_{n-1})$ with $\mathrm{Tor}(H_{n-1}(C), G)$. Verifying the naturality of this sequence is a mental exercise in definition-checking which we leave to the reader.

The splitting is obtained as follows. We observed earlier that the short exact sequence $0 \to Z_n \to C_n \to B_{n-1} \to 0$ splits, so there is a projection $p : C_n \to Z_n$ which restricts to the identity on Z_n. The map p gives an extension of the quotient map $Z_n \to H_n(C)$ to a homomorphism $C_n \to H_n(C)$. Letting n vary, we then have a chain map $C \to H(C)$ where the groups $H_n(C)$ are regarded as a chain complex with trivial boundary maps, so the chain map condition is automatic. Now tensor with G to get a chain map $C \otimes G \to H(C) \otimes G$. Taking homology groups, we then have induced homomorphisms $H_n(C; G) \to H_n(C) \otimes G$ since the boundary maps in the chain complex $H(C) \otimes G$ are trivial. The homomorphisms $H_n(C; G) \to H_n(C) \otimes G$ give the desired splitting since at the level of chains they are the identity on cycles in C, by the definition of p. □

Corollary 3.31. *For each pair of spaces (X, A) there are split exact sequences*

$$0 \longrightarrow H_n(X, A) \otimes G \longrightarrow H_n(X, A; G) \longrightarrow \mathrm{Tor}(H_{n-1}(X, A), G) \longrightarrow 0$$

for all n, and these sequences are natural with respect to maps $(X, A) \to (Y, B)$. □

The splitting is not natural, for if it were, a map $X \to Y$ which induced trivial maps $H_n(X) \to H_n(Y)$ and $H_{n-1}(X) \to H_{n-1}(Y)$ would have to induce the trivial map $H_n(X; G) \to H_n(Y; G)$ for all G, but in §2.2 we saw an example where this fails, namely the quotient map $M(\mathbb{Z}_m, n) \to S^{n+1}$ with $G = \mathbb{Z}_m$.

The basic tools for computing Tor are given by:

Proposition 3.32.

(1) $\mathrm{Tor}(A, B) \approx \mathrm{Tor}(B, A)$.

(2) $\mathrm{Tor}(\bigoplus_i A_i, B) \approx \bigoplus_i \mathrm{Tor}(A_i, B)$.

(3) $\mathrm{Tor}(A, B) = 0$ *if A or B is free, or more generally torsionfree.*

(4) $\mathrm{Tor}(A, B) \approx \mathrm{Tor}(T(A), B)$ *where $T(A)$ is the torsion subgroup of A.*

(5) $\mathrm{Tor}(\mathbb{Z}_n, A) \approx Ker(A \xrightarrow{n} A)$.

(6) *For each short exact sequence $0 \to B \to C \to D \to 0$ there is a naturally associated exact sequence*

$$0 \to \mathrm{Tor}(A, B) \to \mathrm{Tor}(A, C) \to \mathrm{Tor}(A, D) \to A \otimes B \to A \otimes C \to A \otimes D \to 0$$

Proof: (2) is easy since one can choose as a free resolution of $\bigoplus_i A_i$ the direct sum of free resolutions of the A_i's.

For (3), if A is free it has a free resolution with $F_n = 0$ for $n \geq 1$, so $\text{Tor}(A, B) = 0$ for all B. On the other hand, if B is free, then tensoring a free resolution of A with B preserves exactness, since tensoring a sequence with a direct sum of $\mathbb{Z}$'s produces just a direct sum of copies of the given sequence. So $\text{Tor}(A, B) = 0$ in this case too. The generalization to torsionfree A or B will be given below.

Property (5) comes from tensoring the free resolution $0 \to \mathbb{Z} \xrightarrow{n} \mathbb{Z} \to \mathbb{Z}_n \to 0$ with A.

For (6), choose a free resolution $0 \to F_1 \to F_0 \to A \to 0$ and tensor with the given short exact sequence to get a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & F_1 \otimes B & \longrightarrow & F_1 \otimes C & \longrightarrow & F_1 \otimes D & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & F_0 \otimes B & \longrightarrow & F_0 \otimes C & \longrightarrow & F_0 \otimes D & \longrightarrow & 0 \end{array}$$

The rows are exact since tensoring with a free group preserves exactness. Extending the three columns by zeros above and below, we then have a short exact sequence of chain complexes whose associated long exact sequence of homology groups is the desired six-term exact sequence.

To prove (1) we apply (6) to a free resolution $0 \to F_1 \to F_0 \to B \to 0$. Since $\text{Tor}(A, F_1)$ and $\text{Tor}(A, F_0)$ vanish by the part of (3) which we have proved, the six-term sequence in (6) reduces to the first row of the following diagram:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \text{Tor}(A, B) & \longrightarrow & A \otimes F_1 & \longrightarrow & A \otimes F_0 & \longrightarrow & A \otimes B \longrightarrow 0 \\ & & & & \downarrow \approx & & \downarrow \approx & & \downarrow \approx \\ 0 & \longrightarrow & \text{Tor}(B, A) & \longrightarrow & F_1 \otimes A & \longrightarrow & F_0 \otimes A & \longrightarrow & B \otimes A \longrightarrow 0 \end{array}$$

The second row comes from the definition of $\text{Tor}(B, A)$. The vertical isomorphisms come from the natural commutativity of tensor product. Since the squares commute, there is induced a map $\text{Tor}(A, B) \to \text{Tor}(B, A)$, which is an isomorphism by the five-lemma.

Now we can prove (3) in the torsionfree case. For a free resolution $0 \longrightarrow F_1 \xrightarrow{\varphi} F_0 \longrightarrow A \longrightarrow 0$ we wish to show that $\varphi \otimes \mathbb{1} : F_1 \otimes B \to F_0 \otimes B$ is injective if B is torsionfree. Suppose $\sum_i x_i \otimes b_i$ lies in the kernel of $\varphi \otimes \mathbb{1}$. This means that $\sum_i \varphi(x_i) \otimes b_i$ can be reduced to 0 by a finite number of applications of the defining relations for tensor products. Only a finite number of elements of B are involved in this process. These lie in a finitely generated subgroup $B_0 \subset B$, so $\sum_i x_i \otimes b_i$ lies in the kernel of $\varphi \otimes \mathbb{1} : F_1 \otimes B_0 \to F_0 \otimes B_0$.

This kernel is zero since $\text{Tor}(A, B_0) = 0$, B_0 being finitely generated and torsionfree hence free.

Now (4) follows by applying (6) to the short exact sequence $0 \to T(A) \to A \to A/T(A) \to 0$ since $A/T(A)$ is torsionfree. □

In particular, (5) gives $\text{Tor}(\mathbb{Z}_m, \mathbb{Z}_n) \approx \mathbb{Z}_q$ where q is the greatest common divisor of m and n. Thus $\text{Tor}(\mathbb{Z}_m, \mathbb{Z}_n)$ is isomorphic to $\mathbb{Z}_m \otimes \mathbb{Z}_n$, though somewhat by accident. Combining this isomorphism with (2) and (3) we see that for finitely generated A and B, $\text{Tor}(A, B)$ is isomorphic to the tensor product of the torsion subgroups of A and B, or roughly speaking, 'the common torsion of A and B.' This is one reason for the 'Tor' designation, further justification being (3) and (4).

Homology calculations are often simplified by taking coefficients in a field, usually $\mathbb{Q}$ or $\mathbb{Z}_p$ for p prime. In general this gives less information than taking $\mathbb{Z}$ coefficients, but still some of the essential features are retained, as the following result indicates:

Corollary 3.33.

(a) *$H_n(X; \mathbb{Q}) \approx H_n(X; \mathbb{Z}) \otimes \mathbb{Q}$, so when $H_n(X; \mathbb{Z})$ is finitely generated, the dimension of $H_n(X; \mathbb{Q})$ as a vector space over $\mathbb{Q}$ equals the rank of $H_n(X; \mathbb{Z})$.*

(b) *If $H_n(X; \mathbb{Z})$ and $H_{n-1}(X; \mathbb{Z})$ are finitely generated, then for p prime, $H_n(X; \mathbb{Z}_p)$ consists of:*

  (*i*) *a $\mathbb{Z}_p$ summand for each $\mathbb{Z}$ summand of $H_n(X; \mathbb{Z})$*

  (*ii*) *a $\mathbb{Z}_p$ summand for each $\mathbb{Z}_{p^k}$ summand in $H_n(X; \mathbb{Z})$, $k \geq 1$*

  (*iii*) *a $\mathbb{Z}_p$ summand for each $\mathbb{Z}_{p^k}$ summand in $H_{n-1}(X; \mathbb{Z})$, $k \geq 1$.* □

Even in the case of nonfinitely generated homology groups, field coefficients still give good qualitative information:

Corollary 3.34.

(a) *$\widetilde{H}_n(X; \mathbb{Z}) = 0$ for all n iff $\widetilde{H}_n(X; \mathbb{Q}) = 0$ and $\widetilde{H}_n(X; \mathbb{Z}_p) = 0$ for all n and all primes p.*

(b) *A map $f : X \to Y$ induces isomorphisms on homology with $\mathbb{Z}$ coefficients iff it induces isomorphisms on homology with $\mathbb{Q}$ and $\mathbb{Z}_p$ coefficients for all primes p.*

Proof: Statement (b) follows from (a) by passing to the mapping cone of f. The universal coefficient theorem gives the 'only if' half of (a). For the 'if' implication it suffices to show that if an abelian group A is such that $A \otimes \mathbb{Q} = 0$ and $\text{Tor}(A, \mathbb{Z}_p) = 0$ for all primes p, then $A = 0$. For the short exact sequences $0 \to \mathbb{Z} \xrightarrow{p} \mathbb{Z} \to \mathbb{Z}_p \to 0$ and

$0 \to \mathbb{Z} \to \mathbb{Q} \to \mathbb{Q}/\mathbb{Z} \to 0$, the six-term exact sequences in (6) of the proposition become

$$0 \longrightarrow \operatorname{Tor}(A, \mathbb{Z}_p) \longrightarrow A \xrightarrow{p} A \longrightarrow A \otimes \mathbb{Z}_p \longrightarrow 0$$
$$0 \longrightarrow \operatorname{Tor}(A, \mathbb{Q}/\mathbb{Z}) \longrightarrow A \longrightarrow A \otimes \mathbb{Q} \longrightarrow A \otimes \mathbb{Q}/\mathbb{Z} \longrightarrow 0$$

If $\operatorname{Tor}(A, \mathbb{Z}_p) = 0$ for all p, then exactness of the first sequence implies that $A \xrightarrow{p} A$ is injective for all p, so A is torsionfree. Then $\operatorname{Tor}(A, \mathbb{Q}/\mathbb{Z}) = 0$ by (3) or (4) of the proposition, so the second sequence implies that $A \to A \otimes \mathbb{Q}$ is injective, hence $A = 0$ if $A \otimes \mathbb{Q} = 0$. □

The algebra by means of which the Tor functor is derived from tensor products has a very natural generalization in which abelian groups are replaced by modules over a fixed ring R with identity. If R is commutative, the tensor product $A \otimes_R B$ of R-modules A and B is obtained from the quotient of the tensor product $A \otimes B$ of A and B regarded as abelian groups, by imposing the further relations $ra \otimes b = a \otimes rb$ for $r \in R$, $a \in A$, and $b \in B$. In case R is not commutative one assumes A is a right R-module and B is a left R-module, and the relation is written instead $ar \otimes b = a \otimes rb$. It is a small algebra exercise to see that $A \otimes_R B = A \otimes B$ when $R = \mathbb{Z}_m$ or $\mathbb{Q}$. But in general $A \otimes_R B$ is not the same as $A \otimes B$. For example, if $R = \mathbb{Q}(\sqrt{2})$, a 2-dimensional vector space over $\mathbb{Q}$, then $R \otimes_R R = R$ but $R \otimes R$ is a 4-dimensional vector space over $\mathbb{Q}$. And $\mathbb{R} \otimes \mathbb{R}$ is quite huge compared with $\mathbb{R} \otimes_{\mathbb{R}} \mathbb{R} = \mathbb{R}$.

Free resolutions of R-modules are defined just as for abelian groups, using free R-modules, which are direct sums of copies of R. The lemmas 3.28 and 3.29 carry over to this context without change, and so one has functors $\operatorname{Tor}_n^R(A, B)$. However it need not be true that $\operatorname{Tor}_n^R(A, B) = 0$ for $n > 1$. The reason this was true when $R = \mathbb{Z}$ was that subgroups of free groups are free, but submodules of free R-modules need not be free in general. If R is a principal ideal domain, submodules of free R-modules are free, so in this case the rest of the algebra, in particular the universal coefficient theorem, goes through without change.

When R is a field, modules are vector spaces, so every module is free, which simplifies the theory considerably. Tensoring an exact sequence with a free module clearly preserves exactness, since one is just taking the direct sum of copies of the sequence. Thus the sequence (ii) is automatically exact and one immediately has $H_n(C \otimes_F G) \approx H_n(C) \otimes_F G$ if F is a field.

Exercises

1. Show that $\text{Tor}(A, \mathbb{Q}/\mathbb{Z})$ is isomorphic to the torsion subgroup of A. Deduce that A is torsionfree iff $\text{Tor}(A, B) = 0$ for all B.

2. Show that if $\widetilde{H}^n(X; \mathbb{Q})$ and $\widetilde{H}^n(X; \mathbb{Z}_p)$ are zero for all n and all primes p, then $\widetilde{H}_n(X; \mathbb{Z}) = 0$ for all n, and hence $\widetilde{H}^n(X; G) = 0$ for all G and n.

3. Show that $\otimes$ and Tor commute with direct limits: $(\varinjlim A_\alpha) \otimes B = \varinjlim (A_\alpha \otimes B)$ and $\text{Tor}(\varinjlim A_\alpha, B) = \varinjlim \text{Tor}(A_\alpha, B)$.

4. From the fact that $\text{Tor}(A, B) = 0$ if A is free deduce that $\text{Tor}(A, B) = 0$ if A is torsionfree by applying the previous problem to the directed system of finitely generated subgroups A_α of A.

5. Show that $\text{Tor}(A, B)$ is always a torsion group, and that $\text{Tor}(A, B)$ contains an element of order n iff A and B contain elements of order n.

b. The Künneth Formula

Künneth formulas describe the homology or cohomology of a product space in terms of the homology or cohomology of the factors. In nice cases these formulas take the form $H_*(X \times Y; R) \approx H_*(X; R) \otimes H_*(Y; R)$ and $H^*(X \times Y; R) \approx H^*(X; R) \otimes H^*(Y; R)$ for a coefficient ring R. For the case of cohomology, such a formula was given in Theorem 3.9. Our aim in this section is to derive the general Künneth formula for homology, which works somewhat better than cohomology. The new feature of the general case is that an extra Tor term is needed to describe the full homology of a product.

The first big step is to define a **cross product** map

$$H_i(X; R) \times H_j(Y; R) \to H_{i+j}(X \times Y; R)$$

There are two ways to do this. One is a direct definition for singular homology, involving explicit simplicial formulas. More enlightening, however, is the definition in terms of cellular homology. This necessitates assuming X and Y are CW complexes, but this hypothesis can later be removed by the technique of CW approximation in §4.2. We shall focus therefore on the cellular definition, leaving the simplicial definition to the end of the section for those who are curious to see how it goes.

The key ingredient in the definition of the cellular cross product will be the fact that the cellular boundary map satisfies $d(e^i \times e^j) = de^i \times e^j + (-1)^i e^i \times de^j$. If we treat the symbol '$\times$' as a bilinear operation on cellular chains, then this implies more generally that $d(a \times b) = da \times b + (-1)^i a \times db$ where a is a cellular i-chain and b is a cellular j-chain. This formula implies that the cross product of two cycles is a cycle. Also, the product of a boundary and a cycle is a boundary since $da \times b = d(a \times b)$ if $db = 0$, and similarly $a \times db = (-1)^i d(a \times b)$ if $da = 0$. Hence there is an induced homomorphism $H_i(X; R) \times H_j(Y; R) \to H_{i+j}(X \times Y; R)$, which is the cross product in cellular homology. Since it is bilinear, it could also be viewed as a homomorphism $H_i(X; R) \otimes_R H_j(Y; R) \to H_{i+j}(X \times Y; R)$.

Our first task will be to determine the boundary maps in the cellular chain complex $C_*(X \times Y)$ for $X \times Y$ in terms of the boundary maps in the cellular chain complexes $C_*(X)$ and $C_*(Y)$. For simplicity we consider homology with $\mathbb{Z}$ coefficients here, but the case of other coefficients follows immediately from this special case. With $\mathbb{Z}$ coefficients, the cellular chain group $C_i(X)$ is free with basis the i-cells of X, but there is a sign ambiguity for the basis element corresponding to each cell e^i, namely the choice of a generator for the $\mathbb{Z}$ summand of $H_i(X^i, X^{i-1})$ corresponding to e^i. Only when $i = 0$ is this choice canonical. We refer to these choices as 'choosing orientations for the cells.' A

choice of such orientations allows cellular i-chains to be written unambiguously as linear combinations of i-cells.

Proposition 3.35. *Given choices of orientations for the cells of X and Y, then orientations can be chosen for the cells of $X \times Y$ so that the boundary maps in the cellular chain complexes $C_*(X)$ and $C_*(Y)$ determine the boundary map in $C_*(X \times Y)$ by the formula $d(e^i \times e^j) = de^i \times e^j + (-1)^i e^i \times de^j$.*

Here we are treating the symbol $\times$ algebraically as a bilinear operation, so it could just as well be written as $\otimes$.

Proof: Let us first consider the special case of the n-dimensional cube I^n, the product of n copies of the unit interval I. We give I the CW structure with two vertices and one edge, so the i^{th} copy of I has a 1-cell e_i and 0-cells 1_i and 0_i, with $de_i = 1_i - 0_i$. The n-cell in the product I^n is $e_1 \times \cdots \times e_n$, and we claim that

$$(*) \qquad d(e_1 \times \cdots \times e_n) = \sum_i (-1)^i e_1 \times \cdots \times de_i \times \cdots \times e_n$$

if orientations are chosen suitably. Certainly this formula is correct modulo the signs of the $2n$ cells $e_1 \times \cdots \times 0_i \times \cdots \times e_n$ and $e_1 \times \cdots \times 1_i \times \cdots \times e_n$ since these are exactly the $(n-1)$-cells in the boundary sphere ∂I^n of I^n. To obtain the signs in $(*)$, note that switching the two ends of an I factor of I^n produces a reflection of ∂I^n, as does a transposition of two adjacent I factors. Since reflections have degree -1, this implies that $(*)$ is correct up to an overall sign. This sign can be made $+1$ by orienting the cell $e_1 \times \cdots \times e_n$ correctly, assuming inductively that we have already oriented the $(n-1)$-cell of I^{n-1}.

If we write $I^n = I^i \times I^j$ with $i+j = n$ and set $e^i = e_1 \times \cdots \times e_i$ and $e^j = e_{i+1} \times \cdots \times e_n$, then the formula $(*)$ becomes $d(e^i \times e^j) = de^i \times e^j + (-1)^i e^i \times de^j$. We shall use naturality to reduce the general case of the boundary formula to this special case.

Lemma 3.36. *Suppose $f : X \to Z$ and $g : Y \to W$ are cellular maps, taking n-skeleton to n-skeleton for all n. Then the cellular chain maps $f_* : C_*(X) \to C_*(Z)$, $g_* : C_*(Y) \to C_*(W)$, and $(f \times g)_* : C_*(X \times Y) \to C_*(Z \times W)$ satisfy $(f \times g)_* = f_* \times g_*$.*

Proof: The relation $(f \times g)_* = f_* \times g_*$ means that if $f_*(e^i_\alpha) = \sum_\gamma m_{\alpha\gamma} e^i_\gamma$ and $g_*(e^j_\beta) = \sum_\delta n_{\beta\delta} e^j_\delta$, then $(f \times g)_*(e^i_\alpha \times e^j_\beta) = \sum_{\gamma\delta} m_{\alpha\gamma} n_{\beta\delta} (e^i_\gamma \times e^j_\delta)$. The coefficient $m_{\alpha\gamma}$ is the degree of the composition $f_{\alpha\gamma} : S^i \to X^i / X^{i-1} \to Z^i / Z^{i-1} \to S^i$ where the first and third maps are induced by characteristic maps for the cells e^i_α and e^i_γ, and the middle map is induced by the cellular map f. The $n_{\beta\delta}$'s are obtained similarly from maps $g_{\beta\delta} : S^j \to S^j$.

For $f \times g$, the map $(f \times g)_{\alpha\beta,\gamma\delta} : S^{i+j} \to S^{i+j}$ whose degree is the coefficient of $e^i_\gamma \times e^j_\delta$ in $(f \times g)_*(e^i_\alpha \times e^j_\beta)$ is obtained from the product map $f_{\alpha\gamma} \times g_{\beta\delta} : S^i \times S^j \to S^i \times S^j$ by collapsing the $(i+j-1)$-skeleton of $S^i \times S^j$ to a point. In other words, $(f \times g)_{\alpha\beta,\gamma\delta}$ is the smash product map $f_{\alpha\gamma} \wedge g_{\beta\delta}$. What we need to show is the formula $deg(f \wedge g) = deg(f)\deg(g)$ for basepoint-preserving maps $f : S^i \to S^i$ and $g : S^j \to S^j$.

Since $f \wedge g$ is the composition of $f \wedge \mathbb{1}$ and $\mathbb{1} \wedge g$, it suffices to show $deg(f \wedge \mathbb{1}) = deg(f)$ and $deg(\mathbb{1} \wedge g) = deg(g)$. We do this by relating smash products to suspension. The smash product $X \wedge S^1$ can be viewed as $X \times I/(X \times \partial I \cup \{x_0\} \times I)$, so it is the quotient of the suspension SX obtained by collapsing the segment $\{x_0\} \times I$ to a point. (This quotient is also known as the reduced suspension ΣX.) The quotient map $SX \to X \wedge S^1$ induces an isomorphism on homology since it collapses an interval, which is contractible, to a point. Taking $X = S^i$, we then have a commutative diagram

$$\begin{array}{ccc} S(S^i) & \xrightarrow{Sf} & S(S^i) \\ \downarrow & & \downarrow \\ S^i \wedge S^1 & \xrightarrow{f \wedge \mathbb{1}} & S^i \wedge S^1 \end{array}$$

so Sf and $f \wedge \mathbb{1}$ have the same degree. Since suspension preserves degree by Proposition 2.16, we conclude that $deg(f \wedge \mathbb{1}) = deg(f)$. This is for $\mathbb{1}$ the identity on S^1, and by iteration we obtain the same result for $\mathbb{1}$ the identity on S^j since S^j is the smash product of j copies of S^1. Similarly $deg(\mathbb{1} \wedge g) = deg(g)$ since a permutation of coordinates in S^{i+j} does not affect the degree of maps $S^{i+j} \to S^{i+j}$. □

Now to finish the proof of the proposition, let $\Phi : I^i \to X^i$ and $\Psi : I^j \to Y^j$ be characteristic maps of cells $e^i_\alpha \subset X$ and $e^j_\beta \subset Y$. The restriction of Φ to ∂I^i is the attaching map of e^i_α. We may perform a preliminary homotopy of this attaching map $\partial I^i \to X^{i-1}$ to make it cellular. There is no need to appeal to the cellular approximation theorem to do this since a direct argument is easy, namely, first deform the attaching map so that it sends all but one face of I^i to a point, which is possible since the union of these faces is contractible, then do a further deformation so that the image point of this union of faces is a 0-cell. A homotopy of the attaching map $\partial I^i \to X^{i-1}$ does not affect the cellular boundary de^i_α, since de^i_α is determined by the induced map $H_{i-1}(\partial I^i) \to H_{i-1}(X^{i-1}) \to H_{i-1}(X^{i-1}, X^{i-2})$. So we may assume Φ is cellular, and likewise Ψ, hence also $\Phi \times \Psi$. We remind the reader that the map of cellular chain complexes induced by a cellular map between CW complexes is a chain map, commuting with the cellular boundary maps.

If e^i is the i-cell of I^i and e^j the j-cell of I^j, then $\Phi_*(e^i) = e^i_\alpha$, $\Psi_*(e^j) = e^j_\beta$, and

$(\Phi\times\Psi)_*(e^i\times e^j)=e^i_\alpha\times e^j_\beta$, hence

$$\begin{aligned}
d(e^i_\alpha\times e^j_\beta) &= d\big((\Phi\times\Psi)_*(e^i\times e^j)\big)\\
&= (\Phi\times\Psi)_*d(e^i\times e^j) \qquad \text{since } (\Phi\times\Psi)_* \text{ is a chain map}\\
&= (\Phi\times\Psi)_*(de^i\times e^j+(-1)^i e^i\times de^j) \qquad \text{by the special case}\\
&= \Phi_*(de^i)\times\Psi_*(e^j)+(-1)^i\Phi_*(e^i)\times\Psi_*(de^j) \qquad \text{by the lemma}\\
&= d\Phi_*(e^i)\times\Psi_*(e^j)+(-1)^i\Phi_*(e^i)\times d\Psi_*(e^j) \qquad \text{since } \Phi_* \text{ and } \Psi_* \text{ are chain maps}\\
&= de^i_\alpha\times e^j_\beta+(-1)^i e^i_\alpha\times de^j_\beta.
\end{aligned}$$

which completes the proof of the proposition. □

Example 1. Consider $X\times S^k$ where we give S^k its usual CW structure with two cells. The boundary formula in $C_*(X\times S^k)$ takes the form $d(a\times b)=da\times b$ since $d=0$ in $C_*(S^k)$. So the chain complex $C_*(X\times S^k)$ is just the direct sum of two copies of the chain complex $C_*(X)$, one of the copies having its dimension shifted upward by k. Hence $H_n(X\times S^k;\mathbb{Z})\approx H_n(X;\mathbb{Z})\oplus H_{n-k}(X;\mathbb{Z})$ for all i. In particular, we see that all the homology classes in $X\times S^k$ are obtained from 'products' of homology classes in X and S^k.

Example 2. More subtle things can happen when X and Y both have torsion in their homology. To take the simplest case, let X be S^1 with a cell e^2 attached by a map $S^1\to S^1$ of degree m, so $H_1(X;\mathbb{Z})\approx\mathbb{Z}_m$ and $H_i(X;\mathbb{Z})=0$ for $i>1$. Similarly, let $Y=S^1\cup e^2$ with an attaching map of degree n. Since X and Y have no homology in dimension greater than 1, one might guess that $X\times Y$ would have no homology in dimension greater than 2. To see whether this is true, we examine the cellular chain complex for $X\times Y$ in dimensions 2, 3, and 4. This is $\mathbb{Z}\to\mathbb{Z}\oplus\mathbb{Z}\to\mathbb{Z}$ with generators α, β, γ, δ corresponding to the four cells $e^2\times e^2$, $e^1\times e^2$, $e^2\times e^1$, and $e^1\times e^1$, boundary maps being given by $d\alpha=m\beta+n\gamma$, $d\beta=-n\delta$, $d\gamma=m\delta$, and $d\delta=0$. Since $d(k\beta+\gamma)=(m-kn)\delta$, the homology class of δ has order $q=gcd(m,n)$, the greatest common divisor of m and n, and $H_2(X\times Y;\mathbb{Z})\approx\mathbb{Z}_q$, generated by the 'product' of the generators of $H_1(X;\mathbb{Z})$ and $H_1(Y;\mathbb{Z})$. This is what one might expect. However, there are also cellular cycles in dimension 3, the multiples of $(m/q)\beta+(n/q)\gamma$, and the smallest such multiple which is a boundary is $q[(m/q)\beta+(n/q)\gamma]=m\beta+n\gamma$, so $H_3(X\times Y;\mathbb{Z})\approx\mathbb{Z}_q$.

Thus if $q>1$, not all the homology of $X\times Y$ comes from products of homology classes in X and Y. This phenomenon can be traced to $H_1(X;\mathbb{Z})$ and $H_1(Y;\mathbb{Z})$ having elements of the same finite order q, namely the classes represented by $(m/q)e^1$ for X and $(n/q)e^1$

for Y, which give rise to the 3-cycle $(m/q)e^1 \times e^2 + e^2 \times (n/q)e^1 = (m/q)\beta + (n/q)\gamma$. As we shall see, it is the existence of torsion elements of the same order in the homology of X and Y which prevents the homology of $X \times Y$ from being just the product of the homology of X and Y in general.

This example generalizes immediately to higher dimensions, with $X = S^i \cup e^{i+1}$ and $Y = S^j \cup e^{j+1}$, the attaching maps having degrees m and n, respectively. Essentially the same calculation shows that $X \times Y$ has both H_{i+j} and H_{i+j+1} isomorphic to $\mathbb{Z}_q$ where $q = gcd(m, n)$.

The basic question now is whether the map

$$\textstyle\bigoplus_i \big(H_i(X;\mathbb{Z}) \otimes H_{n-i}(Y;\mathbb{Z})\big) \longrightarrow H_n(X \times Y;\mathbb{Z})$$

obtained by adding the various cross products is an isomorphism. Example 2 shows that this is not always the case, though it is true in Example 1. Our main goal in what follows is to show that the map is always injective, with cokernel $\bigoplus_i \operatorname{Tor}\big(H_i(X;\mathbb{Z}), H_{n-i-1}(Y;\mathbb{Z})\big)$. More generally, we consider other coefficients besides $\mathbb{Z}$ and show in particular that with field coefficients it is an isomorphism.

The Algebraic Künneth Formula

For CW complexes X and Y, the relationship between the cellular chain complexes $C_*(X)$, $C_*(Y)$, and $C_*(X \times Y)$ can be expressed nicely in terms of tensor products. Since the n-cells of $X \times Y$ are the products of i-cells of X with $(n-i)$-cells of Y, we have $C_n(X \times Y) \approx \bigoplus_i (C_i(X) \otimes C_{n-i}(Y))$, with $e^i \times e^j$ corresponding to $e^i \otimes e^j$, and the boundary formula of Proposition 3.35 becomes $d(e^i \otimes e^j) = de^i \otimes e^j + (-1)^i e^i \otimes de^j$. Our task now is purely algebraic, to compute the homology of the chain complex $C_*(X \times Y)$ from the homology of $C_*(X)$ and $C_*(Y)$.

Suppose we are given chain complexes $C = \{C_n\}$ and $C' = \{C'_n\}$ of abelian groups, or more generally R-modules over a commutative ring R. The **tensor product chain complex** $C \otimes_R C'$ is then defined by $(C \otimes_R C')_n = \bigoplus_i (C_i \otimes_R C'_{n-i})$, with boundary maps given by $\partial(c \otimes c') = \partial c \otimes c' + (-1)^i c \otimes \partial c'$ for $c \in C_i$ and $c' \in C'_{n-i}$ The sign $(-1)^i$ guarantees that $\partial^2 = 0$ in $C \otimes_R C'$. From the boundary formula $\partial(c \otimes c') = \partial c \otimes c' + (-1)^i c \otimes \partial c'$ it follows that the tensor product of cycles is a cycle, and the tensor product of a cycle and a boundary, in either order, is a boundary, just as for the cross product defined earlier. So there is induced a natural map on homology groups, $h : H_i(C) \otimes_R H_{n-i}(C') \to H_n(C \otimes_R C')$. Summing over i then gives $h : \bigoplus_i \big(H_i(C) \otimes_R H_{n-i}(C')\big) \to H_n(C \otimes_R C')$. This figures in the following algebraic version of the Künneth Formula:

Theorem 3.37. *If R is a principal ideal domain and the R-modules C_i are free, then for each n there is a natural short exact sequence*

$$0 \to \bigoplus\nolimits_i \big(H_i(C) \otimes_R H_{n-i}(C')\big) \to H_n(C \otimes_R C') \to \bigoplus\nolimits_i \big(\mathrm{Tor}^R(H_i(C), H_{n-i-1}(C')\big) \to 0$$

which splits if the groups C_i' are free, or more generally if the cycles Z_i' are direct summands of C_i'.

This is a generalization of the universal coefficient theorem for homology, which is the case that C' consists of just the coefficient group G in dimension zero. The proof will also be a natural generalization of the proof of the universal coefficient theorem.

Proof: First we do the special case that the boundary maps in C are all zero, so $H_i(C) = C_i$. In this case $\partial(c \otimes c') = (-1)^i c \otimes \partial c'$ and the chain complex $C \otimes_R C'$ is simply the direct sum of the complexes $C_i \otimes_R C'$, each of which is a direct sum of copies of C' since C_i is free. Hence $H_n(C_i \otimes_R C') \approx C_i \otimes_R H_{n-i}(C') = H_i(C) \otimes_R H_{n-i}(C')$. Summing over i yields an isomorphism $H_n(C \otimes_R C') \approx \bigoplus_i \big(H_i(C) \otimes_R H_{n-i}(C')\big)$, which is the statement of the theorem since there are no Tor terms, $H_i(C) = C_i$ being free.

In the general case, let $Z_i \subset C_i$ and $B_i \subset C_i$ denote kernel and image of the boundary homomorphisms for C. We can view $Z = \{Z_i\}$ and $B = \{B_i\}$ as subchain complexes of C having trivial boundary maps. Then we have a short exact sequence of chain complexes $0 \to Z \to C \to B \to 0$ made up of the short exact sequences $0 \to Z_i \to C_i \xrightarrow{\partial} B_{i-1} \to 0$ each of which splits since B_{i-1} is free, being a submodule of C_{i-1} which is free by assumption. Because of the splitting, tensoring $0 \to Z \to C \to B \to 0$ with C' produces another short exact sequence of chain complexes, and hence a long exact sequence in homology

$$\cdots \longrightarrow H_n(Z \otimes_R C') \longrightarrow H_n(C \otimes_R C') \longrightarrow H_n(B \otimes_R C') \longrightarrow H_{n-1}(Z \otimes_R C') \longrightarrow \cdots$$

where we have $H_{n-1}(B \otimes_R C')$ instead of the expected $H_n(B \otimes_R C')$ since $\partial : C \to B$ decreases dimension by one. Checking definitions, one sees that the 'boundary' map $H_{n-1}(B \otimes_R C') \to H_{n-1}(Z \otimes_R C')$ in the preceding long exact sequence is just the map induced by the natural map $B \otimes_R C' \to Z \otimes_R C'$ coming from the inclusion $B \subset Z$.

Since Z and B are chain complexes with trivial boundary maps, the special case at the beginning of the proof converts the preceding exact sequence into

$$\cdots \xrightarrow{\iota_n} \bigoplus\nolimits_i \big(Z_i \otimes_R H_{n-i}(C')\big) \longrightarrow H_n(C \otimes_R C') \longrightarrow \bigoplus\nolimits_i \big(B_i \otimes_R H_{n-i-1}(C')\big) \xrightarrow{\iota_{n-1}} \bigoplus\nolimits_i \big(Z_i \otimes_R H_{n-i-1}(C')\big) \longrightarrow \cdots$$

So we have short exact sequences

$$0 \longrightarrow \mathit{Coker}(\iota_n) \longrightarrow H_n(C \otimes_R C') \longrightarrow \mathit{Ker}(\iota_{n-1}) \longrightarrow 0$$

where $\mathit{Coker}(\iota_n) = \bigoplus_i \big(Z_i \otimes_R H_{n-i}(C')\big)/\mathit{Im}(\iota_n) = \bigoplus_i \big(H_i(C) \otimes_R H_{n-i}(C')\big)$ by Lemma 3.28. It remains to identify $\mathit{Ker}(\iota_{n-1})$ with $\bigoplus_i \operatorname{Tor}^R\big(H_i(C), H_{n-i}(C')\big)$.

By the definition of Tor, tensoring the free resolution $0 \to B_i \to Z_i \to H_i(C) \to 0$ with $H_{n-i}(C')$ yields an exact sequence

$$0 \longrightarrow \operatorname{Tor}^R\big(H_i(C), H_{n-i}(C')\big) \longrightarrow B_i \otimes_R H_{n-i}(C') \longrightarrow Z_i \otimes_R H_{n-i}(C') \longrightarrow H_i(C) \otimes_R H_{n-i}(C') \longrightarrow 0$$

Hence, summing over i, $\mathit{Ker}(\iota_n) = \bigoplus_i \operatorname{Tor}^R\big(H_i(C), H_{n-i}(C')\big)$.

Naturality should be obvious, and we leave it for the reader to fill in the details.

To show the short exact sequence in the statement of the theorem splits if C' is free, it suffices to find a homomorphism $k : H_n(C \otimes_R C') \to \bigoplus_i \big(H_i(C) \otimes_R H_{n-i}(C')\big)$ with $kh = \mathbb{1}$. As already noted, the sequence $0 \to Z_i \to C_i \to B_{i-1} \to 0$ splits, so the quotient maps $Z_i \to H_i(C)$ extend to homomorphisms $C_i \to H_i(C)$. Similarly we obtain $C'_j \to H_j(C')$ if C'_{j-1} is free, or if Z'_i is a direct summand of C'_i. Viewing $\{H_i(C)\}$ and $\{H_j(C')\}$ as chain complexes $H(C)$ and $H(C')$ with trivial boundary maps, we thus have chain maps $C \to H(C)$ and $C' \to H(C')$. Taking tensor products gives a chain map $C \otimes_R C' \to H(C) \otimes_R H(C')$, hence an induced map on homology, which is the desired map k since the chain complex $H(C) \otimes_R H(C')$ equals its own homology, the boundary maps being trivial. □

The Topological Künneth Formula

Now we can apply the preceding algebra to obtain the topological statement we are looking for:

Theorem 3.38. *If X and Y are CW complexes and R is a principal ideal domain, then there are natural short exact sequences*

$$0 \longrightarrow \bigoplus_i \big(H_i(X;R) \otimes_R H_{n-i}(Y;R)\big) \longrightarrow H_n(X \times Y; R) \longrightarrow \bigoplus_i \operatorname{Tor}^R\big(H_i(X;R), H_{n-i-1}(Y;R)\big) \longrightarrow 0$$

and these sequences split.

Naturality means that maps $X \to X'$ and $Y \to Y'$ induce a map from the short exact sequence for $X \times Y$ to the corresponding short exact sequence for $X' \times Y'$, with commuting squares. The splitting is not natural, however.

Proof: When dealing with products of CW complexes there is always the bothersome fact that the compactly generated CW topology may not be the same as the product topology. However, in the present context this is not a real problem. Since the two topologies have the same compact sets, they have the same singular simplices, hence the same singular homology groups.

Let $C = C_*(X;R)$ and $C' = C_*(Y;R)$, the cellular chain complexes with coefficients in R, so $C \otimes_R C' = C_*(X;R) \otimes C_*(Y;R)$. By Proposition 3.35 this equals $C_*(X \times Y;R)$, so the algebraic Künneth formula gives the result. □

With field coefficients the Künneth formula simplifies because the Tor terms are always zero over a field:

Corollary 3.39. *If F is a field and X and Y are CW complexes, then the cross product map $h : \bigoplus_i \big(H_i(X;F) \otimes_F H_{n-i}(Y;F)\big) \longrightarrow H_n(X \times Y;F)$ is an isomorphism for all n.* □

There is also a relative version of the Künneth formula for CW pairs (X, A) and (Y, B). This is a split short exact sequence

$$0 \longrightarrow \bigoplus_i \big(H_i(X, A;R) \otimes_R H_{n-i}(Y, B;R)\big) \longrightarrow H_n(X \times Y, A \times Y \cup X \times B;R) \longrightarrow \bigoplus_i \mathrm{Tor}^R\big(H_i(X, A;R), H_{n-i-1}(Y, B;R)\big) \longrightarrow 0$$

for R a principal ideal domain. This too follows from the algebraic Künneth formula since the isomorphism of cellular chain complexes $C_*(X \times Y) \approx C_*(X) \otimes C_*(Y)$ passes down to a quotient isomorphism

$$C_*(X \times Y)/C_*(A \times Y \cup X \times B) \approx C_*(X)/C_*(A) \otimes C_*(Y)/C_*(B)$$

since bases for these three relative cellular chain complexes correspond bijectively with the cells of $(X - A) \times (Y - B)$, $X - A$, and $Y - B$, respectively.

As a special case, suppose A and B are basepoints $x_0 \in X$ and $y_0 \in Y$. Then the subcomplex $A \times Y \cup X \times B$ can be identified with the wedge sum $X \vee Y$ and the quotient $X \times Y / X \vee Y$ is the smash product $X \wedge Y$. Thus we have a reduced Künneth formula

$$0 \longrightarrow \bigoplus_i \big(\widetilde{H}_i(X;R) \otimes_R \widetilde{H}_{n-i}(Y;R)\big) \longrightarrow \widetilde{H}_n(X \wedge Y;R) \longrightarrow \bigoplus_i \mathrm{Tor}^R\big(\widetilde{H}_i(X;R), \widetilde{H}_{n-i-1}(Y;R)\big) \longrightarrow 0$$

For example, when $Y = S^k$, $X \wedge S^k$ is the k-fold reduced suspension of X, and we obtain isomorphisms $\widetilde{H}_n(X;\mathbb{Z}) \approx \widetilde{H}_{n+k}(X \wedge S^k;\mathbb{Z})$. More generally, by taking Y to be a Moore space $M(G,k)$ and then applying the universal coefficient theorem we obtain:

Corollary 3.40. *There are natural isomorphisms* $\widetilde{H}_n(X;G) \approx \widetilde{H}_{n+k}(X \wedge M(G,k);\mathbb{Z})$ *for all CW complexes* X *and abelian groups* G. □

This says that homology with arbitrary coefficients is obtainable from homology with $\mathbb{Z}$ coefficients by a geometric construction as well as by the algebra of tensor products. For general homology theories this formula is used as a definition of homology with coefficients.

The Künneth formula and the universal coefficient theorem can be combined to give a more concise formula $H_n(X \times Y;G) \approx \bigoplus_i H_i(X;H_{n-i}(Y;G))$, at least when $G = \mathbb{Z}$. In fact, with a little more algebra one can show that this formula is valid for arbitrary coefficient groups G; see [Hilton-Wylie p.227] or [Spanier, p. 235]. However the naturality of this isomorphism is problematical since it uses the splittings in the Künneth formulas and universal coefficient theorems.

One might wonder about a cohomology version of the Künneth formula. For example, taking coefficients in a field F and using the natural isomorphism $Hom(A\otimes B,C) \approx Hom(A,Hom(B,C))$, the Künneth formula for homology and the universal coefficient theorem give isomorphisms

$$\begin{aligned} H^n(X\times Y;F) &\approx Hom_F(H_n(X\times Y;F),F) \approx \textstyle\bigoplus_i Hom_F(H_i(X;F)\otimes H_{n-i}(Y;F),F) \\ &\approx \textstyle\bigoplus_i Hom_F\big(H_i(X;F), Hom_F(H_{n-i}(Y;F),F)\big) \\ &\approx \textstyle\bigoplus_i Hom_F\big(H_i(X;F), H^{n-i}(Y;F)\big) \\ &\approx \textstyle\bigoplus_i H^i(X;H^{n-i}(Y;F)) \end{aligned}$$

More generally, there are isomorphisms $H^n(X\times Y;G) \approx \bigoplus_i H^i(X;H^{n-i}(Y;G))$ for any coefficient group G; see [Hilton-Wylie p.227]. However, in practice it usually suffices to apply the Künneth formula for homology and the universal coefficient theorem for cohomology separately. Also, Theorem 3.9 shows that with stronger hypotheses one can draw stronger conclusions using cup products.

The Simplicial Cross Product

It may be of interest to see how the cross product $H_m(X;R)\otimes H_n(Y;R) \to H_{m+n}(X\times Y;R)$ can be defined directly in terms of singular homology. What one wants is a cross product at the level of singular chains, $C_m(X;R)\otimes C_n(Y;R) \to C_{m+n}(X\times Y;R)$. If we are given singular simplices $f:\Delta^m \to X$ and $g:\Delta^n \to Y$, then we have the product map $f\times g:\Delta^m\times\Delta^n \to X\times Y$, and the idea is to subdivide $\Delta^m\times\Delta^n$ into simplices of dimension $m+n$ and then take the sum of the restrictions of $f\times g$ to these simplices, with appropriate signs.

In the special cases that m or n is 1 we have already seen how to subdivide $\Delta^m \times \Delta^n$ into simplices when we constructed prism operators in §2.1. The generalization to $\Delta^m \times \Delta^n$ is not completely obvious, however. Label the vertices of Δ^m as $v_0, v_1, \cdots, v_m$ and the vertices of Δ^n as $w_0, w_1, \cdots, w_n$. Think of the pairs (i, j) with $0 \leq i \leq m$ and $0 \leq j \leq n$ as the vertices of an $m \times n$ rectangular grid in $\mathbb{R}^2$. Let σ be a path formed by a sequence of $m+n$ horizontal and vertical edges in this grid starting at $(0, 0)$ and ending at (m, n), always moving either to the right or upward. To such a path σ we associate a linear map $\ell_\sigma : \Delta^{m+n} \to \Delta^m \times \Delta^n$ sending the k^{th} vertex of Δ^{m+n} to (v_{i_k}, w_{j_k}) where (i_k, j_k) is the k^{th} vertex of the edgepath σ. Then we define a simplicial cross product

$$C_m(X; R) \otimes C_n(Y; R) \xrightarrow{\times} C_{m+n}(X \times Y; R)$$

by the formula

$$f \times g = \sum_\sigma (-1)^{|\sigma|} (f \times g) \ell_\sigma$$

where $|\sigma|$ is the number of squares in the grid lying below the path σ. Note that the symbol '$\times$' means different things on the two sides of the equation. From this definition it is a calculation to show that $\partial(f \times g) = \partial f \times g + (-1)^m f \times \partial g$. This implies that the cross product of two cycles is a cycle, and the cross product of a cycle and a boundary is a boundary, so there is an induced cross product in singular homology.

One can see that the images of the maps ℓ_σ give a simplicial structure on $\Delta^m \times \Delta^n$ in the following way. We can view Δ^m as the subspace of $\mathbb{R}^m$ defined by the inequalities $0 \leq x_1 \leq \cdots \leq x_m \leq 1$, with the vertex v_i the point having coordinates i zeroes followed by $m - i$ ones. Similarly we have $\Delta^n \subset \mathbb{R}^n$ with coordinates $0 \leq y_1 \leq \cdots \leq y_n \leq 1$. The product $\Delta^m \times \Delta^n$ then consists of $(m+n)$-tuples $(x_1, \cdots, x_m, y_1, \cdots, y_n)$ satisfying both sets of inequalities. The inequalities $0 \leq x_1 \leq \cdots \leq x_m \leq y_1 \leq \cdots \leq y_n \leq 1$ define a simplex Δ^{m+n} in $\Delta^m \times \Delta^n$, and every other point of $\Delta^m \times \Delta^n$ satisfies a similar set of inequalities obtained from $0 \leq x_1 \leq \cdots \leq x_m \leq y_1 \leq \cdots \leq y_n \leq 1$ by a permutation of the variables 'shuffling' the y_j's into the x_i's. Each such shuffle corresponds to an edgepath σ consisting of a rightward edge for each x_i and an upward edge for each y_j in the shuffled sequence. Thus we have $\Delta^m \times \Delta^n$ expressed as the union of simplices Δ^{m+n}_σ indexed by the edgepaths σ. One can check that these simplices fit together nicely to form a Δ-complex structure on $\Delta^m \times \Delta^n$, which is also a simplicial complex structure. See [Eilenberg-Steenrod, p.68] or [Gelfand-Manin, p.5]. In fact this construction is sufficiently natural to make the product of any two Δ-complexes into a Δ-complex.

Exercises

1. Compute the groups $H_n(\mathbb{R}P^2 \times \mathbb{R}P^2; G)$ and $H^n(\mathbb{R}P^2 \times \mathbb{R}P^2; G)$ for $G = \mathbb{Z}$ and $\mathbb{Z}_2$

via the cellular chain and cochain complexes.

2. Show by an example that the splitting in the Künneth formula cannot be natural.

3. Show that **slant products**

$$H_n(X \times Y; R) \times H^j(Y; R) \longrightarrow H_{n-j}(Y; R), \quad (e^i \times e^j, \varphi) \mapsto \varphi(e^j)e^i$$
$$H^n(X \times Y; R) \times H_j(Y; R) \longrightarrow H^{n-j}(Y; R), \quad (\varphi, e^j) \mapsto \big(e^i \mapsto \varphi(e^i \times e^j)\big)$$

can be defined via the indicated cellular formulas. [These 'products' are in some ways more like division than multiplication, and this is reflected in the common notation a/b for them, or $a\backslash b$ when the order of the factors is reversed. The first of the two slant products is related to cap product as the cohomology cross product is related to cup product.]

c. H-Spaces and Hopf Algebras

A space X is an **H-space**, 'H' standing for 'Hopf,' if there is a continuous multiplication map $\mu: X \times X \to X$ and an 'identity' element $e \in X$ such that the two maps $x \mapsto \mu(x, e)$ and $x \mapsto \mu(e, x)$ are homotopic to the identity through maps $(X, e) \to (X, e)$. In particular, this implies $\mu(e, e) = e$. For example, the unit spheres S^1 in $\mathbb{C}$, S^3 in the quaternions $\mathbb{H}$, and S^7 in the octonions $\mathbb{O}$ are H-spaces since the multiplications in these division algebras are continuous, being defined by polynomial formulas, and are norm-preserving, $|ab| = |a||b|$, hence restrict to multiplications on the unit spheres, with identity elements. Both S^1 and S^3 are actually groups since the multiplications in $\mathbb{C}$ and $\mathbb{H}$ are associative and inverses exist since $a\overline{a} = |a|^2 = 1$ if $|a| = 1$. Most of the classical examples of H-spaces are in fact **topological groups**, that is, topological spaces X with a group structure such that the multiplication $X \times X \to X$ is continuous, as is also the inversion map $x \mapsto x^{-1}$.

As another example, multiplication of polynomials provides an H-space structure on $\mathbb{C}\mathrm{P}^\infty$ in the following way. A nonzero polynomial $a_0 + a_1 z + \cdots + a_n z^n$ with coefficients $a_i \in \mathbb{C}$ corresponds to a point $(a_0, \cdots, a_n, 0, \cdots) \in \mathbb{C}^\infty - \{0\}$. Multiplication of two such polynomials determines a multiplication $\mathbb{C}^\infty - \{0\} \times \mathbb{C}^\infty - \{0\} \to \mathbb{C}^\infty - \{0\}$ which is associative, commutative, and has an identity element $(1, 0. \cdots)$. Since $\mathbb{C}$ is commutative we can factor out by scalar multiplication by nonzero constants and get an induced product $\mathbb{C}\mathrm{P}^\infty \times \mathbb{C}\mathrm{P}^\infty \to \mathbb{C}\mathrm{P}^\infty$ with the same properties. Thus $\mathbb{C}\mathrm{P}^\infty$ is an associative, commutative H-space with a strict identity. Instead of factoring out by all nonzero scalars, we could factor out only by scalars of the form $\rho e^{2\pi i k/q}$ for a fixed positive integer q, with arbitrary $\rho > 0$ and $k \in \mathbb{Z}$. The quotient of $\mathbb{C}^\infty - \{0\}$ under this identification, an infinite-dimensional lens space L^∞ with $\pi_1(L^\infty) \approx \mathbb{Z}_q$, is therefore also an associative, commutative H-space. This includes $\mathbb{R}\mathrm{P}^\infty$ in particular.

The spaces $J(X)$ defined in §3.2 are also H-spaces, with the multiplication given by $(x_1, \cdots, x_m)(y_1, \cdots, y_n) = (x_1, \cdots, x_m, y_1, \cdots, y_n)$, which is associative and has an identity element (e) where e is the basepoint of X. One could describe $J(X)$ as 'the free associative H-space generated by X.' There is also a commutative analog of $J(X)$ called the **infinite symmetric product** $SP(X)$ defined in the following way. Let $SP_n(X)$ be the quotient space of the n-fold product X^n obtained by identifying all n-tuples $(x_1, \cdots, x_n)$ which differ only by a permutation of their coordinates. The inclusion $X^n \hookrightarrow X^{n+1}$, $(x_1, \cdots, x_n) \mapsto (x_1, \cdots, x_n, e)$ induces an inclusion $SP_n(X) \hookrightarrow SP_{n+1}$, and $SP(X)$ is defined to be the union of this increasing sequence of $SP_n(X)$'s, with the weak topology. Alternatively, $SP(X)$ is the quotient of $J(X)$ obtained by identifying points which differ

only by permutation of coordinates. The H-space structure on $J(X)$ induces an H-space structure on $SP(X)$ which is commutative in addition to being associative and having a strict identity.

An interesting special case is when $X = S^2$ because in this case $SP(S^2)$ can be identified with $\mathbb{C}\mathrm{P}^\infty$ in the following way. We first identify $\mathbb{C}\mathrm{P}^n$ with the nonzero polynomials of degree at most n with coefficients in $\mathbb{C}$, modulo scalar multiplication, as before. The sphere S^2 we view as $\mathbb{C} \cup \{\infty\}$, and then we define $f : (S^2)^n \to \mathbb{C}\mathrm{P}^n$ by setting $f(a_1, \cdots, a_n) = (z + a_1) \cdots (z + a_n)$ with factors $z + \infty$ omitted, so in particular $f(\infty, \cdots, \infty) = 1$. To check that f is continuous, suppose some a_i approaches ∞, say a_n, and all the other a_j's are finite. Then if we write

$$(z + a_1) \cdots (z + a_n) = z^n + (a_1 + \cdots + a_n)z^{n-1} + \cdots + \sum_{i_1 < \cdots < i_k} a_{i_1} \cdots a_{i_k} z^{n-k} + \cdots + a_1 \cdots a_n$$

we see that, dividing through by a_n and letting a_n approach ∞, this polynomial approaches $z^{n-1} + (a_1 + \cdots + a_{n-1})z^{n-2} + \cdots + a_1 \cdots a_{n-1} = (z + a_1) \cdots (z + a_{n-1})$. The same argument would apply if several a_i's approach ∞ simultaneously.

The value $f(a_1, \cdots, a_n)$ is unchanged under permutation of the a_i's, so there is an induced map $SP_n(S^2) \to \mathbb{C}\mathrm{P}^n$ which is a continuous bijection, hence a homeomorphism since both spaces are compact Hausdorff. Letting n go to ∞, we then get a homeomorphism $SP(S^2) \approx \mathbb{C}\mathrm{P}^\infty$.

The spaces $SP(X)$ are studied in more detail in §4.4j.

These examples illustrate the general principle that many of the most important spaces in algebraic topology turn out to be H-spaces.

Hopf Algebras

The homology and cohomology of H-spaces have extra structure arising from the multiplication. In particular, homology has a product operation formally similar to cup product in cohomology, and in many cases the resulting homology ring has a simpler structure than the cohomology ring. For example, we will see that $H_*(J(S^n); \mathbb{Z})$ is just a polynomial ring $\mathbb{Z}[x]$.

Let us look at cohomology first. Choosing a commutative ring R as coefficient ring, we can regard the cohomology ring $H^*(X; R)$ of a space X as an algebra over R rather than merely a ring. Suppose X is an H-space satisfying two conditions:

(1) X is path-connected, hence $H^0(X; R) \approx R$.

(2) $H^n(X; R)$ is a finitely generated free R-module for each n, so the external cup product $H^*(X; R) \otimes H^*(X; R) \to H^*(X \times X; R)$ is an isomorphism.

Let $\mu: X\times X \to X$ be the multiplication map, $e \in X$ the identity element, and $i_1, i_2 : X \to X \times X$ the two inclusions $x \mapsto (x, e)$ and $x \mapsto (e, x)$. The commutative diagram

$$\begin{array}{ccccc} H^*(X;R) & \xrightarrow{\mu^*} & H^*(X\times X;R) & \xrightarrow{i_j^*} & H^*(X;R) \\ & \searrow^{\Delta} & \uparrow\approx & \nearrow_{P_j} & \\ & & H^*(X;R)\otimes H^*(X;R) & & \end{array}$$

defines maps Δ and P_j which are algebra homomorphisms since μ^* and i_j^* are algebra homomorphisms, being induced by maps of spaces. The map P_j is projection onto the j^{th} factor, i.e., $P_1(\alpha\otimes 1) = \alpha$ and $P_1(\alpha\otimes\beta) = 0$ if β has positive dimension, and similarly for P_2. The H-space property says that $\mu i_j \simeq \mathbb{1}$, so $P_j\Delta = \mathbb{1}$. This means that $\Delta(\alpha)$ for $\alpha \in H^n(X;R)$, $n > 0$, has the special form $\Delta(\alpha) = \alpha\otimes 1 + 1\otimes\alpha + \sum_{0<i<n} \alpha_i' \otimes \alpha_{n-i}''$ where $\alpha_j', \alpha_j'' \in H^j(X;R)$.

We introduce some terminology for this algebraic situation. A graded algebra $A = \bigoplus_n A^n$ over a commutative base ring R is said to be **connected** if $A^0 \approx R$. A connected graded algebra A with an identity element $1 \in A^0$ is a **Hopf algebra** if it has a **diagonal** or **coproduct** $\Delta: A \to A\otimes A$ which is a homomorphism of graded algebras satisfying $\Delta(\alpha) = \alpha\otimes 1 + 1\otimes\alpha + \sum_{0<i<n} \alpha_i'\otimes\alpha_{n-i}''$ for $\alpha \in A^n$, $n > 0$, and $\alpha_j', \alpha_j'' \in A^j$. Here $\otimes$ means $\otimes_R$, and the multiplication in $A\otimes A$ is given by $(\alpha\otimes\beta)(\gamma\otimes\delta) = (-1)^{jk}(\alpha\gamma\otimes\beta\delta)$ where $\beta \in A^j$ and $\gamma \in A^k$. The multiplication in A is not assumed to be associative or commutative, where commutativity has its usual graded meaning: $\alpha\beta = (-1)^{ij}\beta\alpha$ for $\alpha \in A^i$ and $\beta \in A^j$.

Example 1. One of the simplest Hopf algebras is a polynomial ring $R[\alpha]$. The coproduct $\Delta(\alpha)$ must equal $\alpha\otimes 1 + 1\otimes\alpha$ since the only elements of $R[\alpha]$ of lower dimension than α are the elements of R in dimension zero, so the terms α_i' and α_{n-i}'' in the coproduct formula $\Delta(\alpha) = \alpha\otimes 1 + 1\otimes\alpha + \sum_{0<i<n} \alpha_i'\otimes\alpha_{n-i}''$ must be zero. The requirement that Δ be an algebra homomorphism then determines Δ completely. To describe Δ explicitly we distinguish two cases. If the dimension of α is even (or if $2 = 0$ in R), then the multiplication in $R[\alpha]\otimes R[\alpha]$ is strictly commutative and $\Delta(\alpha^n) = (\alpha\otimes 1 + 1\otimes\alpha)^n = \sum_i \binom{n}{i}\alpha^i\otimes\alpha^{n-i}$. In the opposite case that α is odd-dimensional, then $\Delta(\alpha^2) = (\alpha\otimes 1 + 1\otimes\alpha)^2 = \alpha^2\otimes 1 + 1\otimes\alpha^2$ since $(\alpha\otimes 1)(1\otimes\alpha) = \alpha\otimes\alpha$ and $(1\otimes\alpha)(\alpha\otimes 1) = -\alpha\otimes\alpha$ if α has odd dimension. Thus if we set $\beta = \alpha^2$, then β is even-dimensional and we have $\Delta(\alpha^{2n}) = \Delta(\beta^n) = (\beta\otimes 1 + 1\otimes\beta)^n = \sum_i \binom{n}{i}\beta^i\otimes\beta^{n-i}$ and $\Delta(\alpha^{2n+1}) = \Delta(\alpha\beta^n) = \Delta(\alpha)\Delta(\beta^n) = \sum_i \binom{n}{i}\alpha\beta^i\otimes\beta^{n-i} + \sum_i \binom{n}{i}\beta^i\otimes\alpha\beta^{n-i}$.

Example 2. The exterior algebra $\Lambda_R[\alpha]$ on an odd-dimensional generator α is a Hopf algebra, with $\Delta(\alpha) = \alpha\otimes 1 + 1\otimes\alpha$. To verify that Δ is an algebra homomorphism we must

check that $\Delta(\alpha^2) = \Delta(\alpha)^2$, or in other words, since $\alpha^2 = 0$, we need to see that $\Delta(\alpha)^2 = 0$. As in the preceding example we have $\Delta(\alpha)^2 = (\alpha\otimes 1 + 1\otimes\alpha)^2 = \alpha^2\otimes 1 + 1\otimes\alpha^2$, so $\Delta(\alpha)^2$ is indeed 0. Note that if α were even-dimensional we would instead have $\Delta(\alpha)^2 = \alpha^2\otimes 1 + 2\alpha\otimes\alpha + 1\otimes\alpha^2$, which would be 0 in $\Lambda_R[\alpha]\otimes\Lambda_R[\alpha]$ only if $2 = 0$ in R.

An element α of a Hopf algebra is called **primitive** if $\Delta(\alpha) = \alpha\otimes 1 + 1\otimes\alpha$. As the preceding examples illustrate, if a Hopf algebra is generated as an algebra by primitive elements, then the coproduct Δ is uniquely determined by the multiplication. This happens in a number of special cases of interest, but certainly not in general.

The existence of the coproduct in a Hopf algebra turns out to restrict the multiplicative structure considerably. Here is an important example illustrating this:

Example 3. Suppose that the truncated polynomial algebra $F[\alpha]/(\alpha^n)$ over a field F is a Hopf algebra. Then α is primitive, just as it is in $F[\alpha]$, so if we assume either that α is even-dimensional or F has characteristic 2, then the relation $\alpha^n = 0$ yields an equation

$$0 = \Delta(\alpha^n) = \alpha^n\otimes 1 + 1\otimes\alpha^n + \sum_{0<i<n}\tbinom{n}{i}\alpha^i\otimes\alpha^{n-i} = \sum_{0<i<n}\tbinom{n}{i}\alpha^i\otimes\alpha^{n-i}$$

which implies $\binom{n}{i} = 0$ in F for each i in the range $0 < i < n$. This is impossible if F has characteristic 0, and if the characteristic of F is $p > 0$ then the next lemma implies that n must be a power of p. Conversely, it is easy to check that $F[\alpha]/(\alpha^{p^i})$ is a Hopf algebra, with α even-dimensional if p is odd.

Lemma 3.41. *If p is a prime, then $\binom{n}{k} \equiv \prod_i \binom{n_i}{k_i} \mod p$ where $n = \sum_i n_i p^i$ and $k = \sum_i k_i p^i$ are the p-adic representations of n and k, with $0 \le n_i < p$ and $0 \le k_i < p$. Hence $\binom{n}{k} \equiv 0 \mod p$ for all k in the interval $0 < k < n$ iff $n = p^\ell$ for some ℓ.*

Here the convention is that $\binom{n}{k} = 0$ if $n < k$, and $\binom{n}{0} = 1$ for all $n \ge 0$.

Proof: In $\mathbb{Z}_p[x]$ there is an identity $(1+x)^p = 1 + x^p$ since p clearly divides $\binom{p}{k} = p!/k!(p-k)!$ for $0 < k < p$. By induction it follows that $(1+x)^{p^i} = 1 + x^{p^i}$. Hence if $n = \sum_i n_i p^i$ is the p-adic representation of n then:

$$\begin{aligned}(1+x)^n &= (1+x)^{n_0}(1+x^p)^{n_1}(1+x^{p^2})^{n_2}\cdots\\ &= \Big[1 + \binom{n_0}{1}x + \binom{n_0}{2}x^2 + \cdots + \binom{n_0}{p-1}x^{p-1}\Big]\\ &\quad\times\Big[1 + \binom{n_1}{1}x^p + \binom{n_1}{2}x^{2p} + \cdots + \binom{n_1}{p-1}x^{(p-1)p}\Big]\\ &\quad\quad\times\Big[1 + \binom{n_2}{1}x^{p^2} + \binom{n_2}{2}x^{2p^2} + \cdots + \binom{n_2}{p-1}x^{(p-1)p^2}\Big]\times\cdots\end{aligned}$$

When this is multiplied out, one sees that no terms combine, and the coefficient of x^k is just $\prod_i \binom{n_i}{k_i}$ where $k = \sum_i k_i p^i$ is the p-adic representation of k.

To deduce the last statement of the lemma, observe that in $\mathbb{Z}_p$, $\prod_i \binom{n_i}{k_i}$ is nonzero exactly when the implication $n_i = 0 \Rightarrow k_i = 0$ holds for all i. By inspection of p-adic representations it is evident that there is a positive $k < n$ satisfying these implications iff n is not a power of p. □

The characteristic 0 case of Example 3 implies that $\mathbb{C}\mathrm{P}^n$ is not an H-space for finite n, in contrast with $\mathbb{C}\mathrm{P}^\infty$ which is an H-space as we saw above. Similarly, taking $F = \mathbb{Z}_2$, we deduce that $\mathbb{R}\mathrm{P}^n$ can be an H-space only if $n+1$ is a power of 2. Indeed, $\mathbb{R}\mathrm{P}^1 = S^1/\pm 1$, $\mathbb{R}\mathrm{P}^3 = S^3/\pm 1$, and $\mathbb{R}\mathrm{P}^7 = S^7/\pm 1$ have quotient H-space structures from S^1, S^3 and S^7 since -1 commutes with all elements of S^1, S^3, or S^7. However, these are the only cases when $\mathbb{R}\mathrm{P}^n$ is an H-space since by an exercise for §3.4c the universal cover of an H-space is an H-space, and S^1, S^3, and S^7 are the only spheres which are H-spaces, by a famous theorem of J.F.Adams; see [VBKT] for a proof.

It is an easy exercise to check that the tensor product of Hopf algebras is again a Hopf algebra, with the obvious coproduct $\Delta(\alpha \otimes \beta) = \Delta(\alpha) \otimes \Delta(\beta)$. So the preceding examples yield many other Hopf algebras, tensor products of polynomial, truncated polynomial, and exterior algebras on any number of generators. The following theorem of Hopf is a partial converse:

Theorem 3.42. *If A is a commutative, associative Hopf algebra over a field F of characteristic 0, and A^n is finite-dimensional over F for each n, then A is isomorphic as an algebra to the tensor product of an exterior algebra on odd-dimensional generators and a polynomial algebra on even-dimensional generators.*

There is an analogous theorem of Borel in case F is a finite field of characteristic p. In this case A is again isomorphic to a tensor product of single-generator Hopf algebras, of one of the following types:

(a) $F[\alpha]$, with α even-dimensional if $p \neq 2$.
(b) $\Lambda_F[\alpha]$ with α odd-dimensional.
(c) $\mathbb{Z}_p[\alpha]/(\alpha^{p^i})$, with α even-dimensional if $p \neq 2$.

For a proof see [Borel] or [Kane].

Proof of 3.42: Say A is generated as an algebra by $x_1, x_2, \cdots$ with $\dim x_i \leq \dim x_{i+1}$. Let A_n be the subalgebra generated by $x_1, \cdots, x_n$. This is a sub Hopf algebra of A, i.e., $\Delta(A_n) \subset A_n \otimes A_n$, since $\Delta(x_i)$ involves only x_i and terms of dimension less than i. We may assume x_n does not lie in A_{n-1}. Then to prove the theorem it suffices by induction

to prove injectivity of the natural surjection $A_{n-1}\otimes F[x_n]\to A_n$ if the dimension d_n of x_n is even, or of $A_{n-1}\otimes\Lambda_F[x_n]\to A_n$ if d_n is odd. Thus in the two cases we must rule out nontrivial relations $\sum_i \alpha_i x_n^i = 0$ and $\alpha_0+\alpha_1 x_n = 0$, respectively, with coefficients $\alpha_i \in A_{n-1}$.

Let I be the ideal in A_n generated by the positive-dimensional elements of A_{n-1}, so I consists of the polynomials $\sum_i \alpha_i x_n^i$ with positive-dimensional coefficients $\alpha_i \in A_{n-1}$. Note that $x_n \notin I$ since elements of I having dimension d_n must lie in A_{n-1}. Consider the composition Δ':

$$A_n \xrightarrow{\Delta} A_n\otimes A_n \longrightarrow A_n\otimes(A_n/I)$$

By the definition of I, Δ' sends $\alpha\in A_{n-1}$ to $\alpha\otimes 1$ and x_n to $x_n\otimes 1+1\otimes x_n'$ where x_n' is the image of x_n in A_n/I. In case d_n is even, applying Δ' to a nontrivial relation $\sum_i \alpha_i x_n^i = 0$ gives

$$0=\sum_i(\alpha_i\otimes 1)(x_n\otimes 1+1\otimes x_n')^i=\big(\sum_i \alpha_i x_n^i\big)\otimes 1+\sum_i i\alpha_i x_n^{i-1}\otimes x_n'+\cdots$$

Since this latter sum is zero in the tensor product $A_n\otimes(A_n/I)$, each summand must be zero for dimension reasons. In particular $\sum_i i\alpha_i x_n^{i-1}\otimes x_n' = 0$, hence $\sum_i i\alpha_i x_n^{i-1} = 0$ since $x_n\notin I$ implies $x_n'\neq 0$. The relation $\sum_i i\alpha_i x_n^{i-1}=0$ has lower degree than the original relation, and is not the trivial relation since F has characteristic 0, $\alpha_i\neq 0$ implying $i\alpha_i\neq 0$. Since we could assume the original relation had minimum degree, we have reached a contradiction.

The case d_n odd is similar. Applying Δ' to a relation $\alpha_0+\alpha_1 x_n=0$ gives $\alpha_0\otimes 1+(\alpha_1\otimes 1)(x_n\otimes 1+1\otimes x_n')=\alpha_1\otimes x_n'=0$, which implies $\alpha_1=0$ and hence $\alpha_0=0$. □

The structure of Hopf algebras over $\mathbb{Z}$ is quite a bit more complicated. Here is an example.

Example 4. We showed in §3.2 that the H-space $J(S^n)$ for n even has $H^*(J(S^n);\mathbb{Z})$ a **divided polynomial algebra**, the algebra $\Gamma_{\mathbb{Z}}[\alpha]$ with additive generators α_i in dimension $2i$ and multiplication given by $\alpha_1^k = k!\alpha_k$, hence $\alpha_i\alpha_j=\binom{i+j}{i}\alpha_{i+j}$. The coproduct in $\Gamma_{\mathbb{Z}}[\alpha]$ is uniquely determined by the multiplicative structure since $\Delta(\alpha_1^k)=(\alpha_1\otimes 1+1\otimes\alpha_1)^k=\sum_i\binom{k}{i}\alpha_1^i\otimes\alpha_1^{k-i}$, hence $\Delta(\alpha_1^k/k!)=\sum_i(\alpha_1^i/i!)\otimes(\alpha_1^{k-i}/(k-i)!)$, i.e., $\Delta(\alpha_k)=\sum_i\alpha_i\otimes\alpha_{k-i}$. Thus the coproduct has a simpler description than the product.

It is interesting to see what happens to the divided polynomial algebra $\Gamma_{\mathbb{Z}}[\alpha]$ when we change to field coefficients. Clearly $\Gamma_{\mathbb{Q}}[\alpha]$ is the same as $\mathbb{Q}[\alpha]$. In contrast with this, $\Gamma_{\mathbb{Z}_p}[\alpha]$, with multiplication defined by $\alpha_i\alpha_j=\binom{i+j}{i}\alpha_{i+j}$ is isomorphic as an algebra to

the infinite tensor product $\otimes_{i\geq 1}\mathbb{Z}_p[\alpha_{p^i}]/(\alpha_{p^i}^p)$, as one can check by looking at binomial coefficients $\mod p$. The coproduct in $\Gamma_{\mathbb{Z}_p}[\alpha]$ is defined by $\Delta(\alpha_k)=\sum_i \alpha_i\otimes\alpha_{k-i}$, so the algebras $\Gamma_{\mathbb{Z}_p}[\alpha]$ and $\otimes_{i\geq 1}\mathbb{Z}_p[\alpha_{p^i}]/(\alpha_{p^i}^p)$ are distinct as Hopf algebras since α_{p^i} is primitive in $\otimes_{i\geq 1}\mathbb{Z}_p[\alpha_{p^i}]/(\alpha_{p^i}^p)$ but not in $\Gamma_{\mathbb{Z}_p}[\alpha]$.

Pontryagin Product

Another special feature of H-spaces is that their homology groups have a product operation, called the **Pontryagin product**. This is the composition

$$H_*(X;R)\otimes H_*(X;R)\xrightarrow{\times} H_*(X\times X;R)\xrightarrow{\mu_*} H_*(X;R)$$

where $\times$ is the homology cross product defined in the preceding section and $\mu: X\times X\to X$ is the H-space multiplication. Thus the Pontryagin product gives maps $H_i(X;R)\times H_j(X;R)\to H_{i+j}(X;R)$. These satisfy the distributive law, and are associative if μ is associative or at least associative up to homotopy, in the sense that the maps $X\times X\times X\to X$, $(x,y,z)\mapsto\mu(x,\mu(y,z))$ and $(x,y,z)\mapsto\mu(\mu(x,y),z)$ are homotopic.

In case X is a CW complex and μ is a cellular map the Pontryagin product can be computed using cellular homology via the cellular chain map

$$C_i(X;R)\times C_j(X;R)\xrightarrow{\times} C_{i+j}(X\times X;R)\xrightarrow{\mu_*} C_{i+j}(X;R)$$

where $\times$ sends generators corresponding to cells e^i and e^j to the generator corresponding to the product cell $e^i\times e^j$, and then μ_* is applied to this product cell.

Example 5. Let us compute the Pontryagin product for $J(S^n)$. Here there is one cell e^{in} for each $i\geq 0$, and μ takes the product cell $e^{in}\times e^{jn}$ homeomorphically onto the cell $e^{(i+j)n}$. This means that $H_*(J(S^n);\mathbb{Z})$ is simply the polynomial ring $\mathbb{Z}[x]$ on an n-dimensional generator x. This holds for n odd as well as for n even, so the Pontryagin product need not satisfy the same general commutativity relation as cup product. In this example the Pontryagin product structure is simpler than the cup product structure, though for some H-spaces it is the other way round. In applications it is often convenient to have the choice of which product structure to use.

This calculation immediately generalizes to $J(X)$ where X is any connected CW complex whose cellular boundary maps are all trivial. The cellular boundary maps in the product X^m of m copies of X are then trivial by induction on m using Proposition 3.35, and therefore the cellular boundary maps in $J(X)$ are all trivial since the quotient map $X^m\to J_m(X)$ is cellular and each cell of $J_m(X)$ is the homeomorphic image of a cell

of X^m. Thus $H_*(J(X);\mathbb{Z})$ is free with additive basis the products $e^{n_1}\times\cdots\times e^{n_k}$ of positive-dimensional cells of X, and the multiplicative structure is that of polynomials in noncommuting variables corresponding to the positive-dimensional cells of X. Another way to describe this algebra is as the tensor algebra $T\widetilde{H}_*(X;\mathbb{Z})$, which is additively the direct sum of the n-fold tensor products of $\widetilde{H}_*(X;\mathbb{Z})$ with itself for all $n\geq 0$, with the obvious multiplication, and with the convention that the 0-fold tensor product is just the coefficient ring, $\mathbb{Z}$ in the present case.

More generally, $H_*(J(X);\mathbb{Z})$ is isomorphic to the tensor algebra $T\widetilde{H}_*(X;\mathbb{Z})$ whenever X is a connected CW complex with $H_*(X;\mathbb{Z})$ free. This holds also for homology with coefficients in a field, when the freeness is automatic. Here is the proof. With either choice of coefficient ring, let $\varphi\colon T\widetilde{H}_*(X)\to H_*(J(X))$ be the homomorphism whose restriction to the n-fold tensor product $\widetilde{H}_*(X)^{\otimes n}$ is the composition

$$\widetilde{H}_*(X)^{\otimes n}\hookrightarrow H_*(X)^{\otimes n}\to H_*(X^n)\to H_*(J_n(X))\to H_*(J(X))$$

where the next-to-last map is induced by the quotient map $X^n\to J_n(X)$. It is clear that φ is a ring homomorphism since the product in $J(X)$ is induced from the natural map $X^m\times X^n\to X^{m+n}$. To show φ is an isomorphism, consider the following commutative diagram of short exact sequences:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & T_{n-1}\widetilde{H}_*(X) & \longrightarrow & T_n\widetilde{H}_*(X) & \longrightarrow & \widetilde{H}_*(X)^{\otimes n} & \longrightarrow & 0\\
 & & \downarrow & & \downarrow & & \downarrow\approx & & \\
0 & \longrightarrow & H_*(J_{n-1}(X)) & \longrightarrow & H_*(J_n(X)) & \longrightarrow & \widetilde{H}_*(X^{\wedge n}) & \longrightarrow & 0
\end{array}$$

In the upper row, $T_m\widetilde{H}_*(X)$ denotes the direct sum of the products $\widetilde{H}_*(X)^{\otimes k}$ for $k\leq m$, so this row is exact. The second row is the homology exact sequence for the pair $(J_n(X),J_{n-1}(X))$, with quotient $J_n(X)/J_{n-1}(X)$ the n-fold smash product $X^{\wedge n}$. This long exact sequence breaks up into short exact sequences as indicated, by commutativity of the right-hand square and the fact that the right-hand vertical map is an isomorphism by the Künneth formula, using the hypothesis that $H_*(X)$ is free over the given coefficient ring. By induction on n and the five-lemma we deduce from the diagram that $\varphi\colon T_n\widetilde{H}_*(X)\to H_*(J_n(X))$ is an isomorphism for all n. Passing to the direct limit as n goes to ∞, this implies that $\varphi\colon T\widetilde{H}_*(X)\to H_*(J(X))$ is an isomorphism since in any given dimension $T_n\widetilde{H}_*(X)$ is independent of n when n is sufficiently large, and the same is true of $H_*(J_n(X))$ by the second row of the diagram above.

This calculation can be paraphrased as saying that 'the homology of the free H-space generated by a space with free homology is the free algebra generated by the homology of the space.'

Dual Hopf Algebras

There is a close connection between the Pontryagin product in homology and the Hopf algebra structure on cohomology. Suppose that X is an H-space such that, with coefficients in a field R, the vector spaces $H_n(X;R)$ are finite-dimensional for all n. Alternatively, we could take $R = \mathbb{Z}$ and assume $H_n(X;\mathbb{Z})$ is finitely-generated and free for all n. In either case we have $H^n(X;R) = Hom_R(H_n(X;R),R)$, and the Pontryagin product $H_*(X;R)\otimes H_*(X;R) \to H_*(X;R)$ and the coproduct $\Delta : H^*(X;R) \to H^*(X;R)\otimes H^*(X;R)$ are dual to each other, both being induced by the H-space product $\mu : X\times X \to X$. Therefore the coproduct in cohomology determines the Pontryagin product in homology, and vice versa. Specifically, the component $\Delta_{ij} : H^{i+j}(X;R) \to H^i(X;R)\otimes H^j(X;R)$ of Δ is dual to the product $H_i(X;R)\otimes H_j(X;R) \to H_{i+j}(X;R)$.

Example 6. Consider $J(S^n)$ with n even, so $H^*(J(S^n);\mathbb{Z})$ is the divided polynomial algebra $\Gamma_{\mathbb{Z}}[\alpha]$. In Example 4 we derived the coproduct formula $\Delta(\alpha_k) = \sum_i \alpha_i \otimes \alpha_{k-i}$. Thus Δ_{ij} takes α_{i+j} to $\alpha_i \otimes \alpha_j$, so if x_i is the generator of $H_{in}(J(S^n);\mathbb{Z})$ dual to α_i, then $x_i x_j = x_{i+j}$. This says that $H_*(J(S^n);\mathbb{Z})$ is the polynomial ring $\mathbb{Z}[x]$. We showed this in Example 5 using the cell structure of $J(S^n)$, but the present proof deduces it purely algebraically from the cup product structure.

Now we wish to show that the relation between $H^*(X;R)$ and $H_*(X;R)$ is perfectly symmetric: they are *dual Hopf algebras*. This is a purely algebraic fact:

Proposition 3.43. *Let A be a Hopf algebra over R which is a finitely-generated free R-module in each dimension. Then the product $\pi : A\otimes A \to A$ and coproduct $\Delta : A \to A\otimes A$ have duals $\pi^* : A^* \to A^*\otimes A^*$ and $\Delta^* : A^*\otimes A^* \to A^*$ which give A^* the structure of a Hopf algebra.*

Proof: This will be apparent if we reinterpret the Hopf algebra structure on A more formally as a pair of graded R-module homomorphisms $\pi : A\otimes A \to A$ and $\Delta : A \to A\otimes A$ together with an element $1 \in A^0$ satisfying:

(1) The two compositions $A \xrightarrow{i_\ell} A\otimes A \xrightarrow{\pi} A$ and $A \xrightarrow{i_r} A\otimes A \xrightarrow{\pi} A$ are the identity, where $i_\ell(a) = a\otimes 1$ and $i_r(a) = 1\otimes a$. This says that 1 is a two-sided identity for the multiplication in A.

(2) The two compositions $A \xrightarrow{\Delta} A\otimes A \xrightarrow{p_\ell} A$ and $A \xrightarrow{\Delta} A\otimes A \xrightarrow{p_r} A$ are the identity, where $p_\ell(a\otimes 1) = a$, $p_\ell(a\otimes b) = 0$ if $b \in A^j$ with $j > 0$, $p_r(1\otimes a) = a$, and $p_r(a\otimes b) = 0$ if $a \in A^j$ with $j > 0$. This is the formula $\Delta(a) = a\otimes 1 + 1\otimes a + \sum_i a'_i \otimes \alpha''_{n-i}$.

(3) The diagram

$$\begin{array}{ccccc} A\otimes A & \xrightarrow{\pi} & A & \xrightarrow{\Delta} & A\otimes A \\ \downarrow{\scriptstyle \Delta\otimes\Delta} & & & & \uparrow{\scriptstyle \pi\otimes\pi} \\ A\otimes A\otimes A\otimes A & & \xrightarrow{\tau} & & A\otimes A\otimes A\otimes A \end{array}$$

commutes, where $\tau(a\otimes b\otimes c\otimes d) = (-1)^{ij}a\otimes c\otimes b\otimes d$ for $b\in A^i$ and $c\in A^j$. This is the condition that Δ is an algebra homomorphism since if we follow an element $a\otimes b\in A^m\otimes A^n$ across the top of the diagram we get $\Delta(ab)$, while via the lower route we get first $\Delta(a)\otimes\Delta(b) = \big(\sum_i a'_i\otimes a''_{m-i}\big)\otimes\big(\sum_j b'_j\otimes b''_{n-j}\big)$, then after applying τ this becomes $\sum_{i,j}(-1)^{(m-i)j}a'_ib'_j\otimes a''_{m-i}b''_{n-j}$ which equals $\big(\sum_i a'_i\otimes a''_{m-i}\big)\big(\sum_j b'_j\otimes b''_{n-j}\big) = \Delta(a)\Delta(b)$.

Condition (1) for A dualizes to (2) for A^*, and similarly (2) for A dualizes to (1) for A^*. Condition (3) for A dualizes to (3) for A^*. □

Example 7. Let us compute the dual of a polynomial algebra $R[x]$. Suppose first that x has even dimension. Then $\Delta(x^n) = (x\otimes 1 + 1\otimes x)^n = \sum_i \binom{n}{i}x^i\otimes x^{n-i}$, so if α_i is dual to x^i, the term $\binom{n}{i}x^i\otimes x^{n-i}$ in $\Delta(x^n)$ gives the product relation $\alpha_i\alpha_{n-i} = \binom{n}{i}\alpha_n$. This is the rule for multiplication in a divided polynomial algebra, so the dual of $R[x]$ is $\Gamma_R[\alpha]$ if the dimension of x is even. This also holds if $2 = 0$ in R, since the even-dimensionality of x was used only to deduce that $R[x]\otimes R[x]$ is strictly commutative.

In case x is odd-dimensional, then as we saw in Example 1, if we set $y = x^2$, we have $\Delta(y^n) = (y\otimes 1 + 1\otimes y)^n = \sum_i \binom{n}{i}y^i\otimes y^{n-i}$ and $\Delta(xy^n) = \Delta(x)\Delta(y^n) = \sum_i \binom{n}{i}xy^i\otimes y^{n-i} + \sum_i \binom{n}{i}y^i\otimes xy^{n-i}$. These formulas for Δ say that the dual of $R[x]$ is $\Lambda_R[\alpha]\otimes\Gamma_R[\beta]$ where α is dual to x and β is dual to y.

This algebra allows us to deduce the cup product structure on $H^*(J(S^n);R)$ from the geometric calculation $H_*(J(S^n);R)\approx R[x]$ in Example 5. As another application, recall from earlier in this section that $\mathbb{R}\mathrm{P}^\infty$ and $\mathbb{C}\mathrm{P}^\infty$ are H-spaces, so from their cup product structures we can conclude that the Pontryagin rings $H_*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$ and $H_*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ are divided polynomial algebras.

However, one has to remember that the product in a Hopf algebra does not usually determine the coproduct uniquely. Thus in Example 7 one starts with a polynomial algebra $R[x]$, where the coproduct is unique, but in the dual algebra $\Gamma_R[\alpha]$ or $\Lambda_R[\alpha]\otimes\Gamma_R[\beta]$ this may no longer be true. For example, in $\Gamma_R[\alpha]$ the coproduct is unique if $R = \mathbb{Z}$ but not if $R = \mathbb{Z}_p$ as we observed in Example 4 and the remarks following it. In particular, the algebra $\otimes_{i\ge 1}\mathbb{Z}_p[\alpha_{p^i}]/(\alpha_{p^i}^p)$ arises as the dual not only of $\mathbb{Z}_p[x]$ but also of $\otimes_{i\ge 1}\mathbb{Z}_p[x_{p^i}]/(x_{p^i}^p)$ where the latter tensor product means tensor product of Hopf algebras,

and we use that fact that the Hopf algebra $\mathbb{Z}_p[\alpha]/(\alpha^p)$ is its own dual, one of the exercises below.

Exercises

1. Suppose that X is a CW complex with basepoint $e \in X$ a 0-cell. Show that X is an H-space if there is a map $\mu : X \times X \to X$ such that the maps $X \to X$, $x \mapsto \mu(x, e)$ and $x \mapsto \mu(e, x)$, are homotopic to the identity. [Sometimes this is taken as the definition of an H-space, rather than the slightly more restrictive condition in the definition we have given.] With the same hypotheses, show also that μ can be homotoped so that e is a strict two-sided identity.

2. Show that a retract of an H-space is an H-space if it contains the identity element.

3. Show that if X is an H-space such that the set of path components of X is a group with respect to the multiplication induced by the H-space structure, then all the path components of X are homotopy equivalent.

4. Show that an H-space or topological group structure on a path-connected, locally path-connected space can be lifted to such a structure on its universal cover. [In the case of $SO(n)$, the universal cover for $n > 2$ is a 2-sheeted cover, a group called $Spin(n)$.]

5. Show that if (X, e) is an H-space then $\pi_1(X, e)$ is abelian. [Compare the usual composition $f \cdot g$ of loops with the product $\mu(f(t), g(t))$ coming from the H-space multiplication μ.]

6. Show that S^n is an H-space iff the attaching map of the $2n$-cell of $J_2(S^n)$ is homotopically trivial.

7. What are the primitive elements of the Hopf algebra $\mathbb{Z}_p[x]$ for p prime?

8. Show that the tensor product of two Hopf algebras is a Hopf algebra.

9. Apply the theorems of Hopf and Borel to show that a finite CW complex H-space X such that $\widetilde{H}_*(X; \mathbb{Z}) \neq 0$ has Euler characteristic $\chi(X) = 0$.

10. Let X be a path-connected H-space with $H^*(X; R)$ free and finitely-generated in each dimension. For maps $f, g : X \to X$, the product $fg : X \to X$ is defined by $(fg)(x) = f(x)g(x)$, using the H-space product.

(a) Show that $(fg)^*(\alpha) = f^*(\alpha) + g^*(\alpha)$ for primitive elements $\alpha \in H^*(X; R)$.

(b) Deduce that the k^{th}-power map $x \mapsto x^k$ induces the map $\alpha \mapsto k\alpha$ on primitive elements α. In particular the quaternionic k^{th}-power map $S^3 \to S^3$ has degree k.

(c) Show that every polynomial $a_n x^n b_n + \cdots + a_1 x b_1 + a_0$ with coefficients in $\mathbb{H}$ has a root in $\mathbb{H}$ if $n > 0$. [See Theorem 1.9.]

11. If T^n is the n-dimensional torus, the product of n circles, show that the Pontryagin ring $H_*(T^n;\mathbb{Z})$ is the exterior algebra $\Lambda_{\mathbb{Z}}[x_1,\cdots,x_n]$ on one-dimensional generators x_i.

12. Compute the Pontryagin product structure in $H_*(L;\mathbb{Z}_p)$ where L is an infinite-dimensional lens space $S^\infty/\mathbb{Z}_p$, for p an odd prime, using the coproduct in $H^*(L;\mathbb{Z}_p)$.

13. Verify that the Hopf algebras $\Lambda_R[\alpha]$ and $\mathbb{Z}_p[\alpha]/(\alpha^p)$ are self-dual.

14. Show that the coproduct in the Hopf algebra $H_*(X;R)$ dual to $H^*(X;R)$ is induced by the diagonal map $X \to X\times X$, $x\mapsto(x,x)$.

15. Suppose that X is a path-connected H-space such that $H^*(X;\mathbb{Z})$ is free and finitely generated in each dimension, and $H^*(X;\mathbb{Q})$ is a polynomial ring $\mathbb{Q}[\alpha]$. Show that the Pontryagin ring $H_*(X;\mathbb{Z})$ is commutative and associative, with a structure uniquely determined by the ring $H^*(X;\mathbb{Z})$.

16. Classify algebraically the Hopf algebras A over $\mathbb{Z}$ such that $A\otimes\mathbb{Q}\approx\mathbb{Q}[\alpha]$. In particular, determine which Hopf algebras $A\otimes\mathbb{Z}_p$ arise from such A's.

d. The Cohomology of SO(n)

An interesting object to consider from the viewpoint of algebraic topology is the orthogonal group $O(n)$, consisting of the isometries of $\mathbb{R}^n$ which fix the origin. Equivalently, this is the group of $n \times n$ matrices A with entries in $\mathbb{R}$ satisfying $AA^t = I$ where A^t is the transpose of A. From this viewpoint, $O(n)$ is topologized as a subspace of $\mathbb{R}^{n^2}$, with coordinates the n^2 entries of an $n \times n$ matrix. Since the columns of a matrix in $O(n)$ are unit vectors, $O(n)$ can also be regarded as a subspace of the product of n copies of S^{n-1}. It is a closed subspace since the conditions that columns be orthogonal are defined by polynomial equations. Hence $O(n)$ is compact. The map $O(n) \times O(n) \to O(n)$ given by matrix multiplication is continuous since it is defined by polynomials. The inversion map $A \mapsto A^{-1} = A^t$ is clearly continuous, so $O(n)$ is a topological group, and in particular an H-space.

The determinant map $O(n) \to \{\pm 1\}$ is a homomorphism, so its kernel $SO(n)$, the 'special orthogonal group,' is a subgroup of index two. The two cosets $SO(n)$ and $O(n) - SO(n)$ are homeomorphic to each other since for fixed $B \in O(n) - SO(n)$, the maps $A \mapsto AB$ and $A \mapsto AB^{-1}$ are inverse homeomorphisms between these two cosets. As we shall see below, $SO(n)$ is path-connected, so it can be described as the identity component of $O(n)$, the component containing the identity matrix.

The group $O(n)$ is a subgroup of $GL_n(\mathbb{R})$, the 'general linear group' of all invertible $n \times n$ matrices with entries in $\mathbb{R}$, topologized as a subspace of $\mathbb{R}^{n^2}$. In fact, $GL_n(\mathbb{R})$ is an open subset of $\mathbb{R}^{n^2}$ since it is the subspace consisting of $n \times n$ matrices with nonzero determinant, and the determinant map $\mathbb{R}^{n^2} \to \mathbb{R}$ is continuous. The Gram-Schmidt orthogonalization process applied to the columns of matrices in $GL_n(\mathbb{R})$ provides a retraction $r : GL_n(\mathbb{R}) \to O(n)$, continuity of r being evident from the explicit formulas for the Gram-Schmidt process. By inserting appropriate scalar factors into these formulas it is easy to see that $O(n)$ is in fact a deformation retract of $GL_n(\mathbb{R})$. (Using a bit more linear algebra, namely the polar decomposition, it is possible to show that $GL_n(\mathbb{R})$ is homeomorphic to $O(n) \times \mathbb{R}^k$ for $k = n(n+1)/2$.)

The topological structure of $SO(n)$ for small values of n can be described in terms of more familiar spaces:

— $SO(1)$ is a point.

— $SO(2)$, the rotations of $\mathbb{R}^2$, is both homeomorphic and isomorphic as a group to S^1, thought of as the unit complex numbers.

— $SO(3)$ is homeomorphic to $\mathbb{R}\mathrm{P}^3$. To see this, let $\varphi : D^3 \to SO(3)$ send a nonzero vector x to the rotation through angle $\pi|x|$ about the line determined by x. An ori-

entation convention such as the 'right-hand rule' is needed to make this unambiguous. By continuity, φ must send 0 to the identity. Antipodal points of $S^2 = \partial D^3$ are sent to the same rotation through angle π, so φ induces a map $\overline{\varphi}\colon \mathbb{R}\mathrm{P}^3 \to SO(3)$, regarding $\mathbb{R}\mathrm{P}^3$ as D^3 with antipodal boundary points identified. The map $\overline{\varphi}$ is clearly injective since the axis of a nontrivial rotation is uniquely determined as its fixed point set, and $\overline{\varphi}$ is surjective since by easy linear algebra each non-identity element of $SO(3)$ is a rotation about some axis. It follows that $\overline{\varphi}$ is a homeomorphism $\mathbb{R}\mathrm{P}^3 \approx SO(3)$.

— One can go one step further and show that $SO(4)$ is homeomorphic to $S^3 \times SO(3)$. Identifying $\mathbb{R}^4$ with the quaternions $\mathbb{H}$ and S^3 with the group of unit quaternions, the quaternion multiplication $v \mapsto vw$ for fixed $w \in S^3$ defines an isometry $\rho_w \in O(4)$ since $|vw| = |v||w|$. Points of $O(4)$ are 4-tuples $(v_1, \cdots, v_4)$ of orthonormal vectors $v_i \in \mathbb{H} = \mathbb{R}^4$, and $O(3)$ can be viewed as the subspace with $v_4 = 1$. Now define a map $S^3 \times O(3) \to O(4)$ by sending $(v, (v_1, v_2, v_3))$ to $\rho_v(v_1, v_2, v_3, 1)$. This is a homeomorphism since it has an inverse given by $(v_1, v_2, v_3, v) \mapsto (v, \rho_{v^{-1}}(v_1, v_2, v_3, v))$. Restricting to identity components gives a homeomorphism $S^3 \times SO(3) \approx SO(4)$. This is not a group isomorphism, however. It can be shown, though we will not digress to do so here, that the homomorphism $\psi\colon S^3 \times S^3 \to SO(4)$ sending a pair (u, v) of unit quaternions to the isometry $w \mapsto uwv^{-1}$ of $\mathbb{H}$ is surjective with kernel $\mathbb{Z}_2 = \{\pm(1,1)\}$, and that ψ is a covering space projection, representing $S^3 \times S^3$ as a 2-sheeted cover of $SO(4)$, the universal cover. Further, restricting ψ to the diagonal $S^3 = \{(u,u)\} \subset S^3 \times S^3$ gives the universal cover $S^3 \to SO(3)$, so that $SO(3)$ is isomorphic to the quotient group of S^3 by the normal subgroup $\{\pm 1\}$.

Using octonions one can construct in the same way a homeomorphism $S^7 \times SO(7) \approx SO(8)$. But in all other cases there is only a fiber bundle $SO(n-1) \to SO(n) \to S^{n-1}$ expressing $SO(n)$ as a 'twisted product' of $SO(n-1)$ and S^{n-1}; see §4.3.

Cell Structure

We shall compute the homology and cohomology of $SO(n)$ by constructing a very nice CW structure on it, together with a cellular map $\rho\colon P^{n-1} \times P^{n-2} \times \cdots \times P^1 \to SO(n)$ where $P^i = \mathbb{R}\mathrm{P}^i$.

For a nonzero vector $v \in \mathbb{R}^n$ let $r(v) \in O(n)$ be the reflection across the hyperplane orthogonal to v, so that $r(v) = -v$ and $r(w) = w$ if $w \perp v$. Let $\rho(v) = r(v)r(e_1) \in SO(n)$ where $e_1 = (1, 0, \cdots, 0)$, the first standard basis vector of $\mathbb{R}^n$. Thus $\rho(v)$ rotates the plane containing e_1 and v through twice the angle from e_1 to v, and $\rho(v)$ is the identity on the orthogonal subspace. Since $\rho(v)$ depends only on the line spanned by v, ρ defines

a map $P^{n-1} \to SO(n)$. It is easy to see that this map is injective, so we may think of P^{n-1} as a subspace of $SO(n)$. More generally, for a sequence $I = (i_1, \cdots, i_m)$ with each $i_j < n$, a map $\rho : P^I = P^{i_1} \times \cdots \times P^{i_m} \to SO(n)$ is defined by $\rho(v_1, \cdots, v_m) = \rho(v_1) \cdots \rho(v_m)$, composing the rotations $\rho(v_j)$. Let $\varphi^i : D^i \to P^i$ be a characteristic map for the i-dimensional cell of P^i, identifying antipodal points of ∂D^i, and let $\varphi^I : D^I \to P^I$ be the product of these maps, a characteristic map for the top-dimensional cell of P^I.

Proposition 3.44. *The maps $\rho\varphi^I : D^I \to SO(n)$ for $n > i_1 > \cdots > i_m > 0$, together with $\rho\varphi^0 : D^0 \to SO(n)$, are the characteristic maps of a CW structure on $SO(n)$ for which the map $\rho : P^{n-1} \times P^{n-2} \times \cdots \times P^1 \to SO(n)$ is cellular.*

Thus there is a single 0-cell $e^0 = \{\mathbb{1}\}$, the image of $\rho\varphi^0$. The other cells $e^I = e^{i_1} \cdots e^{i_m}$ are products, via the group operation in $SO(n)$, of the cells $e^i \subset P^{n-1} \subset SO(n)$. Note that under the standard inclusion of $SO(n-1)$ in $SO(n)$ as the isometries fixing the last standard basis vector $e_n = (0, \cdots, 0, 1)$, $SO(n-1)$ consists of the cells $e^{i_1} \cdots e^{i_m}$ with $i_1 < n-1$. Since the theorem also holds for $SO(n-1)$, it follows that $SO(n-1)$ is a subcomplex, and $SO(n)$ is obtained from $SO(n-1)$ by attaching the cells $e^{i_1} \cdots e^{i_m}$ with $i_1 = n-1$, i.e., the products of e^{n-1} with the cells of $SO(n-1)$. Thus $SO(n)$ has twice as many cells as $SO(n-1)$, hence, inductively, 2^{n-1} cells.

Proof: To obtain the CW structure there are three things to show:

(1) For each decreasing sequence I, $\rho\varphi^I$ is a homeomorphism from the interior of D^I onto its image.
(2) The resulting image cells e^I are all disjoint and cover $SO(n)$.
(3) For each e^I, $\rho\varphi^I(\partial D^I)$ is contained in a union of cells of lower dimension than e^I.

Define $p : SO(n) \to S^{n-1}$ by $p(\alpha) = \alpha(e_n)$. Elements of the subspace $P^{n-2} \subset SO(n)$ fix e_n, so $p(P^{n-2}) = \{e_n\}$. We claim that p is a homeomorphism from $P^{n-1} - P^{n-2}$ onto $S^{n-1} - \{e_n\}$. This can be seen as follows. We may view $P^{n-1} - P^{n-2}$ as the open upper hemisphere U of S^{n-1}. For a vector $v \in U$ orthogonal to e_1 and elevated an angle θ from the equator S^{n-2}, the 180 degree rotation $\rho(v)$ sends e_n to the point $p(v)$ in the plane of v and e_n at an angle 2θ down from e_n, as indicated in the figure. The point e_n is rotated to $p(e_n) = -e_n$. Thus the subspace of U orthogonal to e_1 is stretched homeomorphically over the subspace of $S^{n-1} - \{e_n\}$ orthogonal to e_1. Moreover, great

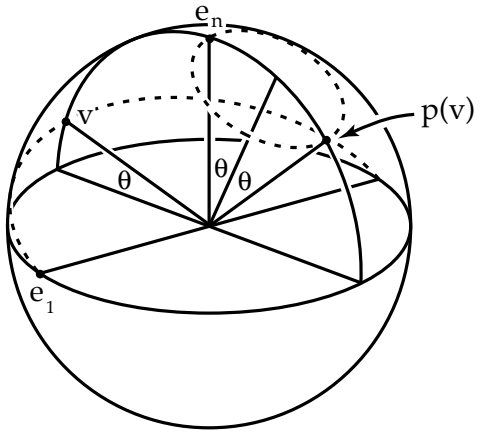

semicircles in U with endpoints e_1 and $-e_1$ are mapped by p to circular arcs with both endpoints at e_n. So it is clear that p is a homeomorphism from $P^{n-1} - P^{n-2}$ onto $S^{n-1} - \{e_n\}$.

The next statement is that the map

$$h : (P^{n-1} \times SO^{n-1}, P^{n-2} \times SO^{n-1}) \to (SO(n), SO(n-1)), \quad h(v, \alpha) = \rho(v)\alpha$$

is a homeomorphism from $(P^{n-1} - P^{n-2}) \times SO(n-1)$ onto $SO(n) - SO(n-1)$. For let $\beta \in SO(n) - SO(n-1)$ be given, and let $w_i = \beta(e_i)$. Then $w_n \neq e_n$ so by the preceding paragraph there is a unique $v \in P^{n-1} - P^{n-2}$ with $\rho(v)e_n = w_n$. The vectors $\rho(v)^{-1}(w_1), \cdots, \rho(v)^{-1}(w_n)$ form a positively oriented orthonormal basis for $\mathbb{R}^n$ so there is a unique $\alpha \in SO(n)$ taking the standard basis to this one. Since $\rho(v)^{-1}(w_n) = e_n$, α lies in $SO(n-1) \subset SO(n)$. Thus we have expressed $\beta \in SO(n) - SO(n-1)$ uniquely as $h(v, \alpha) = \rho(v)\alpha$ for $v \in P^{n-1} - P^{n-2}$ and $\alpha \in SO(n-1)$. This means h is a continuous bijection between the desired subspaces. Since h takes compact sets to compact sets, it takes closed sets to closed sets, so h^{-1} is continuous on $SO(n) - SO(n-1)$.

Statements (1) and (2) now follow by induction on n. To prove (3) it suffices to show that $P^iP^i \subset P^iP^{i-1}$ in $SO(n)$ since for a strictly decreasing sequence I, the map $\rho : P^I \to SO(n)$ takes the boundary of the top-dimensional cell to the image of products P^J with J just a weakly decreasing sequence. As a preliminary to showing $P^iP^i \subset P^iP^{i-1}$, observe that for $\alpha \in O(n)$ we have $r(\alpha(v)) = \alpha r(v)\alpha^{-1}$. Hence $\rho(v)\rho(w) = r(v)r(e_1)r(w)r(e_1) = r(v)r(w')$ where $w' = r(e_1)w$. Thus to show $P^iP^i \subset P^iP^{i-1}$ it suffices, for $v, w \in \mathbb{R}^{i+1}$, to find $x \in \mathbb{R}^{i+1}$ and $y \in \mathbb{R}^i$ with $r(v)r(w) = r(x)r(y)$.

Let $V \subset \mathbb{R}^{i+1}$ be a 2-plane containing v and w. Since $V \cap \mathbb{R}^i$ is at least 1-dimensional, we can choose a unit vector $y \in V \cap \mathbb{R}^i$. Let $\alpha \in O(i+1)$ take V to $\mathbb{R}^2$ and y to e_1. Then the conjugate $\alpha r(v)r(w)\alpha^{-1} = r(\alpha(v))r(\alpha(w))$ lies in $SO(2)$, hence has the form $\rho(z) = r(z)r(e_1)$ for some $z \in \mathbb{R}^2$. Therefore $r(v)r(w) = r(\alpha^{-1}(z))r(\alpha^{-1}(e_1)) = r(x)r(y)$ for $x = \alpha^{-1}(z)$.

It remains to show that $\rho : P^{n-1} \times P^{n-2} \times \cdots \times P^1 \to SO(n)$ is cellular. This follows from the inclusions $P^iP^i \subset P^iP^{i-1}$ derived above, together with inclusions $P^iP^j \subset P^jP^i$ for $i < j$. To prove the latter we have the formulas

$$\begin{aligned}
\rho(v)\rho(w) &= r(v)r(w') \qquad \text{where } w' = r(e_1)w, \text{ as earlier} \\
&= r(v)r(w')r(v)r(v) \\
&= r\big(r(v)w'\big)r(v) \qquad \text{from } r\big(\alpha(v)\big) = \alpha r(v)\alpha^{-1} \\
&= r\big(r(v)r(e_1)w\big)r(v) = r\big(\rho(v)w\big)r(v) \\
&= \rho\big(\rho(v)w\big)\rho(v') \qquad \text{where } v' = r(e_1)v, \text{ hence } v = r(e_1)v'
\end{aligned}$$

Thus if $v \in \mathbb{R}^{i+1}$ and $w \in \mathbb{R}^{j+1}$ with $i < j$ then $\rho(v)w \in \mathbb{R}^{j+1}$ so the composition $\rho(v)\rho(w) = \rho(\rho(v)w)\rho(v') \in P^iP^j$ lies in P^jP^i. □

Mod 2 Homology and Cohomology

The fact that each cell of $SO(n)$ is the homeomorphic image of a cell in P implies that for $\mathbb{Z}_2$ coefficients, the cellular boundary maps for $SO(n)$ are all trivial, since this is true in P^i hence in $P^{n-1} \times P^{n-2} \times \cdots \times P^1$ by Proposition 3.35, and a cellular map induces a chain map of cellular chain complexes. Thus $H_*(SO(n); \mathbb{Z}_2)$ has a $\mathbb{Z}_2$ summand for each cell of $SO(n)$. One can rephrase this as saying that there are isomorphisms $H_i(SO(n); \mathbb{Z}_2) \approx H_i(S^1 \times S^2 \times \cdots \times S^{n-1}; \mathbb{Z}_2)$ for all i. The full structure of the $\mathbb{Z}_2$ homology and cohomology rings is given by:

Theorem 3.45.

(a) $H^*(SO(n); \mathbb{Z}_2) \approx \bigotimes_{i \text{ odd}} \mathbb{Z}_2[\beta_i]/(\beta_i^{p_i})$ *where* $|\beta_i| = i$ *and* p_i *is the smallest power of* 2 *such that* $|\beta_i^{p_i}| \geq n$.

(b) *The Pontryagin ring* $H_*(SO(n); \mathbb{Z}_2)$ *is the exterior algebra* $\Lambda_{\mathbb{Z}_2}[e^1, \cdots, e^{n-1}]$.

Here β_i is the dual class to e^i, represented by the cellular cochain assigning the value 1 to the cell e^i and 0 to all other i-cells.

Proof: The cup product structure will be determined using the map $\rho\colon P^{n-1} \times \cdots \times P^1 \to SO(n)$. Since all cellular boundary maps in both these spaces are zero with $\mathbb{Z}_2$ coefficients, and each cell of $SO(n)$ is the homeomorphic image of a cell of $P^{n-1} \times \cdots \times P^1$, ρ_* is surjective on $H_*(-; \mathbb{Z}_2)$, hence ρ^* is injective on $H^*(-; \mathbb{Z}_2)$. We know that $H^*(P^{n-1} \times \cdots \times P^1; \mathbb{Z}_2)$ is the polynomial ring $\mathbb{Z}_2[\alpha_1, \cdots, \alpha_{n-1}]$ truncated by the relations $\alpha_i^{i+1} = 0$. For $\beta_i \in H^i(SO(n); \mathbb{Z}_2)$ the dual class to e^i, we have $\rho^*(\beta_i) = \sum_j \alpha_j^i$, the class assigning 1 to each i-cell in a factor P^j of $P^{n-1} \times \cdots \times P^1$ and 0 to all other i-cells, which are products of lower-dimensional cells and hence map to cells in $SO(n)$ disjoint from e^i.

First we show that the monomials $\beta_{i_1} \cdots \beta_{i_m}$ with $n > i_1 > \cdots > i_m$ are linearly independent in $H^*(SO(n); \mathbb{Z}_2)$, hence are a basis. Since ρ^* is injective, we may identify each β_i with $\sum_j \alpha_j^i$ in the truncated polynomial ring $\mathbb{Z}_2[\alpha_1, \cdots, \alpha_{n-1}]/(\alpha_1^2, \cdots, \alpha_{n+1}^n)$. Suppose we have a linear relation $\sum_I b_I \beta_I = 0$ among the monomials $\beta_{i_1} \cdots \beta_{i_m}$ with $n > i_1 > \cdots > i_m$. If β_1 appears in this relation, write the relation in the form $x\beta_1 + y = 0$ where x and y do not involve β_1. Since α_1 occurs only in the term β_1 of $x\beta_1 + y$, where it has exponent 1, we have $x\beta_1 + y = x\alpha_1 + z$ where neither x nor z involve α_1. The relation $x\alpha_1 + z = 0$ in $\mathbb{Z}_2[\alpha_1, \cdots, \alpha_{n-1}]/(\alpha_1^2, \cdots, \alpha_{n-1}^n)$ then implies $x = 0$. Thus we may assume the original relation does not involve β_1. Now we repeat the argument for β_2. Write the relation in the form $x\beta_2 + y = 0$ where neither x nor y involve β_2 or β_1.

The variable α_2 now occurs only in the term β_2 of $x\beta_2 + y$, where it has exponent 2, so we have $x\beta_2 + y = x\alpha_2^2 + z$ where x and z do not involve α_1 or α_2. Then $x\alpha_2^2 + z = 0$ implies $x = 0$ and we have a relation involving neither β_1 nor β_2. Continuing inductively, we eventually deduce that all coefficients b_I in the original relation $\sum_I b_I\beta_I = 0$ must be zero.

Observe now that $\beta_i^2 = \beta_{2i}$ if $2i < n$ and $\beta_i^2 = 0$ if $2i \geq n$, since $\left(\sum_j \alpha_j^i\right)^2 = \sum_j \alpha_j^{2i}$. This implies that $H^*(SO(n);\mathbb{Z}_2)$ is generated as an algebra by the β_i's with i odd, and in fact that $H^*(SO(n);\mathbb{Z}_2)$ is the quotient of $\mathbb{Z}_2[\beta_1, \beta_2, \cdots]$ by the relations $\beta_i^2 = \beta_{2i}$ and $\beta_j = 0$ for $j \geq n$. Writing $j = ip_i$ with i odd and p_i a power of 2, the relation $\beta_j = 0$ becomes $\beta_{ip_i} = \beta_i^{p_i} = 0$. For a given i, this relation holds iff $ip_i \geq n$. This proves (a).

For part (b), note first that the group multiplication $SO(n) \times SO(n) \to SO(n)$ is cellular in view of the inclusions $P^iP^i \subset P^iP^{i-1}$ and $P^iP^j \subset P^jP^i$ for $i < j$. So we can compute Pontryagin products at the cellular level. We know that there is at least an additive isomorphism $H_*(SO(n);\mathbb{Z}_2) \approx \Lambda_{\mathbb{Z}_2}[e^1, \cdots, e^{n-1}]$ since the products $e^{i_1} \cdots e^{i_m}$ with $n > i_1 > \cdots > i_m > 0$, $m \geq 0$, form a basis for $H_*(SO(n);\mathbb{Z}_2)$. The inclusion $P^iP^i \subset P^iP^{i-1}$ then implies that the Pontryagin product $(e^i)^2$ is 0. It remains only to see the commutativity relation $e^ie^j = e^je^i$. The inclusion $P^iP^j \subset P^jP^i$ for $i < j$ was obtained from the formula $\rho(v)\rho(w) = \rho(\rho(v)w)\rho(v')$ for $v \in \mathbb{R}^{i+1}$ and $w \in \mathbb{R}^{j+1}$. The map $f : P^i \times P^j \to P^j \times P^i$, $f(v, w) = (\rho(v)w, v')$, is a homeomorphism since it is the composition of homeomorphisms $(v, w) \mapsto (v, \rho(v)w) \mapsto (v', \rho(v)w) \mapsto (\rho(v)w, v')$. Also, f obviously takes $(i+j-1)$-skeleton to $(i+j-1)$-skeleton, hence $e^i \times e^j$ homeomorphically to $e^j \times e^i$. Therefore $e^ie^j = e^je^i$ in $H_*(SO(n);\mathbb{Z}_2)$. □

The cup product and Pontryagin product structures in this theorem may seem at first glance to be unrelated, but in fact the relationship is fairly direct. As we saw in the previous section, the dual of a polynomial algebra $\mathbb{Z}_2[x]$ is a divided polynomial algebra $\Gamma_{\mathbb{Z}_2}[\alpha]$, and with $\mathbb{Z}_2$ coefficients the latter is an exterior algebra $\Lambda_{\mathbb{Z}_2}[\alpha_0, \alpha_1, \cdots]$ where $|\alpha_i| = 2^i|x|$. If we truncate the polynomial algebra by a relation $x^{2^n} = 0$, then this just eliminates the generators α_i for $i \geq n$. In view of this, if it were the case that the generators β_i for the algebra $H^*(SO(n);\mathbb{Z}_2)$ happened to be primitive, then $H^*(SO(n);\mathbb{Z}_2)$ would be isomorphic as a Hopf algebra to the tensor product of the single-generator Hopf algebras $\mathbb{Z}_2[\beta_i]/(\beta_i^{p_i})$, $i = 1, 3, \cdots$, hence the dual Hopf algebra $H_*(SO(n);\mathbb{Z}_2)$ would be the tensor product of the corresponding truncated divided polynomial algebras, in other words an exterior algebra as we explained above. This is in fact the structure of $H_*(SO(n);\mathbb{Z}_2)$, so since the Pontryagin product in $H_*(SO(n);\mathbb{Z}_2)$ determines the coproduct in $H^*(SO(n);\mathbb{Z}_2)$ uniquely, it follows that the β_i's must indeed be primitive.

It is not difficult to give a direct argument that each β_i is primitive. The coproduct $\Delta : H^*(SO(n);\mathbb{Z}_2) \to H^*(SO(n);\mathbb{Z}_2) \otimes H^*(SO(n);\mathbb{Z}_2)$ is induced by the group multiplication $\mu : SO(n) \times SO(n) \to SO(n)$. We need to show that the value of $\Delta(\beta_i)$ on $e^I \otimes e^J$, which we denote $\langle \Delta(\beta_i), e^I \otimes e^J \rangle$, is the same as the value $\langle \beta_i \otimes 1 + 1 \otimes \beta_i, e^I \otimes e^J \rangle$ for all cells e^I and e^J whose dimensions add up to i. Since $\Delta = \mu^*$, we have $\langle \Delta(\beta_i), e^I \otimes e^J \rangle = \langle \beta_i, \mu_*(e^I \otimes e^J) \rangle$. Because μ is the multiplication map, $\mu(e^I \times e^J)$ is contained in $P^I P^J$, and if we use the relations $P^j P^j \subset P^j P^{j-1}$ and $P^j P^k \subset P^k P^j$ for $k > j$ to rearrange the factors P^j of $P^I P^J$ so that their dimensions are in decreasing order, then the only way we will end up with a term P^i is if we start with $P^I P^J$ equal to $P^i P^0$ or $P^0 P^i$. Thus $\langle \beta_i, \mu_*(e^I \otimes e^J) \rangle = 0$ unless $e^I \otimes e^J$ equals $e^i \otimes e^0$ or $e^0 \otimes e^i$. Hence $\Delta(\beta_i)$ contains no other terms besides $\beta_i \otimes 1 + 1 \otimes \beta_i$, and β_i is primitive.

Integer Homology and Cohomology

Since there are homeomorphisms $SO(3) \approx \mathbb{R}\mathrm{P}^3$ and $SO(4) \approx S^3 \times \mathbb{R}\mathrm{P}^3$, we see that the homology and cohomology of $SO(n)$ with $\mathbb{Z}$ coefficients are not going to be completely simple. The additive structures however turn out to be not a whole lot more complicated than in these two cases. In particular one has a qualitative statement:

Proposition 3.46. *$H_*(SO(n);\mathbb{Z})$ is a direct sum of $\mathbb{Z}$'s and $\mathbb{Z}_2$'s.*

Proof: We compute the cellular chain complex of $SO(n)$, showing that it splits as a tensor product of simpler complexes. For a cell $e^i \subset P^{n-1} \subset SO(n)$ the cellular boundary de^i is $2e^{i-1}$ for even $i > 0$ and 0 for odd i. To compute the cellular boundary of a cell $e^{i_1} \cdots e^{i_m}$ we can pull it back to a cell $e^{i_1} \times \cdots \times e^{i_m}$ of $P^{n-1} \times \cdots \times P^1$ whose cellular boundary, by Proposition 3.35, is $\sum_j (-1)^{\sigma_j} e^{i_1} \times \cdots \times de^{i_j} \times \cdots \times e^{i_m}$ where $\sigma_j = i_1 + \cdots + i_{j-1}$. Hence $d(e^{i_1} \cdots e^{i_m}) = \sum_j (-1)^{\sigma_j} e^{i_1} \cdots de^{i_j} \cdots e^{i_m}$ where it is understood that $e^{i_1} \cdots de^{i_j} \cdots e^{i_m}$ is zero if $i_j - 1 = i_{j-1}$ since $P^{i_j - 1} P^{i_j - 1} \subset P^{i_j - 1} P^{i_j - 2}$, in a lower-dimensional skeleton.

Now to split the cellular chain complex $C_*(SO(n))$ as a tensor product of smaller chain complexes, let C^{2i} be the subcomplex of $C_*(SO(n))$ with basis the cells e^0, e^{2i}, e^{2i-1}, and $e^{2i}e^{2i-1}$. This is a subcomplex since $de^{2i-1} = 0$, $de^{2i} = 2e^{2i-1}$, and, in $P^{2i} \times P^{2i-1}$, $d(e^{2i} \times e^{2i-1}) = de^{2i} \times e^{2i-1} + e^{2i} \times de^{2i-1} = 2e^{2i-1} \times e^{2i-1}$, hence $d(e^{2i}e^{2i-1}) = 0$ since $P^{2i-1}P^{2i-1} \subset P^{2i-1}P^{2i-2}$. In view of the boundary formula for $C_*(SO(n))$ it is then clear that there are chain complex isomorphisms

$$C_*(SO(2k+1)) \approx C^2 \otimes C^4 \otimes \cdots \otimes C^{2k}$$
$$C_*(SO(2k+2)) \approx C^2 \otimes C^4 \otimes \cdots \otimes C^{2k} \otimes C^{2k+1}$$

where C^{2k+1} has basis e^0 and e^{2k+1}. Since $H_*(C^{2i})$ consists of $\mathbb{Z}$'s in dimensions 0 and $4i-1$ and a $\mathbb{Z}_2$ in dimension $2i-1$, while $H_*(C^{2k+1})$ consists of $\mathbb{Z}$'s in dimensions 0

and $2k+1$, we conclude from the algebraic Künneth formula that $H_*(SO(n))$ is a direct sum of $\mathbb{Z}$'s and $\mathbb{Z}_2$'s. □

Note that the calculation shows that $SO(2k)$ and $SO(2k-1)\times S^{2k-1}$ have isomorphic homology groups in all dimensions.

In view of the preceding proposition, one can get rather complete information about $H_*(SO(n);\mathbb{Z})$ by considering the natural maps to $H_*(SO(n);\mathbb{Z}_2)$ and to the quotient of $H_*(SO(n);\mathbb{Z})$ by its torsion subgroup. Let us denote this quotient by $H_*^{free}(SO(n);\mathbb{Z})$. The same strategy applies equally well to cohomology, and the universal coefficient theorem gives an isomorphism $H^*_{free}(SO(n);\mathbb{Z}) \approx H_*^{free}(SO(n);\mathbb{Z})$.

The proof of the proposition shows that the additive structure of $H_*^{free}(SO(n);\mathbb{Z})$ is fairly simple:

$$H_*^{free}(SO(2k+1)) \approx H_*(S^3 \times S^7 \times \cdots \times S^{4k-1})$$
$$H_*^{free}(SO(2k+2)) \approx H_*(S^3 \times S^7 \times \cdots \times S^{4k-1} \times S^{2k+1})$$

The multiplicative structure is also as simple as it could be:

Proposition 3.47. *The Pontryagin ring $H_*^{free}(SO(n);\mathbb{Z})$ is an exterior algebra,*

$$H_*^{free}(SO(2k+1);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[a_3, a_7, \cdots, a_{4k-1}]$$
$$H_*^{free}(SO(2k+2);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[a_3, a_7, \cdots, a_{4k-1}, a_{2k+1}]$$

where subscripts denote dimension. The generators a_i are primitive, so the dual Hopf algebra $H^*_{free}(SO(n);\mathbb{Z})$ is an exterior algebra on the dual generators α_i.

Proof: As in the case of $\mathbb{Z}_2$ coefficients we can work at the level of cellular chains since the multiplication in $SO(n)$ is cellular. Let E^i be the cycle $e^{2i}e^{2i-1}$ generating a $\mathbb{Z}$ summand of $H_*(SO(n);\mathbb{Z})$. As sets, we have $E^iE^j = e^{2i}e^{2i-1}e^{2j}e^{2j-1} = e^{2i}e^{2j}e^{2i-1}e^{2j-1}$ for $i \le j$ since the map $f : P^i \times P^j \to P^j \times P^i$ considered in the proof of Proposition 3.45 induces a homeomorphism $e^{2i-1}e^{2j} \approx e^{2j}e^{2i-1}$. So if $i = j$ we get $(E^i)^2 = 0$ in $H_*(SO(n);\mathbb{Z})$ since $P^{2i}P^{2i} \subset P^{2i}P^{2i-1}$. If $i < j$ we claim that $E^iE^j = -E^jE^i$, so the E^i's generate an exterior subalgebra $\Lambda_{\mathbb{Z}}[E^1, \cdots, E^k] \subset H_*(SO(n);\mathbb{Z})$ for $n = 2k+1$ or $2k+2$. The relation $E^iE^j = -E^jE^i$ follows from the fact that f induces a homeomorphism $e^ie^j \approx e^je^i$ of local degree $(-1)^{ij+1}$, namely, f is the composition $(v,w) \mapsto (v, \rho(v)w) \mapsto (v', \rho(v)w) \mapsto (\rho(v)w, v')$ with local degrees $+1, -1$, and $(-1)^{ij}$, so in commuting $e^{2i}e^{2i-1}e^{2j}e^{2j-1}$ into $e^{2j}e^{2j-1}e^{2i}e^{2i-1}$ the total number of sign changes is odd. Similar reasoning shows that when $n = 2k+2$, so that e^{2k+1} generates a $\mathbb{Z}$ summand of $H_*(SO(n);\mathbb{Z})$, we have an exterior subalgebra $\Lambda_{\mathbb{Z}}[E^1, \cdots, E^k, e^{2k+1}]$ in $H_*(SO(n);\mathbb{Z})$.

These exterior subalgebras account for all the nontorsion in $H_*(SO(n);\mathbb{Z})$, so the product structure in $H_*^{free}(SO(n);\mathbb{Z})$ is as stated in the proposition.

Now we show that the generators E^i and e^{2k+1} are primitive in $H_*^{free}(SO(n);\mathbb{Z})$. Let $X_2 \subset SO(n)$ be the union of the cells e^i and e^ie^j. This is a subcomplex since it is the union of the subcomplexes P^iP^j. The cellular chain complex of X_2 is the direct sum of the subcomplexes $C(0)$, $C(1)$, and $C(2)$ with bases $\{e^0\}$, $\{e^i \mid i>0\}$, and $\{e^ie^j \mid i>j>0\}$, respectively. Mod torsion, the homology of $C(1)$ and $C(2)$ is concentrated in odd dimensions, hence $H_*^{free}(X_2;\mathbb{Z})$ is also concentrated in odd dimensions, apart from a $\mathbb{Z}$ in dimension zero. This forces all positive-dimensional elements of $H_*^{free}(X_2;\mathbb{Z})$ to be primitive with respect to the diagonal map $X_2 \to X_2 \times X_2$. By naturality of the diagonal map, the generators E^i must then be primitive in $H_*^{free}(SO(n);\mathbb{Z})$, as well as the generator e^{2k+1} when $n = 2k+2$.

Since the exterior algebra generators of $H_*^{free}(SO(n);\mathbb{Z})$ are primitive, this algebra splits, as a Hopf algebra, into a tensor product of single-generator exterior algebras $\Lambda_{\mathbb{Z}}[a_i]$. The dual Hopf algebra $H^*_{free}(SO(n);\mathbb{Z})$ therefore splits as the tensor product of the dual exterior algebras $\Lambda_{\mathbb{Z}}[\alpha_i]$, hence $H^*_{free}(SO(n);\mathbb{Z})$ is also an exterior algebra. □

The exact ring structure of $H^*(SO(n);\mathbb{Z})$ can be deduced from these results via Bockstein homomorphisms, as we shall see in Example 4 in the next section, though the process is somewhat laborious and the answer not very neat.

Stiefel Manifolds

Consider the **Stiefel manifold** $V_{n,k}$, whose points are the k-tuples of orthonormal vectors in $\mathbb{R}^n$, i.e., **orthonormal k-frames**. Thus $V_{n,k}$ is a subset of the product of k copies of S^{n-1}, and it is given the subspace topology. As special cases, $V_{n,n} = O(n)$ and $V_{n,1} = S^{n-1}$. Also, $V_{n,2}$ can be identified with the space of unit tangent vectors to S^{n-1} since a vector v at the point $x \in S^{n-1}$ is tangent to S^{n-1} iff it is orthogonal to x.

There is a natural projection $p: O(n) \to V_{n,k}$ sending $\alpha \in O(n)$ to the k-frame consisting of the first k columns of α. This projection is onto, and the pre-images of points are precisely the cosets $\alpha O(n-k)$ where we embed $O(n-k)$ in $O(n)$ as the orthogonal transformation of the last $n-k$ coordinates of $\mathbb{R}^n$. Thus $V_{n,k}$ can be viewed as the space $O(n)/O(n-k)$ of such cosets, with the quotient topology from $O(n)$. This is the same as the previously-defined topology on $V_{n,k}$ since the projection $O(n) \to V_{n,k}$ is a surjection of compact Hausdorff spaces.

When $k < n$ the projection $p: SO(n) \to V_{n,k}$ is surjective, and we can use this to induce a CW structure on $V_{n,k}$ with a 0-cell $e^0 = p(SO(n-k))$ and cells $e^{i_1,\cdots,i_m} =$

$p(e^{i_1} \cdots e^{i_m})$ for $n > i_1 > \cdots > i_m \geq n-k$. It is an easy exercise, left for the reader, to check that these subsets of $V_{n,k}$ do indeed form the cells of a CW structure. Further, since p is a cellular map, the structure of the cellular chain complex of $V_{n,k}$ can easily be deduced from that of $SO(n)$. For example, the cellular chain complex of $V_{2k+1,2}$ is just the complex C^{2k} defined earlier, while for $V_{2k,2}$ the cellular boundary maps are all trivial, hence:

$$H_i(V_{2k+1,2};\mathbb{Z}) = \begin{cases} \mathbb{Z} & \text{for } i = 0,\ 4k-1 \\ \mathbb{Z}_2 & \text{for } i = 2k-1 \\ 0 & \text{otherwise} \end{cases}$$

$$H_i(V_{2k,2};\mathbb{Z}) = \begin{cases} \mathbb{Z} & \text{for } i = 0,\ 2k-2,\ 2k-1,\ 4k-3 \\ 0 & \text{otherwise} \end{cases}$$

e. Bockstein Homomorphisms

Even though homology or cohomology with $\mathbb{Z}$ coefficients may be the object of primary interest, it often turns out that cohomology with coefficients in a field, particularly $\mathbb{Z}_p$ with p prime, has more structure and is easier to compute. In this section we shall see how, in favorable cases at least, information about $\mathbb{Z}_p$ cohomology can be used to determine $\mathbb{Z}$ cohomology. We shall concentrate on cohomology since it generally has more structure than homology, but the basic constructions work equally well for homology.

If we take a short exact sequence $0 \to G \to H \to K \to 0$ of abelian groups and apply the covariant functor $Hom(C_n(X), -)$, we obtain

$$0 \longrightarrow C^n(X;G) \longrightarrow C^n(X;H) \longrightarrow C^n(X;K) \longrightarrow 0$$

which is exact since $C_n(X)$ is free. Letting n vary, we have a short exact sequence of chain complexes, so there is an associated long exact sequence of homology groups

$$\cdots \longrightarrow H^n(X;G) \longrightarrow H^n(X;H) \longrightarrow H^n(X;K) \longrightarrow H^{n+1}(X;G) \longrightarrow \cdots$$

whose 'boundary' map $H^n(X;K) \to H^{n+1}(X;G)$ is called a **Bockstein homomorphism**.

Of special interest are the Bockstein homomorphisms $\widetilde{\beta}: H^n(X;\mathbb{Z}_p) \to H^{n+1}(X;\mathbb{Z})$ and $\beta: H^n(X;\mathbb{Z}_p) \to H^{n+1}(X;\mathbb{Z}_p)$ associated to the two coefficient sequences

$$0 \longrightarrow \mathbb{Z} \xrightarrow{p} \mathbb{Z} \longrightarrow \mathbb{Z}_p \longrightarrow 0$$
$$0 \longrightarrow \mathbb{Z}_p \xrightarrow{p} \mathbb{Z}_{p^2} \longrightarrow \mathbb{Z}_p \longrightarrow 0$$

Usually p is taken to be a prime. These two Bockstein homomorphisms are related by $\beta = \rho\widetilde{\beta}$ where $\rho: H^*(X;\mathbb{Z}) \to H^*(X;\mathbb{Z}_p)$ is the homomorphism induced by the projection $\mathbb{Z} \to \mathbb{Z}_p$, reducing coefficients $\mod p$. Hence we have inclusions

$$Im\,\beta \subset Im\,\rho \subset Ker\,\beta$$

where the latter inclusion holds since $\widetilde{\beta}\rho = 0$ in the long exact sequence containing $\widetilde{\beta}$. In view of the inclusion $Im\,\beta \subset Ker\beta$, we can form the associated **Bockstein homology groups** $Ker\beta/Im\,\beta$, which we denote by $\beta^n(X)$ in dimension n, letting the prime p be implicit.

In the case that $H_*(X;\mathbb{Z})$ is finitely generated in each dimension, we shall analyze the relation between $\beta^*(X)$ and $H^*(X;\mathbb{Z})$ via the algebraic notion of **minimal chain complexes**. Suppose that $C = \{C_n, \partial\}$ is a chain complex whose homology groups

$H_n(C)$ are finitely generated for each n. Choose a splitting of $H_n(C)$ into a product of cyclic groups. To each cyclic factor of $H_n(C)$ with generator g_i define a corresponding chain complex $M(g_i)$ by the prescription: If g_i is of infinite order, let $M(g_i)$ consist of just a $\mathbb{Z}$ in dimension n, with generator z_i, while if g_i has order k, let $M(g_i)$ consist of $\mathbb{Z}$'s in dimensions n and $n+1$, generated by x_i and y_i respectively, with $\partial y_i = kx_i$. Let M be the direct sum of the chain complexes $M(g_i)$. Define a chain map $\sigma : M \to C$ by sending z_i and x_i to cycles ζ_i and ξ_i representing the corresponding homology classes g_i, and y_i to a chain η_i with $\partial\eta_i = k\xi_i$. The chain map σ induces an isomorphism on homology, hence also on cohomology with any coefficients by Corollary 3.4. The dual cochain complex M^* splits as the direct sum of the dual complexes $M^*(g_i)$. So in cohomology with $\mathbb{Z}$ coefficients the dual basis element z_i^* generates a $\mathbb{Z}$ summand in dimension n, while y_i^* generates a $\mathbb{Z}_k$ summand in dimension $n+1$ since $\delta x_i^* = ky_i^*$. One can also see from this what happens with other choices of coefficients. For example, with $\mathbb{Z}_p$ coefficients, p prime, z_i^* gives a $\mathbb{Z}_p$ summand of $H^n(C;\mathbb{Z}_p)$, while x_i^* and y_i^* give $\mathbb{Z}_p$ summands of $H^n(C;\mathbb{Z}_p)$ and $H^{n+1}(C;\mathbb{Z}_p)$ if p divides k and otherwise they give nothing.

The chain map σ induces an isomorphism between the associated Bockstein long exact sequences of cohomology groups, so we can use M^* to compute $\widetilde{\beta}$ and β. We have $\widetilde{\beta}(y_i^*) = 0 = \widetilde{\beta}(z_i^*)$ and $\widetilde{\beta}(x_i^*) = p^{k-1}y_i^*$. Hence $\beta(y_i^*) = 0 = \beta(z_i^*)$, while $\beta(x_i^*) = y_i^*$ if $k = 1$ and $\beta(x_i^*) = 0$ if $k > 1$. The following is then obvious:

Proposition 3.48.

(a) *Each $\mathbb{Z}$ summand of $H^n(X;\mathbb{Z})$ contributes a $\mathbb{Z}_p$ summand to $\beta^n(X)$.*

(b) *Each $\mathbb{Z}_{p^k}$ summand of $H^n(X;\mathbb{Z})$ with $k > 1$ contributes $\mathbb{Z}_p$ summands to $\beta^{n-1}(X)$ and $\beta^n(X)$.*

(c) *A $\mathbb{Z}_p$ summand of $H^n(X;\mathbb{Z})$ produces $\mathbb{Z}_p$ summands of $H^{n-1}(X;\mathbb{Z}_p)$ and $H^n(X;\mathbb{Z}_p)$ with β an isomorphism between these two summands, hence there is no contribution to $\beta^*(X)$.*

(d) *The image of β in $H^n(X;\mathbb{Z}_p)$ is the injective image under $\rho : H^n(X;\mathbb{Z}) \to H^n(X;\mathbb{Z}_p)$ of the subgroup of $H^n(X;\mathbb{Z})$ consisting of 0 and the elements which generate $\mathbb{Z}_p$ summands. This is the full p-torsion subgroup of $H^n(X;\mathbb{Z})$ iff $H^n(X;\mathbb{Z})$ contains no elements of order p^2.* □

We will work out four examples, exploiting the fact that β satisfies the derivation property

$$(*) \qquad \beta(a \smile b) = \beta(a) \smile b + (-1)^{|a|} a \smile \beta(b)$$

which comes from the corresponding formula for ordinary coboundary. Namely, let φ and ψ be $\mathbb{Z}_p$-cocycles representing a and b, and let $\widetilde{\varphi}$ and $\widetilde{\psi}$ be lifts of these to $\mathbb{Z}_{p^2}$-cochains. Concretely, one can view φ and ψ as functions on singular simplices with values in $\{0, 1, \cdots, p-1\}$, and then $\widetilde{\varphi}$ and $\widetilde{\psi}$ can be taken to be the same functions, but with $\{0, 1, \cdots, p-1\}$ regarded as a subset of $\mathbb{Z}_{p^2}$. Then $\delta\widetilde{\varphi} = p\eta$ and $\delta\widetilde{\psi} = p\mu$ for $\mathbb{Z}_p$-cocycles η and μ representing $\beta(a)$ and $\beta(b)$. Taking cup products, $\widetilde{\varphi} \smile \widetilde{\psi}$ is a $\mathbb{Z}_{p^2}$-cochain lifting the $\mathbb{Z}_p$-cocycle $\varphi \smile \psi$, and

$$\delta(\widetilde{\varphi} \smile \widetilde{\psi}) = \delta\widetilde{\varphi} \smile \widetilde{\psi} \pm \widetilde{\varphi} \smile \delta\widetilde{\psi} = p\eta \smile \widetilde{\psi} \pm \widetilde{\varphi} \smile p\mu = p(\eta \smile \psi) \pm p(\varphi \smile \mu)$$

where the sign $\pm$ is $(-1)^{|a|}$. Hence $\eta \smile \psi + (-1)^{|a|}\varphi \smile \mu$ represents $\beta(a \smile b)$, giving the formula $(*)$. Note that p did not have to be a prime for this argument.

Example 1: Lens Spaces. The cup product structure for lens spaces was computed in §3.3 using Poincaré duality, but using Bocksteins we can deduce it from the cup product structure in $\mathbb{C}\mathrm{P}^\infty$. Consider first the infinite-dimensional lens space $L = S^\infty/\mathbb{Z}_m$ where $\mathbb{Z}_m$ acts on S^∞, the unit sphere in $\mathbb{C}^\infty$, by scalar multiplication. Thus the action is generated by the rotation $v \mapsto e^{2\pi i/m}v$. The quotient map $S^\infty \to \mathbb{C}\mathrm{P}^\infty$ factors through L, and in particular we have a projection $L \to \mathbb{C}\mathrm{P}^\infty$. Looking at the cell structure on L described in §2.2, we see that each even-dimensional cell of L projects homeomorphically onto the corresponding cell of $\mathbb{C}\mathrm{P}^\infty$. Namely, the $2n$-cell of L is the image of points in $S^{2n+1} \subset \mathbb{C}^{n+1}$ of the form $\cos\theta(z_1, \cdots, z_n, 0) + \sin\theta(0, \cdots, 0, 1)$ with $\sum z_i^2 = 1$ and $0 < \theta \leq \pi$, and the same holds for the $2n$-cell of $\mathbb{C}\mathrm{P}^\infty$. From cellular cohomology it then follows that the map $L \to \mathbb{C}\mathrm{P}^\infty$ induces isomorphisms on even-dimensional cohomology with $\mathbb{Z}_m$ coefficients. Since $H^*(\mathbb{C}\mathrm{P}^\infty; \mathbb{Z}_m)$ is a polynomial ring, we deduce that if $y \in H^2(L; \mathbb{Z}_m)$ is a generator, then y^k generates $H^{2k}(L; \mathbb{Z}_m)$ for all k.

From the additive structure of $H_*(L; \mathbb{Z})$ it follows, via the analysis leading up to the previous proposition, that the Bockstein β for $\mathbb{Z}_m$ coefficients is zero on even-dimensional classes in $H^*(L; \mathbb{Z}_m)$ and takes a generator in dimension $2k-1$ to a generator in dimension $2k$. Thus there is a generator $x \in H^1(L; \mathbb{Z}_m)$ with $\beta(x) = y$. The product formula $(*)$ holds for arbitrary m, so $\beta(xy^k) = \beta(x)y^k - x\beta(y^k) = y^{k+1}$. Thus β takes xy^k to a generator, hence xy^k must be a generator of $H^{2k+1}(L; \mathbb{Z}_m)$.

This completely determines the cup product structure in $H^*(L; \mathbb{Z}_m)$ if m is odd since the commutativity property of cup product implies that $x^2 = 0$ in this case. Thus $H^*(L; Z_m) \approx \Lambda_{\mathbb{Z}_m}[x] \otimes \mathbb{Z}_m[y]$ for odd m. When m is even this statement needs to be modified slightly by inserting the relation that x^2 is the unique element of order 2 in $H^2(L; \mathbb{Z}_m) \approx \mathbb{Z}_m$, as we showed in §3.2 by an explicit calculation in the 2-skeleton of L.

The cup product structure in finite-dimensional lens spaces follows from this since a finite-dimensional lens space embeds as a skeleton in an infinite-dimensional lens space, and the homotopy type of an infinite-dimensional lens space is determined by its fundamental group since it is a $K(\pi,1)$. It follows that the cup product structure on a lens space $S^{2n+1}/\mathbb{Z}_m$ is obtained from the preceding calculation by truncating via the relation $y^{n+1}=0$.

Example 2. We compute $H^*(\mathbb{R}\mathrm{P}^\infty\times\mathbb{R}\mathrm{P}^\infty;\mathbb{Z})$. In principle this can be done from the Künneth formula and the universal coefficient theorem, but we will obtain a more concise answer by using Bocksteins instead of the Künneth formula.

From cellular cohomology we know that $\widetilde{H}^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_p)=0$ for p an odd prime. Theorem 3.12 then implies that $\widetilde{H}^*(\mathbb{R}\mathrm{P}^\infty\times\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_p)=0$, so by the universal coefficient theorem, $\widetilde{H}^*(\mathbb{R}\mathrm{P}^\infty\times\mathbb{R}\mathrm{P}^\infty;\mathbb{Z})$ consists entirely of elements of order a power of 2. The Bockstein β in $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)\approx\mathbb{Z}_2[x]$ is easy to compute. Since $\widetilde{H}_*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z})$ has $\mathbb{Z}_2$'s in odd dimensions and 0's in even dimensions, it follows from the discussion before the proposition that in $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$, $\beta(x^n)$ is x^{n+1} for n odd and 0 for n even. In $H^*(\mathbb{R}\mathrm{P}^\infty\times\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)\approx\mathbb{Z}_2[x,y]$ we have $\beta(x^my^n)=(\beta x^m)y^n+x^m(\beta y^n)$. Whether we get zero, one, or two terms here depends on the parities of m and n. The result can be expressed graphically by the following figure:

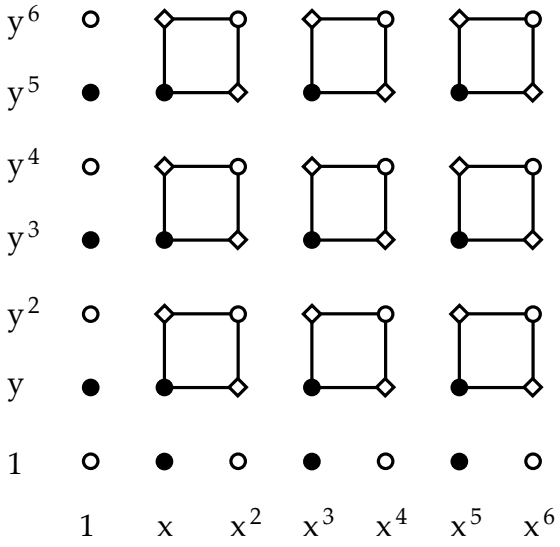

Here the dot, diamond, or circle in the (m,n) position represents the monomial x^my^n and line segments indicate nontrivial Bocksteins. For example, the lower left square records the formulas $\beta(xy)=x^2y+xy^2$, $\beta(x^2y)=x^2y^2=\beta(xy^2)$, and $\beta(x^2y^2)=0$. Thus in this square we see that $Ker\beta=Im\,\beta$, with generators the 'diagonal' sum x^2y+xy^2 and x^2y^2. The same thing happens in all the other squares, so it is apparent that $Ker\beta=Im\,\beta$

except for the zero-dimensional class '1.' Furthermore, $Ker\beta$ consists of the subring $\mathbb{Z}_2[x^2, y^2]$, indicated by the circles, together with the multiples of $x^2y + xy^2$ by elements of $\mathbb{Z}_2[x^2, y^2]$. It follows that there is a ring isomorphism

$$H^*(\mathbb{R}\mathrm{P}^\infty \times \mathbb{R}\mathrm{P}^\infty; \mathbb{Z}) \approx \mathbb{Z}[\lambda, \mu, \nu]/(2\lambda, 2\mu, 2\nu, \nu^2 + \lambda^2\mu + \lambda\mu^2)$$

where $\rho(\lambda) = x^2$, $\rho(\mu) = y^2$, $\rho(\nu) = x^2y + xy^2$, and the relation $\nu^2 + \lambda^2\mu + \lambda\mu^2 = 0$ holds since $(x^2y + xy^2)^2 = x^4y^2 + x^2y^4$.

This calculation is an instance of the general phenomenon that cup product structures with $\mathbb{Z}$ coefficients tend to be considerably messier than with field coefficients. If one computes $H^*(\mathbb{R}\mathrm{P}^\infty \times \mathbb{R}\mathrm{P}^\infty \times \mathbb{R}\mathrm{P}^\infty; \mathbb{Z})$ by the same technique one sees even more striking evidence of this.

Example 3. We use Bocksteins to produce finite CW complexes X_1, X_2, and Y such that the rings $H^*(X_1; \mathbb{Z})$ and $H^*(X_2; \mathbb{Z})$ are isomorphic but the rings $H^*(X_1 \times Y; \mathbb{Z})$ and $H^*(X_2 \times Y; \mathbb{Z})$ are not isomorphic. According to Theorem 3.12 this can happen only if all three of X_1, X_2, and Y have torsion in their $\mathbb{Z}$-cohomology. The space X_1 is obtained from $S^2 \times S^2$ by attaching a 3-cell e^3 to the second S^2 factor by a map of degree 2. Thus X_1 has a CW structure with cells e^0, e_1^2, e_2^2, e^3, e^4 with e^3 attached to the 2-sphere $e_0 \cup e_2^2$. The space X_2 is obtained from $S^2 \vee S^2 \vee S^4$ by attaching a 3-cell to the second S^2 summand by a map of degree 2, so it has a CW structure with the same collection of five cells, the only difference being that in X_2 the 4-cell is attached trivially. The space Y is a Moore space $M(\mathbb{Z}_2, 2)$, with three cells which we label f^0, f^2, f^3, the last cell being attached by a map of degree 2.

From cellular cohomology we see that $H^*(X_1) \approx H^*(X_2)$ consists of $\mathbb{Z}$'s in dimensions 0, 2, 4 and a $\mathbb{Z}_2$ in dimension 3. The cup product structure for both X_1 and X_2 is trivial since for dimension reasons the only possible nontrivial product is the square of the class e_1^2, but this is zero as one sees by restricting to the subcomplex $S^2 \times S^2$ or $S^2 \vee S^2 \vee S^4$. Also from cellular cohomology we have $H^*(Y)$ consisting of a $\mathbb{Z}$ in dimension 0 and a $\mathbb{Z}_2$ in dimension 3.

The restriction map $H^*(X_i \times Y; \mathbb{Z}) \to H^*(X_i; \mathbb{Z})$ is an isomorphism if 2-torsion is ignored. This can seen by applying 3.12 with coefficients in $\mathbb{Z}_p$ for p an odd prime. To determine the 2-torsion in $H^*(X_i \times Y; \mathbb{Z})$ we compute $\mathbb{Z}_2$ Bocksteins. With $\mathbb{Z}_2$ coefficients the cellular cochain complexes for X_i, Y, and $X_i \times Y$ are all trivial, so we can identify the cells with a basis for $\mathbb{Z}_2$-cohomology. In X_i and Y the only nontrivial Bocksteins are $\beta(e_2^2) = e^3$ and $\beta(f^2) = f^3$. The Bocksteins in $X_i \times Y$ can then be computed using the product formula for β. The results are shown in the following table, where an arrow

denotes a nontrivial Bockstein.

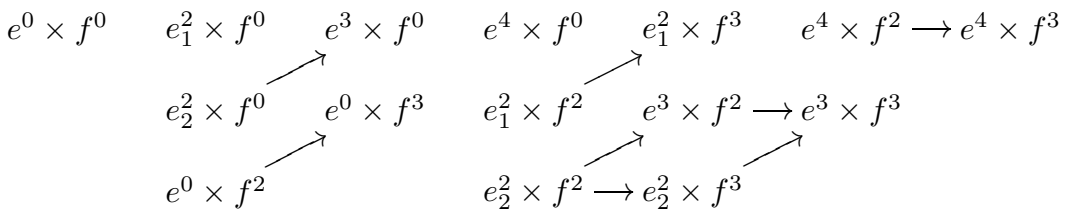

The two arrows from $e_2^2 \times f^2$ mean that $\beta(e_2^2 \times f^2) = e^3 \times f^2 + e_2^2 \times f^3$. The three $\mathbb{Z}$'s in $H^*(X_i \times Y; \mathbb{Z})$ appear as the three cells which are neither sources nor targets of Bocksteins. Apart from these three cells we have $Ker\beta = Im\,\beta$, so the torsion in $H^*(X_i \times Y; \mathbb{Z})$ consists entirely of $\mathbb{Z}_2$'s mapping isomorphically to $Im\,\beta$ under the coefficient homomorphism $H^*(X_i \times Y; \mathbb{Z}) \to H^*(X_i \times Y; \mathbb{Z}_2)$. Thus $H^*(X_i \times Y; \mathbb{Z})$ consists of $\mathbb{Z}$'s in dimensions 0, 2, and 4, $\mathbb{Z}_2 \times \mathbb{Z}_2$'s in dimensions 3 and 5, and $\mathbb{Z}_2$'s in dimensions 6 and 7. In particular, there is a $\mathbb{Z}_2$ corresponding to $e^3 \times f^2 + e_2^2 \times f^3$ in dimension 5.

We can observe now that with $\mathbb{Z}$ coefficients, the cup product $H^2 \times H^5 \to H^7$ is nontrivial for $X_1 \times Y$ but trivial for $X_2 \times Y$. For, in $H^*(X_i \times Y; \mathbb{Z}_2)$ we have, using the relation $(a \times b) \smile (c \times d) = (a \smile c) \times (b \smile d)$ which follows immediately from the definition of cross product,

$$e_1^2 \times f^0 \smile e_1^2 \times f^3 = (e_1^2 \smile e_1^2) \times (f^0 \smile f^3) = 0 \qquad \text{since} \quad e_1^2 \smile e_1^2 = 0$$

$$\begin{aligned} e_1^2 \times f^0 \smile (e^3 \times f^2 + e_2^2 \times f^3) &= (e_1^2 \smile e^3) \times (f^0 \smile f^2) + (e_1^2 \smile e_2^2) \times (f^0 \smile f^3) \\ &= (e_1^2 \smile e_2^2) \times f^3 \qquad \text{since} \quad e_1^2 \smile e^3 = 0 \end{aligned}$$

and in $H^7(X_i \times Y; \mathbb{Z}_2) \approx H^7(X_i \times Y; \mathbb{Z})$, $(e_1^2 \smile e_2^2) \times f^3 = e^4 \times f^3 \neq 0$ for $i = 1$ but $(e_1^2 \smile e_2^2) \times f^3 = 0 \times f^3 = 0$ for $i = 2$.

Thus the cohomology ring of a product space is not always determined by the cohomology rings of the factors.

Example 4. Bockstein homomorphisms can be used to get a more complete picture of the structure of $H^*(SO(n); \mathbb{Z})$ than we obtained in the preceding section. Continuing the notation employed there, we know from the calculations above that $\beta\big(\sum_j \alpha_j^{2i-1}\big) = \sum_j \alpha_j^{2i}$ and $\beta\big(\sum_j \alpha_j^{2i}\big) = 0$, hence $\beta(\beta_{2i-1}) = \beta_{2i}$ and $\beta(\beta_{2i}) = 0$. Taking the case $n = 5$ as an example, where $H^*(SO(5); \mathbb{Z}_2) \approx \mathbb{Z}_2[\beta_1, \beta_3]/(\beta_1^8, \beta_3^2)$, the table below shows the nontrivial Bocksteins. Here two arrows from an element mean 'sum,' e.g., $\beta(\beta_1\beta_3) = \beta(\beta_1)\beta_3 + \beta_1\beta(\beta_3) = \beta_2\beta_3 + \beta_1\beta_4 = \beta_1^2\beta_3 + \beta_1^5$. By Proposition 3.48, the image of $\rho : H^*(SO(5); \mathbb{Z}) \to H^*(SO(5); \mathbb{Z}_2)$ is the kernel of β, and the image of β is the image of the torsion in

$H^*(SO(5);\mathbb{Z})$ under the coefficient map $H^*(SO(5);\mathbb{Z}) \to H^*(SO(5);\mathbb{Z}_2)$. This allows the calculation of the rest of the table, giving the groups $H^i(SO(5);\mathbb{Z})$.

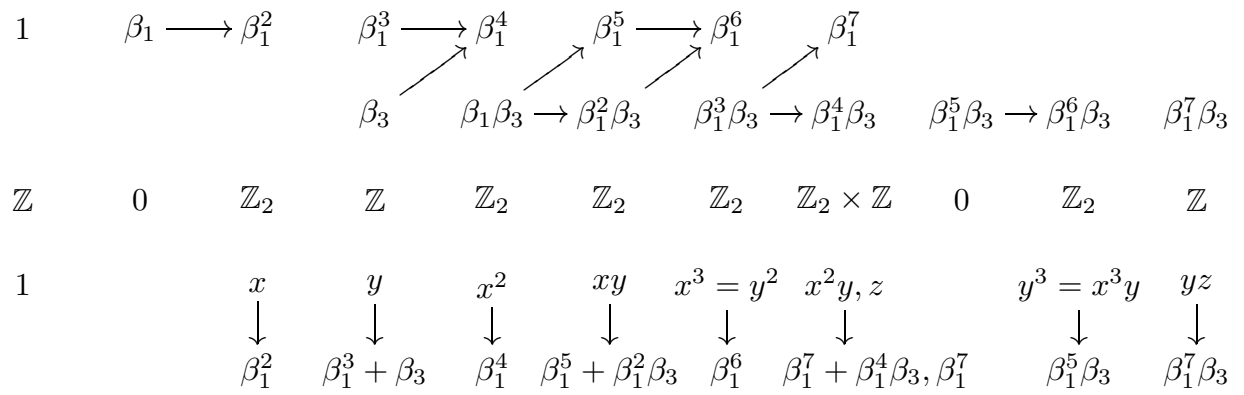

It is interesting that the generator $y \in H^3(SO(5);\mathbb{Z}) \approx \mathbb{Z}$ has y^2 nontrivial, since this implies that the ring structures of $H^*(SO(5);\mathbb{Z})$ and $H^*(\mathbb{R}\mathrm{P}^7 \times S^3;\mathbb{Z})$ are not isomorphic, even though the cohomology groups and the $\mathbb{Z}_2$ cohomology rings of these two spaces are the same. An exercise at the end of the section is to show that in fact $SO(5)$ is not homotopy equivalent to the product of any two CW complexes with nontrivial cohomology.

A natural way to describe $H^*(SO(5);\mathbb{Z})$ would be as a quotient of a free graded-commutative associative algebra $F[x,y,z]$ over $\mathbb{Z}$ with generators x, y, z in dimensions 2, 3, 7. Elements of $F[x,y,z]$ are representable as polynomials $p(x,y,z)$, subject only to the relations imposed by graded-commutativity. In particular, since y and z are odd-dimensional, $yz = -zy$, and y^2 and z^2 are nonzero elements of order 2 in $F[x,y,z]$. Any monomial containing y^2 or z^2 as a factor also has order 2. In these terms, the calculation of $H^*(SO(5);\mathbb{Z})$ can be written:

$$H^*(SO(5);\mathbb{Z}) \approx F[x,y,z]/(2x, x^4, y^4, z^2, xz, x^3 - y^2)$$

The next figure shows the nontrivial Bocksteins for $H^*(SO(7);\mathbb{Z}_2)$. Here the numbers across the top indicate dimension, stopping with 21, the dimension of $SO(7)$. The labels on the dots refer to the basis of products of distinct β_i's, e.g., the dot labelled 135 is $\beta_1\beta_3\beta_5$. The left-right symmetry of the figure displays Poincaré duality quite graphically. Note also that the corresponding diagram for $SO(5)$, drawn in a slightly different way from the preceding figure, occurs in the upper left corner as the subdiagram with labels 1 through 4.

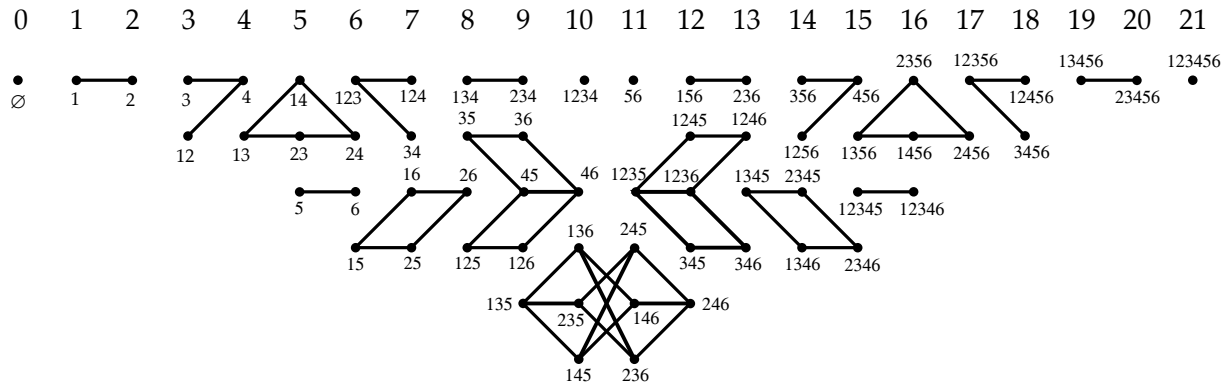

From the diagram one can with some effort work out the cup product structure in $H^*(SO(7);\mathbb{Z})$, but the answer is rather complicated, just as the diagram is:

$$F[x,y,z,v,w]/(2x, 2v, x^4, y^4, z^2, v^2, w^2, xz, vz, vw, y^2w, x^3y^2v, y^2z - x^3v, xw - y^2v - x^3v)$$

where x, y, z, v, w have dimensions 2,3,7,7,11, respectively. It is curious that the relation $x^3 = y^2$ in $H^*(SO(5);\mathbb{Z})$ no longer holds in $H^*(SO(7);\mathbb{Z})$.

Exercises

1. Show that $H^*(K(\mathbb{Z}_m,1);\mathbb{Z}_k)$ is isomorphic as a ring to $H^*(K(\mathbb{Z}_m,1);\mathbb{Z}_m)\otimes\mathbb{Z}_k$ if k divides m. In particular, if m/k is even this is $\Lambda_{\mathbb{Z}_k}[x]\otimes\mathbb{Z}_k[y]$.

2. In this problem we will derive one half of the classification of lens space up to homotopy equivalence, by showing that if $L_m(\ell_1,\cdots,\ell_n) \simeq L_m(\ell'_1,\cdots,\ell'_n)$ then $\ell_1\cdots\ell_n \equiv \pm\ell'_1\cdots\ell'_n k^n \mod m$ for some integer k. The converse is an exercise for §4.1.

(a) Let $L = L_m(\ell_1,\cdots,\ell_n)$ and let $\mathbb{Z}_m^*$ be the multiplicative group of invertible elements of $\mathbb{Z}_m$. Define $t \in \mathbb{Z}_m^*$ by the equation $xy^{n-1} = tz$ where x is a generator of $H^1(L;\mathbb{Z}_m)$, $y = \beta(x)$, and $z \in H^{2n-1}(L;\mathbb{Z}_m)$ is the image of a generator of $H^{2n-1}(L;\mathbb{Z})$. Show that the image $\tau(L)$ of t in the quotient group $\mathbb{Z}_m^*/\pm(\mathbb{Z}_m^*)^n$ depends only on the homotopy type of L.

(b) Given nonzero integers $k_1,\cdots,k_n$, define a map $\widetilde{f}: S^{2n-1} \to S^{2n-1}$ by $re^{i\theta} \mapsto re^{ik_j\theta}$ in the j^{th} coordinate of $\mathbb{C}^n$. Show:

(i) $\widetilde{f}$ has degree $k_1\cdots k_n$.

(ii) $\widetilde{f}$ induces a quotient map $f: L \to L'$ for $L' = L_m(\ell'_1,\cdots,\ell'_n)$ provided that $k_j\ell_j \equiv \ell'_j \mod m$ for each j.

(iii) f induces an isomorphism on π_1, hence on $H^1(-;\mathbb{Z}_m)$.

(iv) f has degree $k_1\cdots k_n$, i.e., induces multiplication by $k_1\cdots k_n$ on $H_{2n-1}(-;\mathbb{Z})$.

(c) Using the f in (b), show that $\tau(L) = k_1 \cdots k_n \tau(L')$.

(d) Deduce that if $L_m(\ell_1, \cdots, \ell_n) \simeq L_m(\ell'_1, \cdots, \ell'_n)$, then $\ell_1 \cdots \ell_n \equiv \pm \ell'_1 \cdots \ell'_n k^n \mod m$ for some integer k.

3. Let X be the smash product of k copies of a Moore space $M(\mathbb{Z}_p, n)$ with p prime. Compute the Bockstein homomorphisms in $H^*(X; \mathbb{Z}_p)$ and use this to describe $H^*(X; \mathbb{Z})$.

4. Using the cup product structure in $H^*(SO(5); \mathbb{Z})$, show that $SO(5)$ is not homotopy equivalent to the product of any two CW complexes with nontrivial cohomology.

5. Compute Bocksteins in $H_*(SO(5); \mathbb{Z}_2)$ and then determine the Pontryagin ring structure in $H_*(SO(5); \mathbb{Z})$.

f. Limits

It often happens that one has a CW complex X expressed as a union of an increasing sequence of subcomplexes $X_0 \subset X_1 \subset X_2 \subset \cdots$. For example, X_i could be the i-skeleton of X, or the X_i's could be finite complexes whose union is X. In situations of this sort, Proposition 3.17 says that $H_n(X;G)$ is the direct limit $\varinjlim H_n(X_i;G)$. Our goal in this section is to show this holds more generally for any homology theory, and to derive the corresponding formula for cohomology theories, which is a bit more complicated even for ordinary cohomology with $\mathbb{Z}$ coefficients. For ordinary homology and cohomology the results apply somewhat more generally than just to CW complexes, since if a space X is the union of an increasing sequence of subspaces X_i with the property that each compact set in X is contained in some X_i, then the singular complex of X is the union of the singular complexes of the X_i's, and so this gives a reduction to the CW case.

Let us begin by recalling that the direct limit $\varinjlim G_i$ of a sequence of homomorphisms of abelian groups

$$G_1 \xrightarrow{\alpha_1} G_2 \xrightarrow{\alpha_2} G_3 \longrightarrow \cdots$$

is defined to be the quotient of the direct sum $\bigoplus_i G_i$ by the subgroup of elements of the form $(g_1, g_2 - \alpha_1(g_1), g_3 - \alpha_2(g_2), \cdots)$. It is easy to see from this definition that every element of $\varinjlim G_i$ is represented by an element $g_i \in G_i$ for some i, and two such representatives $g_i \in G_i$ and $g_j \in G_j$ define the same element of $\varinjlim G_i$ iff they have the same image in some G_k under the appropriate composition of α_ℓ's. In case all the α_i's are injective and are viewed as inclusions of subgroups, then $\varinjlim G_i$ is just $\bigcup_i G_i$.

Example 1. For a prime p, consider the sequence $\mathbb{Z} \xrightarrow{p} \mathbb{Z} \xrightarrow{p} \mathbb{Z} \longrightarrow \cdots$ with all maps multiplication by p. Then $\varinjlim G_i$ can be identified with the subgroup $\mathbb{Z}[1/p]$ of $\mathbb{Q}$ consisting of rational numbers with denominator a power of p. More generally, we can realize any subgroup of $\mathbb{Q}$ as the direct limit of a sequence $\mathbb{Z} \longrightarrow \mathbb{Z} \longrightarrow \mathbb{Z} \longrightarrow \cdots$ with an appropriate choice of maps. For example, if the n^{th} map is multiplication by n, then the direct limit is $\mathbb{Q}$ itself.

Example 2. The sequence of injections $\mathbb{Z}_p \xrightarrow{p} \mathbb{Z}_{p^2} \xrightarrow{p} \mathbb{Z}_{p^3} \longrightarrow \cdots$, with p prime, has direct limit a group which we may call $\mathbb{Z}_{p^\infty}$. This is isomorphic to the quotient $\mathbb{Z}[1/p]/\mathbb{Z}$, a subgroup of $\mathbb{Q}/\mathbb{Z}$. More generally, we can realize any subgroup of $\mathbb{Q}/\mathbb{Z}$ as a direct limit of finite cyclic groups and injective maps between them. Note that $\mathbb{Q}/\mathbb{Z}$ is isomorphic to the direct sum of its subgroups $\mathbb{Z}[1/p]/\mathbb{Z} \approx \mathbb{Z}_{p^\infty}$ for all primes p.

We can realize these algebraic examples topologically by the following construction.

The **mapping telescope** of a sequence of maps

$$X_0 \xrightarrow{f_0} X_1 \xrightarrow{f_1} X_2 \longrightarrow \cdots$$

is the union of the mapping cylinders M_{f_i} with the copies of X_i in M_{f_i} and $M_{f_{i-1}}$ identified for all i. Thus the mapping telescope is the quotient space of the disjoint union $\coprod_i (X_i \times [i, i+1])$ in which $(x_i, i+1) \in X_i \times [i, i+1]$ is identified with $(f_i(x_i), i+1) \in X_{i+1} \times [i+1, i+2]$.

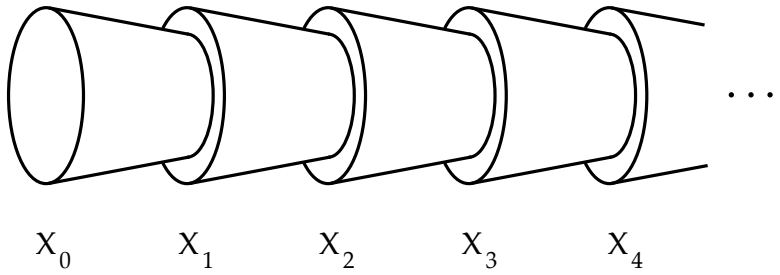

In the mapping telescope T, let T_i be the union of the first i mapping cylinders. This deformation retracts onto X_i by deformation retracting each mapping cylinder onto its right end in turn. If the maps f_i are cellular, each mapping cylinder is a CW complex and the telescope T is the increasing union of the subcomplexes $T_i \simeq X_i$. Then Proposition 3.17, or Theorem 3.50 below, implies that $H_n(T; G) \approx \varinjlim H_n(X_i; G)$.

Example 3. Suppose each f_i is a map $S^n \to S^n$ of degree p for a fixed prime p. Then $H_n(T)$ is the direct limit of the sequence $\mathbb{Z} \xrightarrow{p} \mathbb{Z} \xrightarrow{p} \mathbb{Z} \longrightarrow \cdots$ considered in Example 1 above, and $\widetilde{H}_k(T) = 0$ for $k \neq n$, so T is a Moore space $M(\mathbb{Z}[1/p], n)$.

Example 4. In the preceding example, if we attach a cell e^{n+1} to the first S^n in T via the identity map of S^n, we obtain a Moore space $M(\mathbb{Z}_{p^\infty}, n)$ since the space $X = T \cup e^{n+1}$ can be regarded as the direct limit of its subspaces $X_i = T_i \cup e^{n+1}$, which are $M(\mathbb{Z}_{p^i}, n)$'s, and the inclusion $X_i \subset X_{i+1}$ induces the inclusion $\mathbb{Z}_{p^i} \subset \mathbb{Z}_{p^{i+1}}$ on H_n.

Generalizing these two examples, if $f_i : S^n \to S^n$ has degree d_i then by suitable choice of the d_i's we may obtain Moore spaces $M(G, n)$ for any subgroup G of $\mathbb{Q}$ or $\mathbb{Q}/\mathbb{Z}$.

The behavior of cohomology groups is more complicated. If X is the increasing union of subcomplexes X_i, then the cohomology groups $H^n(X_i; G)$, for fixed n and G, form a sequence of homomorphisms

$$\cdots \longrightarrow G_2 \xrightarrow{\alpha_2} G_1 \xrightarrow{\alpha_1} G_0$$

Given such a sequence of group homomorphisms, the **inverse limit** $\varprojlim G_i$ is defined to be the subgroup of $\prod_i G_i$ consisting of sequences (g_i) with $\alpha_i(g_i) = g_{i-1}$ for all i. There

is a natural map $\lambda : H^n(X;G) \to \varprojlim H^n(X_i;G)$ sending an element of $H^n(X;G)$ to its sequence of images in $H^n(X_i;G)$ under the maps $H^n(X;G) \to H^n(X_i;G)$ induced by inclusion. One might hope that λ is an isomorphism, but this is not true in general, as we shall see. However, for some choices of G it is:

Proposition 3.49. *If the CW complex X is the union of an increasing sequence of subcomplexes X_i and if G is one of the fields $\mathbb{Q}$ or $\mathbb{Z}_p$, then $\lambda : H^n(X;G) \to \varprojlim H^n(X_i;G)$ is an isomorphism for all n.*

Proof: First we have an easy algebraic fact: Given a sequence $G_1 \xrightarrow{\alpha_1} G_2 \xrightarrow{\alpha_2} G_3 \longrightarrow \cdots$ of homomorphisms of abelian groups, then $Hom(\varinjlim G_i, G) = \varprojlim Hom(G_i, G)$ for any G. Namely, it follows from the definition of $\varinjlim G_i$ that a homomorphism $\varphi : \varinjlim G_i \to G$ is the same thing as a sequence of homomorphisms $\varphi_i : G_i \to G$ with $\varphi_i = \varphi_{i+1}\alpha_i$ for all i. Such a sequence (φ_i) is exactly an element of $\varprojlim Hom(G_i, G)$.

Now if G is a field $\mathbb{Q}$ or $\mathbb{Z}_p$ we have

$$\begin{aligned} H^n(X;G) &= Hom(H_n(X;G), G) \\ &= Hom(\varinjlim H_n(X_i;G), G) \\ &= \varprojlim Hom(H_n(X_i;G), G) \\ &= \varprojlim H^n(X_i;G) \end{aligned}$$

□

Let us analyze what happens for cohomology with an arbitrary coefficient group, or more generally for any cohomology theory. Given a sequence of homomorphisms

$$\cdots \longrightarrow G_2 \xrightarrow{\alpha_2} G_1 \xrightarrow{\alpha_1} G_0$$

define a 'difference' map $\delta : \prod_i G_i \to \prod_i G_i$ by $\delta(\cdots, g_i, \cdots) = (\cdots, g_i - \alpha_{i+1}(g_{i+1}), \cdots)$, so that $\varprojlim G_i$ is the kernel of δ. Denoting the cokernel of δ by $\varprojlim{}^1 G_i$, we have then an exact sequence

$$0 \longrightarrow \varprojlim G_i \longrightarrow \textstyle\prod_i G_i \xrightarrow{\delta} \prod_i G_i \longrightarrow \varprojlim{}^1 G_i \longrightarrow 0$$

This may be compared with the corresponding situation for the direct limit of a sequence $G_1 \xrightarrow{\alpha_1} G_2 \xrightarrow{\alpha_2} G_3 \longrightarrow \cdots$. Here one has a short exact sequence

$$0 \longrightarrow \textstyle\bigoplus_i G_i \xrightarrow{\delta} \bigoplus_i G_i \longrightarrow \varinjlim G_i \longrightarrow 0$$

where $\delta(\cdots, g_i, \cdots) = (\cdots, g_i - \alpha_{i-1}(g_{i-1}), \cdots)$, hence δ is injective. So there is no $\varinjlim{}^1 G_i$.

A few simple observations:

— If all the α_i's are isomorphisms then $\varprojlim G_i \approx G_1$ and $\varprojlim{}^1 G_i = 0$. More generally, $\varprojlim{}^1 G_i = 0$ if each α_i is surjective, for to realize a given element $(h_i) \in \prod_i G_i$ as $\delta(g_i)$ we can take $g_1 = 0$, then solve $\alpha_2(g_2) = -h_1$, $\alpha_3(g_3) = g_2 - h_2$, etc.

— If all the α_i's are zero then $\varprojlim G_i = \varprojlim{}^1 G_i = 0$.

— A finite number of terms can be deleted from the end of the sequence $\cdots \to G_2 \to G_1$ without affecting $\varprojlim G_i$ or $\varprojlim{}^1 G_i$. More generally, $\varprojlim G_i$ and $\varprojlim{}^1 G_i$ are unchanged if we replace the sequence $\cdots \to G_2 \to G_1$ by a subsequence, with the appropriate compositions of α_j's as the maps between the G_i's in the subsequence.

Example 5. For a prime p, consider the sequence of natural surjections $\cdots \to \mathbb{Z}_{p^3} \to \mathbb{Z}_{p^2} \to \mathbb{Z}_p$. The inverse limit of this sequence is a famous object in number theory, called the ***p*-adic integers**. Our notation for it will be $\widehat{\mathbb{Z}}_p$. It is actually a commutative ring, not just a group, since the projections $\mathbb{Z}_{p^{i+1}} \to \mathbb{Z}_{p^i}$ are ring homomorphisms, but we shall only be interested in the additive group structure. Elements of $\widehat{\mathbb{Z}}_p$ are infinite sequences $(\cdots, a_2, a_1)$ with $a_i \in \mathbb{Z}_{p^i}$ such that a_i is the $\mod p^i$ reduction of a_{i+1}. For each choice of a_i there are exactly p choices for a_{i+1}, so $\widehat{\mathbb{Z}}_p$ is uncountable. There is a natural inclusion $\mathbb{Z} \subset \widehat{\mathbb{Z}}_p$ as the constant sequences $a_i = n \in \mathbb{Z}$.

There is another way of looking at $\widehat{\mathbb{Z}}_p$. An element $(\cdots, a_2, a_1) \in \widehat{\mathbb{Z}}_p$ has a unique such expression with $0 \le a_i < p^i$ for all i. We can write each a_i uniquely in the form $b_{i-1}p^{i-1} + \cdots + b_1 p + b_0$ with $0 \le b_j < p$. The fact that a_{i+1} reduces $\mod p^i$ to a_i means that the numbers b_j depend only on the element $(\cdots, a_2, a_1) \in \widehat{\mathbb{Z}}_p$, so we can view the elements of $\widehat{\mathbb{Z}}_p$ as the 'base p infinite numbers' $\cdots b_1 b_0$ with $0 \le b_i < p$ for all i, with the familiar rule for addition in base p notation. From this viewpoint, the subgroup $\mathbb{Z} \subset \widehat{\mathbb{Z}}_p$ consists of the finite expressions $b_n \cdots b_1 b_0$. It is also clear from this representation of $\widehat{\mathbb{Z}}_p$ that $\widehat{\mathbb{Z}}_p$ is torsionfree.

Since the maps $\mathbb{Z}_{p^{i+1}} \to \mathbb{Z}_{p^i}$ are surjective, $\varprojlim{}^1 \mathbb{Z}_{p^i} = 0$. The next example shows how p-adic integers can also give rise to a nonvanishing $\varprojlim{}^1$ term.

Example 6. Consider the sequence $\cdots \longrightarrow \mathbb{Z} \xrightarrow{p} \mathbb{Z} \xrightarrow{p} \mathbb{Z}$ for p prime. In this case the inverse limit is zero since a given integer can only be divided by p finitely often. The $\varprojlim{}^1$ term is the cokernel of the map $\delta : \prod_\infty \mathbb{Z} \to \prod_\infty \mathbb{Z}$, $\delta(y_1, y_2, \cdots) = (y_1 - py_2, y_2 - py_3, \cdots)$. We claim that the map $\widehat{\mathbb{Z}}_p/\mathbb{Z} \to Coker(\delta)$ sending a p-adic number $\cdots b_1 b_0$ as above to $(b_0, b_1, \cdots)$ is an isomorphism. To see this, note first that the image of δ consists of the sums $y_1(1, 0, \cdots) + y_2(-p, 1, 0, \cdots) + y_3(0, -p, 1, 0, \cdots) + \cdots$. The terms after $y_1(1, 0, \cdots)$ give exactly the relations which hold among the p-adic numbers $\cdots b_1 b_0$, and in particular

allow one to reduce an arbitrary sequence $(b_0, b_1, \cdots)$ to a unique sequence with $0 \le b_i < p$ for all i. And the term $y_1(1, 0, \cdots)$ accounts for the subgroup $\mathbb{Z} \subset \widehat{\mathbb{Z}}_p$.

We come now to the main result of this section:

Theorem 3.50. *For a CW complex X which is the union of subcomplexes $X_0 \subset X_1 \subset \cdots$ there is an exact sequence*

$$0 \longrightarrow \varprojlim{}^1 h^{n-1}(X_i; G) \longrightarrow h^n(X; G) \xrightarrow{\lambda} \varprojlim h^n(X_i; G) \longrightarrow 0$$

where h^* is any reduced or unreduced cohomology theory. For any homology theory h_*, reduced or unreduced, the natural maps $\varinjlim h_n(X_i) \to h_n(X)$ are isomorphisms.

Proof: Let T be the mapping telescope of the sequence of inclusions $X_0 \hookrightarrow X_1 \hookrightarrow \cdots$. This is a subcomplex of $X \times [0, \infty)$ when $[0, \infty)$ is given the CW structure with the integer points as 0-cells. We have $T \simeq X$ since T is a deformation retract of $X \times [0, \infty)$. This follows from the general fact that for a CW pair (Y, B), $Y \times I$ deformation retracts onto $Y \times \{0\} \cup B \times I$. Hence $X \times [0, \infty)$ deformation retracts onto $T \cup (X \times [1, \infty))$, which deformation retracts onto $T \cup (X \times [2, \infty))$, etc., and the infinite concatenation of these deformation retractions is a deformation retraction of $X \times [0, \infty)$ onto T.

Let $T_1 \subset T$ be the subcomplex which is the union of the products $X_i \times [i, i+1]$ for i odd, and let T_2 be the corresponding union for i even. Thus $T_1 \cap T_2 = \coprod_i X_i$ and $T_1 \cup T_2 = T$. For an unreduced cohomology theory h^* we have then a Mayer-Vietoris sequence

$$\begin{array}{ccccccccc}
h^{n-1}(T_1) \oplus h^{n-1}(T_2) & \longrightarrow & h^{n-1}(T_1 \cap T_2) & \longrightarrow & h^n(T) & \longrightarrow & h^n(T_1) \oplus h^n(T_2) & \longrightarrow & h^n(T_1 \cap T_2) \\
\wr \uparrow & & \wr \uparrow & & \wr \uparrow & & \wr \uparrow & & \wr \uparrow \\
\prod_i h^{n-1}(X_i) & \xrightarrow{\varphi} & \prod_i h^{n-1}(X_i) & \longrightarrow & h^n(X) & \longrightarrow & \prod_i h^n(X_i) & \xrightarrow{\varphi} & \prod_i h^n(X_i)
\end{array}$$

The maps φ making the diagram commute are given by $\varphi(\cdots, g_i, \cdots) = (\cdots, (-1)^{i-1}(g_i - \rho(g_{i+1})), \cdots)$, the ρ's being the appropriate restriction maps. This differs from δ only in the sign of its even coordinates, so if we change the isomorphism $h^k(T_1 \cap T_2) \approx \prod_i h^k(X_i)$ by inserting a minus sign in the even coordinates, we can replace φ by δ in the second row of the diagram. This row then yields a short exact sequence $0 \to Coker(\delta) \to H^n(X; G) \to Ker(\delta) \to 0$, finishing the proof for unreduced cohomology.

The same argument works for reduced cohomology if we use the reduced telescope obtained from T by collapsing $\{x_0\} \times [0, \infty)$ to a point, for x_0 a basepoint 0-cell of X_0. Then $T_1 \cap T_2 = \bigvee_i X_i$ rather than $\coprod_i X_i$, and the rest of the argument goes through unchanged. The proof also applies for homology theories, with direct products replaced

by direct sums in the second row of the diagram. As we noted earlier, $Ker(\delta) = 0$ in the direct limit case, and $Coker(\delta) = \varinjlim$. □

Now we consider the earlier Examples 3 and 4 from the viewpoint of cohomology.

Example 7. Consider the mapping telescope $T = \bigcup_i T_i$ from Example 3. The inclusions $T_1 \hookrightarrow T_2 \hookrightarrow \cdots$ induce on $H^n(-;\mathbb{Z})$ the sequence $\cdots \longrightarrow \mathbb{Z} \xrightarrow{p} \mathbb{Z}$ in Example 6. From the proposition we deduce that $H^{n+1}(T;\mathbb{Z}) \approx \widehat{\mathbb{Z}}_p/\mathbb{Z}$ and $\widetilde{H}^k(T;\mathbb{Z}) = 0$ for $k \neq n+1$. This is a somewhat curious state of affairs. The CW complex T is the union of subcomplexes T_i each having cohomology consisting only of a $\mathbb{Z}$ in dimension n, but T itself has no cohomology in dimension n and instead a huge uncountable group $\widehat{\mathbb{Z}}_p/\mathbb{Z}$ in dimension $n+1$. This contrasts sharply with what happens on homology, where the groups $H_n(T_i) \approx \mathbb{Z}$ fit together nicely to give $H_n(T) \approx \mathbb{Z}[1/p]$.

Example 8. A more reasonable behavior is exhibited if we consider the space $X = M(\mathbb{Z}_{p^\infty}, n)$ in Example 4 expressed as the union of its subspaces X_i. By the universal coefficient theorem, the reduced cohomology of X_i with $\mathbb{Z}$ coefficients consists of a $\mathbb{Z}_{p^i} = Ext(\mathbb{Z}_{p^i}, \mathbb{Z})$ in dimension $n+1$. The inclusion $X_i \subset X_{i+1}$ induces the inclusion $\mathbb{Z}_{p^i} \subset \mathbb{Z}_{p^{i+1}}$ on H_n, and on Ext this induced map is a surjection $\mathbb{Z}_{p^{i+1}} \to \mathbb{Z}_{p^i}$ as one can see by looking at the diagram of free resolutions

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathbb{Z} & \xrightarrow{p^i} & \mathbb{Z} & \longrightarrow & \mathbb{Z}_{p^i} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle 1\!\!1} & & \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle p} & & \\
0 & \longrightarrow & \mathbb{Z} & \xrightarrow{p^{i+1}} & \mathbb{Z} & \longrightarrow & \mathbb{Z}_{p^{i+1}} & \longrightarrow & 0
\end{array}$$

The right-hand vertical map is the inclusion $\mathbb{Z}_{p^i} \hookrightarrow \mathbb{Z}_{p^{i+1}}$. When we apply $Hom(-,\mathbb{Z})$ to this diagram we get

$$\begin{array}{ccccccccc}
0 & \longleftarrow & Ext(\mathbb{Z}_{p^i},\mathbb{Z}) & \longleftarrow & Hom(\mathbb{Z},\mathbb{Z}) & \xleftarrow{p^i} & Hom(\mathbb{Z},\mathbb{Z}) \\
 & & \uparrow & & \uparrow{\scriptstyle 1\!\!1} & & \uparrow{\scriptstyle p} \\
0 & \longleftarrow & Ext(\mathbb{Z}_{p^{i+1}},\mathbb{Z}) & \longleftarrow & Hom(\mathbb{Z},\mathbb{Z}) & \xleftarrow{p^{i+1}} & Hom(\mathbb{Z},\mathbb{Z})
\end{array}$$

and the left-hand vertical map is a surjection since the middle one is. Thus the sequence $\cdots \to H^{n+1}(X_2;\mathbb{Z}) \to H^{n+1}(X_1;\mathbb{Z})$ is the sequence in Example 5, and we deduce that $H^{n+1}(X;\mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$, the p-adic integers, and $\widetilde{H}^k(X;\mathbb{Z}) = 0$ for $k \neq n+1$.

These two examples mesh together nicely: If we view X as the mapping cone of the inclusion $S^n \hookrightarrow T$ of the first S^n, then the long exact sequences of homology and

cohomology groups for the pair (T, S^n) reduce to

$$\begin{array}{ccccccccc} 0 \to H_n(S^n) & \to & H_n(T) & \to & H_n(X) & \to 0 \\ \| & & \| & & \| \\ \mathbb{Z} & & \mathbb{Z}[1/p] & & \mathbb{Z}_{p^\infty} \end{array} \qquad \begin{array}{ccccccccc} 0 \to H^n(S^n) & \to & H^{n+1}(X) & \to & H^{n+1}(T) & \to 0 \\ \| & & \| & & \| \\ \mathbb{Z} & & \widehat{\mathbb{Z}}_p & & \widehat{\mathbb{Z}}_p/\mathbb{Z} \end{array}$$

From these examples and the universal coefficient theorem we obtain isomorphisms $Ext(\mathbb{Z}_{p^\infty}, \mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$ and $Ext(\mathbb{Z}[1/p], \mathbb{Z}) \approx \widehat{\mathbb{Z}}_p/\mathbb{Z}$. These can also be derived directly from the definition of Ext. A free resolution of $\mathbb{Z}_{p^\infty}$ is

$$0 \longrightarrow \mathbb{Z}^\infty \xrightarrow{\varphi} \mathbb{Z}^\infty \longrightarrow \mathbb{Z}_{p^\infty} \longrightarrow 0$$

where $\mathbb{Z}^\infty$ is the direct sum of an infinite number of $\mathbb{Z}$'s, the sequences $(x_1, x_2, \cdots)$ of integers all but finitely many of which are zero, and $\varphi(x_1, x_2, \cdots) = (px_1 - x_2, px_2 - x_3, \cdots)$, with infinite matrix consisting of p's on the diagonal and -1's just above the diagonal. Clearly $Ker(\varphi) = 0$ since integers cannot be divided by p infinitely often. The image of φ is generated by the vectors $(p, 0, \cdots), (-1, p, 0, \cdots), (0, -1, p, 0, \cdots), \cdots$ so $Coker(\varphi) \approx \mathbb{Z}_{p^\infty}$. Dualizing by taking $Hom(-, \mathbb{Z})$, we have $Hom(\mathbb{Z}^\infty, \mathbb{Z})$ the infinite direct product of $\mathbb{Z}$'s, and $\varphi^*(y_1, y_2, \cdots) = (py_1, py_2 - y_1, py_3 - y_2, \cdots)$, corresponding to the transpose of the matrix of φ. By definition, $Ext(\mathbb{Z}_{p^\infty}, \mathbb{Z}) = Coker(\varphi^*)$. The image of φ^* consists of the infinite sums $y_1(p, -1, 0 \cdots) + y_2(0, p, -1, 0, \cdots) + \cdots$, so $Coker(\varphi^*)$ can be identified with $\widehat{\mathbb{Z}}_p$ by rewriting a sequence $(z_1, z_2, \cdots)$ as the p-adic number $\cdots z_2 z_1$.

The calculation $Ext(\mathbb{Z}[1/p], \mathbb{Z}) \approx \widehat{\mathbb{Z}}_p/\mathbb{Z}$ is quite similar. A free resolution of $\mathbb{Z}[1/p]$ can be obtained from the free resolution of $\mathbb{Z}_{p^\infty}$ by deleting the first column of the matrix of φ and, for convenience, changing sign, so $\varphi(x_1, x_2, \cdots) = (x_1, x_2 - px_1, x_3 - px_2, \cdots)$ with image generated by $(1, -p, 0, \cdots), (0, 1, -p, 0, \cdots), \cdots$. Then $\varphi^*(y_1, y_2, \cdots) = (y_1 - py_2, y_2 - py_3, \cdots)$ with image the sums $y_1(1, 0 \cdots) + y_2(-p, 1, 0, \cdots) + y_3(0, -p, 1, 0, \cdots) + \cdots$, so $Ext(\mathbb{Z}[1/p], \mathbb{Z}) = Coker(\varphi^*) \approx \widehat{\mathbb{Z}}_p/\mathbb{Z}$. Note that φ^* is exactly the map δ in Example 6.

It is interesting to note also that the map $\varphi : \mathbb{Z}^\infty \to \mathbb{Z}^\infty$ in the two cases $\mathbb{Z}_{p^\infty}$ and $\mathbb{Z}[1/p]$ is precisely the cellular boundary map $H_{n+1}(X^{n+1}, X^n) \to H_n(X^n, X^{n-1})$ for the Moore space $M(\mathbb{Z}_{p^\infty}, n)$ or $M(\mathbb{Z}[1/p], n)$, $n > 1$, constructed earlier as the mapping telescope of the sequence of degree p maps $S^n \to S^n \to \cdots$, with a cell e^{n+1} attached to the first S^n in the case of $\mathbb{Z}_{p^\infty}$.

Exercises

1. Given maps $f_i : X_i \to X_{i+1}$ for integers $i < 0$, show that the 'reverse mapping telescope' obtained by glueing together the mapping cylinders of the f_i's in the obvious way

deformation retracts onto X_0. Similarly, if maps $f_i : X_i \to X_{i+1}$ are given for all $i \in \mathbb{Z}$, show that the resulting 'double mapping telescope' deformation retracts onto any of the ordinary mapping telescopes contained in it, the union of the mapping cylinders of the f_i's for i greater than a given number n.

g. More About Ext

The functors Hom and Ext behave fairly simply for finitely generated groups, when cohomology and homology are essentially the same except for a dimension shift in the torsion. But matters are more complicated in the nonfinitely-generated case. A useful tool for getting a handle on this complication is the following:

Proposition 3.51. *Given an abelian group G and a short exact sequence of abelian groups $0 \to A \to B \to C \to 0$, there are exact sequences*

$$0 \to Hom(G,A) \to Hom(G,B) \to Hom(G,C) \to$$
$$Ext(G,A) \to Ext(G,B) \to Ext(G,C) \to 0$$

$$0 \to Hom(C,G) \to Hom(B,G) \to Hom(A,G) \to$$
$$Ext(C,G) \to Ext(B,G) \to Ext(A,G) \to 0$$

Proof: A free resolution $0 \to F_1 \to F_0 \to G \to 0$ gives rise to a commutative diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & Hom(F_0,A) & \longrightarrow & Hom(F_0,B) & \longrightarrow & Hom(F_0,C) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & Hom(F_1,A) & \longrightarrow & Hom(F_1,B) & \longrightarrow & Hom(F_1,C) & \longrightarrow & 0
\end{array}$$

Since F_0 and F_1 are free, the two rows are exact, being just direct products of copies of the exact sequence $0 \to A \to B \to C \to 0$ since in general $Hom(\bigoplus_i G_i, H) = \prod_i Hom(G_i, H)$. Enlarging the diagram by zeros above and below, it can be viewed as a short exact sequence of chain complexes, and the associated long exact sequence of homology groups is the first of the two six-term exact sequences of the proposition.

To obtain the other exact sequence we shall construct a diagram

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & F_1 & \longrightarrow & F_1' & \longrightarrow & F_1'' & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & F_0 & \longrightarrow & F_0' & \longrightarrow & F_0'' & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

The construction begins with free resolutions of A and C for the two outer columns. Then let $F_0' = F_0 \oplus F_0''$, with the obvious maps in the second row, inclusion and projection. The

map $F_0 \to B$ is defined on the summand F_0 to make the lower left square commute, and on the summand F_0'' it is defined by sending basis elements of F_0'' to elements of B mapping to the images of the basis elements in C, so the lower right square commutes also. It is easy to see that $F_0' \to B$ is surjective with this definition. Define F_1 to be the kernel of $F_0' \to B$, so the middle column is exact. The maps in the upper row are induced from those in the middle row, and we claim the upper row is also exact. This can be seen by regarding the three rows as chain complexes, and the whole diagram as a short exact sequence of chain complexes. In the associated long exact sequence of homology the homology groups from the second and third rows are zero since these rows are exact, so it follows that the homology groups from the upper row are zero also.

The first two rows of the diagram split by freeness, so applying $Hom(-, G)$ yields a diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & Hom(F_0'', G) & \longrightarrow & Hom(F_0', G) & \longrightarrow & Hom(F_0, G) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & Hom(F_1'', G) & \longrightarrow & Hom(F_1', G) & \longrightarrow & Hom(F_1, G) & \longrightarrow & 0 \end{array}$$

with exact rows. The desired six-term exact sequence now follows as in the first paragraph of the proof. □

The second sequence in the proposition says in particular that an injection $A \to B$ induces a surjection $Ext(B, C) \to Ext(A, C)$ for any C. For example, if A has torsion, this says $Ext(A, \mathbb{Z})$ is nonzero since it maps onto $Ext(\mathbb{Z}_n, \mathbb{Z}) \approx \mathbb{Z}_n$ for some $n > 1$. The calculation $Ext(\mathbb{Z}_{p^\infty}, \mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$ in the preceding section shows that torsion in A does not necessarily yield torsion in $Ext(A, \mathbb{Z})$, however.

Also useful are the formulas

$$Ext(\textstyle\bigoplus_i A_i, B) \approx \prod_i Ext(A_i, B) \qquad\qquad Ext(A, \bigoplus_i B_i) \approx \bigoplus_i Ext(A, B_i)$$

whose proof we leave to the exercises. For example, since $\mathbb{Q}/\mathbb{Z} = \bigoplus_p \mathbb{Z}_{p^\infty}$ we obtain $Ext(\mathbb{Q}/\mathbb{Z}, \mathbb{Z}) \approx \prod_p \widehat{\mathbb{Z}}_p$ from the calculation $Ext(\mathbb{Z}_{p^\infty}, \mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$. Then from the exact sequence $0 \to \mathbb{Z} \to \mathbb{Q} \to \mathbb{Q}/\mathbb{Z} \to 0$ we get $Ext(\mathbb{Q}, \mathbb{Z}) \approx (\prod_p \widehat{\mathbb{Z}}_p)/\mathbb{Z}$ using the second exact sequence in the preceding proposition.

In these examples the groups $Ext(A, \mathbb{Z})$ are rather large, and this is part of a general pattern:

Proposition 3.52. *If A is not finitely generated then either $Hom(A, \mathbb{Z})$ or $Ext(A, \mathbb{Z})$ is uncountable. Hence if $H_n(X; \mathbb{Z})$ is not finitely generated, then either $H^n(X; \mathbb{Z})$ or $H^{n+1}(X; \mathbb{Z})$ is uncountable.*

Both possibilities can occur, e.g., $Hom(\bigoplus_\infty \mathbb{Z}, \mathbb{Z}) \approx \prod_\infty \mathbb{Z}$ and $Ext(\mathbb{Z}_{p^\infty}, \mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$.

The proposition has some interesting consequences. First, it implies that if a space X has $\widetilde{H}^*(X;\mathbb{Z}) = 0$, then $\widetilde{H}_*(X;\mathbb{Z}) = 0$, since the case of finitely generated homology groups follows from our earlier results. And second, it says that one cannot always construct a space X with prescribed cohomology groups $H^n(X;\mathbb{Z})$, unlike the situation for homology. For example there is no space whose only nonvanishing $\widetilde{H}^n(X;\mathbb{Z})$ is a countable nonfinitely-generated group such as $\mathbb{Q}$ or $\mathbb{Q}/\mathbb{Z}$. Even in the finitely generated case the dimension $n = 1$ is somewhat special since the group $H^1(X;\mathbb{Z}) \approx Hom(H_1(X),\mathbb{Z})$ is always torsionfree.

Proof: Consider the map $A \xrightarrow{p} A$, $a \mapsto pa$, multiplication by the positive integer p. Denote the kernel, image, and cokernel of this map by ${}_pA$, pA, and A_p, respectively. The short exact sequences $0 \to {}_pA \to A \to pA \to 0$ and $0 \to pA \to A \to A_p \to 0$ give two six-term exact sequences involving $Hom(-,\mathbb{Z})$ and $Ext(-,\mathbb{Z})$. The parts of these exact sequences we need are

$$0 \longrightarrow Hom(pA,\mathbb{Z}) \xrightarrow{\approx} Hom(A,\mathbb{Z}) \longrightarrow Hom({}_pA,\mathbb{Z}) = 0$$

$$Hom(pA,\mathbb{Z}) \longrightarrow Ext(A_p,\mathbb{Z}) \longrightarrow Ext(A,\mathbb{Z})$$

where the term $Hom({}_pA,\mathbb{Z})$ in the first sequence is zero since ${}_pA$ is a torsion group.

Let p be any prime, so A_p is a vector space over $\mathbb{Z}_p$. If this vector space is infinite-dimensional, it is an infinite direct sum of $\mathbb{Z}_p$'s, and $Ext(A_p,\mathbb{Z})$ is the direct product of an infinite numbers of $\mathbb{Z}_p$'s, hence uncountable. Exactness of the second sequence above then implies that one of the two adjacent terms $Ext(A,\mathbb{Z})$ and $Hom(pA,\mathbb{Z}) \approx Hom(A,\mathbb{Z})$ must be uncountable, so we are done when A_p is infinite.

At the other extreme is the possibility that $A_p = 0$. This means that $A = pA$, so every element of A is divisible by p. Hence if A is nontrivial, it then contains a subgroup isomorphic to either $\mathbb{Z}[1/p]$ or $\mathbb{Z}_{p^\infty}$. We have seen that $Ext(\mathbb{Z}[1/p],\mathbb{Z}) \approx \widehat{\mathbb{Z}}_p/\mathbb{Z}$ and $Ext(\mathbb{Z}_{p^\infty},\mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$, an uncountable group in either case. An inclusion $B \hookrightarrow A$ induces a surjection $Ext(A,\mathbb{Z}) \to Ext(B,\mathbb{Z})$, so it follows that $Ext(A,\mathbb{Z})$ is uncountable when $A_p = 0$ and $A \neq 0$.

The remaining case that A_p is a finite direct sum of $\mathbb{Z}_p$'s will be reduced to the case $A_p = 0$. Choose finitely many elements of A whose images in A_p are a set of generators, and let $B \subset A$ be the subgroup generated by these elements. Thus the map $B_p \to A_p$ induced by the inclusion $B \hookrightarrow A$ is surjective. The functor $A \mapsto A_p$ is the same as $A \mapsto A \otimes \mathbb{Z}_p$, so exactness of $B \to A \to A/B \to 0$ implies exactness of $B_p \to A_p \to (A/B)_p \to 0$, and hence $(A/B)_p = 0$. If A is not finitely generated, A/B is nonzero, so the preceding case implies that $Ext(A/B,\mathbb{Z})$ is uncountable. This implies that $Ext(A,\mathbb{Z})$

is also uncountable via the exact sequence $Hom(B,\mathbb{Z}) \to Ext(A/B,\mathbb{Z}) \to Ext(A,\mathbb{Z})$, since $Hom(B,\mathbb{Z})$ is finitely generated hence countable. □

From this proposition one might conjecture that cohomology groups with $\mathbb{Z}$ coefficients are either finitely generated or uncountable.

Exercises

1. Show that $Ext(A,\mathbb{Q}) = 0$ for all A. [Consider the homology with $\mathbb{Q}$ coefficients of a Moore space $M(A,n)$.]

2. An abelian group G is *divisible* if the map $G \xrightarrow{n} G$, $g \in ng$, is surjective for all $n > 1$. Show that a group is divisible iff it is a quotient of a direct sum of $\mathbb{Q}$'s. Deduce from the previous problem that if G is divisble then $Ext(A,G) = 0$ for all A.

3. Show that $Ext(A,\mathbb{Z})$ is isomorphic to the cokernel of $Hom(A,\mathbb{Q}) \to Hom(A,\mathbb{Q}/\mathbb{Z})$, the map induced by the quotient map $\mathbb{Q} \to \mathbb{Q}/\mathbb{Z}$. Use this to get another proof that $Ext(\mathbb{Z}_{p^\infty},\mathbb{Z}) \approx \widehat{\mathbb{Z}}_p$ for p prime.

4. Show that $Ext(\mathbb{Z}_{p^\infty},\mathbb{Z}_p) \approx \mathbb{Z}_p$.

5. Show that, given a short exact sequence of abelian groups $0 \to A \to B \to C \to 0$, a Moore space $M(C,n)$ can be realized as a quotient $M(B,n)/M(A,n)$. Applying the long exact sequence of cohomology for the pair $(M(B,n), M(A,n))$ with any coefficient group G, deduce an exact sequence

$$0 \to Hom(C,G) \to Hom(B,G) \to Hom(A,G) \to$$
$$Ext(C,G) \to Ext(B,G) \to Ext(A,G) \to 0$$

6. As in the proof of Proposition 3.52, for an abelian group A let $p\colon A \to A$ be multiplication by p, with ${}_pA = Ker(p)$, $pA = Im(p)$, and $A_p = Coker(p)$. Show that the six-term exact sequences involving $Hom(-,\mathbb{Z})$ and $Ext(-,\mathbb{Z})$ associated to the short exact sequences $0 \to {}_pA \to A \to pA \to 0$ and $0 \to pA \to A \to A_p \to 0$ can be spliced together to yield the exact sequence across the top of the diagram

$$\begin{array}{ccccccccc} Hom(pA,\mathbb{Z}) & \to & Ext(A_p,\mathbb{Z}) & \to & Ext(A,\mathbb{Z}) & \xrightarrow{p} & Ext(A,\mathbb{Z}) & \to & Ext({}_pA,\mathbb{Z}) \to 0 \\ & & & & & \searrow & \uparrow & & \\ & & & & & & Ext(pA,\mathbb{Z}) & & \\ & & & & & \nearrow & & \searrow & \\ 0 \to Hom(pA,\mathbb{Z}) & \xrightarrow{\approx} & Hom(A,\mathbb{Z}) & \to & 0 & & & & 0 \end{array}$$

where the map labelled 'p' is multiplication by p. Use this to show

(a) $Ext(A,\mathbb{Z})$ is divisible iff A is torsionfree.

(b) $Ext(A,\mathbb{Z})$ is torsionfree if A is divisible, and the converse holds if $Hom(A,\mathbb{Z}) = 0$.

7. Show that the Bockstein long exact sequence in cohomology associated to a short exact sequence of coefficient groups $0 \to A \to B \to C \to 0$ reduces to the exact sequence

$$0 \to Hom(G, A) \to Hom(G, B) \to Hom(G, C) \to$$
$$Ext(G, A) \to Ext(G, B) \to Ext(G, C) \to 0$$

of Proposition 3.51 when the space in question is a Moore space $M(G, n)$.

h. Transfer Homomorphisms

There is a simple construction called 'transfer' which gives very useful information about homology and cohomology of finite-sheeted covering spaces. Let $\pi:\widetilde{X}\to X$ be an n-to-1 covering space projection. In addition to the induced map on singular chains $\pi_\sharp:C_k(\widetilde{X})\to C_k(X)$ there is also a homomorphism in the opposite direction $\tau:C_k(X)\to C_k(\widetilde{X})$ which assigns to a singular simplex $\sigma:\Delta^k\to X$ the sum of the n distinct lifts $\widetilde{\sigma}:\Delta^k\to\widetilde{X}$. This is obviously a chain map, commuting with boundary homomorphisms, so it induces **transfer** homomorphisms $\tau_*:H_k(X;G)\to H_k(\widetilde{X};G)$ and $\tau^*:H^k(\widetilde{X};G)\to H^k(X;G)$ for any coefficient group G. We focus on cohomology in what follows, but similar statements hold for homology as well.

The composition $\pi_\sharp\tau$ is clearly multiplication by n, hence $\tau^*\pi^*=n$. This has the following consequence:

— The map $\pi^*:H^k(X;G)\to H^k(\widetilde{X};G)$ is injective if $G=\mathbb{Q}$ or if $G=\mathbb{Z}_p$ with p a prime not dividing n. If $G=\mathbb{Z}$, the kernel of π^* consists of torsion elements of order dividing n.

Thus the cohomology of $\widetilde{X}$ is 'larger' than that of X except possibly for torsion of order dividing n. This can be a genuine exception as one sees from the examples of S^m covering $\mathbb{R}\mathrm{P}^m$ and, when m is odd, lens spaces. More generally, if $S^m\to X$ is any n-sheeted covering space, then the relation $\tau^*\pi^*=n$ implies that $\widetilde{H}^*(X;\mathbb{Z})$ consists entirely of torsion elements of order dividing n, apart from a possible $\mathbb{Z}$ in dimension m.

By studying the other composition $\pi^*\tau^*$ we will prove:

Proposition 3.53. *Let $\pi:\widetilde{X}\to X$ be an n-sheeted covering space defined by an action of a group Γ on $\widetilde{X}$. Then with coefficients in a field F whose characteristic is 0 or a prime not dividing n, the map $\pi^*:H^k(X;F)\to H^k(\widetilde{X};F)$ is injective with image the subgroup $H^*(\widetilde{X};F)^\Gamma$ consisting of classes α such that $\gamma^*(\alpha)=\alpha$ for all $\gamma\in\Gamma$.*

Proof: We have already seen that π^* is injective, so it remains only to describe its image. Note first that $\tau\pi_\sharp$ sends a singular simplex $\Delta^k\to\widetilde{X}$ to the sum of all its images under the Γ-action. Hence $\pi^*\tau^*(\alpha)=\sum_{\gamma\in\Gamma}\gamma^*(\alpha)$ for $\alpha\in H^k(X;F)$. If α is fixed under the action of Γ on $H^k(\widetilde{X};F)$, the sum $\sum_{\gamma\in\Gamma}\gamma^*(\alpha)$ equals $n\alpha$, so if the coefficient field F has characteristic 0 or a prime not dividing n, we can write $\alpha=\pi^*\tau^*(\alpha/n)$ and thus α lies in the image of π^*. Conversely, since $\pi\gamma=\pi$ for all $\gamma\in\Gamma$, we have $\gamma^*\pi^*(\alpha)=\pi^*(\alpha)$ for all α, and so the image of π^* is contained in $H^*(\widetilde{X};F)^\Gamma$. □

As a simple example, take $X=S^1\vee S^k$, $k>1$, with $\widetilde{X}$ the n-sheeted cover corresponding to the index n subgroup of $\pi_1(X)\approx\mathbb{Z}$. Thus $\widetilde{X}$ is a circle with n k-spheres

attached at equally-spaced points. The deck transformation group $\mathbb{Z}_n$ acts by rotating the circle, permuting the spheres cyclically. Hence for any coefficient group G, the invariant cohomology $H^*(\widetilde{X};G)^{\mathbb{Z}_n}$ is all of H^0 and H^1, plus a copy of G in dimension k, the cellular cohomology classes assigning the same element of G to each of the n k-spheres. Thus $H^i(\widetilde{X};G)^{\mathbb{Z}_n}$ is exactly the image of π^* for $i = 0$ and k, while the image of π^* in dimension 1 is the subgroup $nH^1(\widetilde{X};G)$, and whether this equals $H^1(\widetilde{X};G)^{\mathbb{Z}_n}$ or not depends on G. For $G = \mathbb{Q}$ or $\mathbb{Z}_p$ with p not dividing n, we have equality, but not for $G = \mathbb{Z}$ or $\mathbb{Z}_p$ with p dividing n. In this last case the map π^* is not injective on H^1.

Spaces with Polynomial mod p Cohomology

An interesting special case of the general problem of realizing graded commutative rings as cup product rings of spaces is the case of polynomial rings $\mathbb{Z}_p[x_1,\cdots,x_n]$ over the coefficient field $\mathbb{Z}_p$, p prime. The basic question here is, which sets of numbers $d_1,\cdots,d_n$ are realizable as the dimensions $|x_i|$ of the generators x_i? From Chapter 3 we have the examples of products of $\mathbb{C}\mathrm{P}^\infty$'s and $\mathbb{H}\mathrm{P}^\infty$'s with d_i's equal to 2 or 4, for arbitrary p, and when $p = 2$ we can also take $\mathbb{R}\mathrm{P}^\infty$'s with d_i's equal to 1. As an illustration of transfer homomorphisms we shall now construct some examples with larger d_i's. In the case of polynomials in one variable the method gives examples which turn out to realize everything that can be realized. But for two or more variables, more sophisticated techniques are necessary to realize all the realizable cases; see the end of this section for further remarks on this.

The construction can be outlined as follows. Start with a space Y already known to have polynomial cohomology $H^*(Y;\mathbb{Z}_p) = \mathbb{Z}_p[y_1,\cdots,y_n]$, and suppose there is an action of a finite group Γ on Y. A simple trick called the Borel construction shows that without loss of generality we may assume the action is free, defining a covering space $Y \to Y/\Gamma$. Then by Proposition 3.53 above, if p does not divide the order of Γ, $H^*(Y/\Gamma;\mathbb{Z}_p)$ is isomorphic to the subring of $\mathbb{Z}_p[y_1,\cdots,y_n]$ consisting of polynomials which are invariant under the induced action of Γ on $H^*(Y;\mathbb{Z}_p)$. And in some cases this subring is itself a polynomial ring.

For example if Y is the product of n copies of $\mathbb{C}\mathrm{P}^\infty$ and the symmetric group $\Gamma = \Sigma_n$ acts on Y by permuting the factors, so the induced action on $H^*(Y;Z_p) \approx \mathbb{Z}_p[y_1,\cdots,y_n]$ permutes the y_i's, then the invariant polynomials form a polynomial ring $\mathbb{Z}_p[\sigma_1,\cdots,\sigma_n]$ where σ_i is the i^{th} elementary symmetric polynomial, the sum of all products of i distinct y_j's, so σ_i is a polynomial of degree i. The order of Γ is $n!$ so the condition that p not divide the order of Γ amounts to $p > n$. Thus we realize the polynomial ring $\mathbb{Z}_p[x_1,\cdots,x_n]$ with $|x_i| = 2i$, provided that $p > n$. Actually, this example is less than optimal since

there happens to be another space, the Grassmannian of n-dimensional linear subspaces of $\mathbb{C}^\infty$, whose cohomology with any coefficient ring R is $R[x_1, \cdots, x_n]$ with $|x_i| = 2i$, as we show in §4.4d, so the restriction $p > n$ is unnecessary.

To get further examples the idea is to replace $\mathbb{C}\mathrm{P}^\infty$ by a space with the same $\mathbb{Z}_p$ cohomology but with 'more symmetry,' allowing for larger groups Γ to act. The constructions will be made using $K(\pi, 1)$ spaces, which were introduced in §1.4b. For a group π we constructed there a Δ-complex $B\pi$ with contractible universal cover $E\pi$. The construction is functorial: A homomorphism $\varphi : \pi \to \pi'$ induces a map $B\varphi : B\pi \to B\pi'$, $B\varphi([g_1|\cdots|g_n]) = [\varphi(g_1)|\cdots|\varphi(g_n)]$, satisfying the functor properties $B(\varphi\psi) = B\varphi B\psi$ and $B\mathbb{1} = \mathbb{1}$. In particular, if Γ is a group of automorphisms of π, then Γ acts on $B\pi$.

The other ingredient we shall need is the **Borel construction**, which converts an action of a group Γ on a space Y into a free action on a space $Y' \simeq Y$. Namely, take $Y' = Y \times E\Gamma$ with the diagonal action of Γ, $\gamma(y, z) = (\gamma y, \gamma z)$ where Γ acts on $E\Gamma$ as deck transformations. The diagonal action is free, in fact a covering space action, since this is true for the action in the second coordinate. The orbit space of this diagonal action is denoted $Y \times_\Gamma E\Gamma$.

Example 1. Let $\pi = \mathbb{Z}_p$ and let Γ be the full automorphism group $Aut(\mathbb{Z}_p)$. Automorphisms of $\mathbb{Z}_p$ have the form $x \mapsto mx$ for $(m, p) = 1$, so Γ is the multiplicative group of invertible elements in the field $\mathbb{Z}_p$. By elementary field theory this is a cyclic group, of order $p-1$. The preceding constructions then give a covering space $K(\mathbb{Z}_p, 1) \to K(\mathbb{Z}_p, 1)/\Gamma$ with $H^*(K(\mathbb{Z}_p, 1)/\Gamma; \mathbb{Z}_p) \approx H^*(K(\mathbb{Z}_p, 1); \mathbb{Z}_p)^\Gamma$. We may assume we are in the nontrivial case $p > 2$. From the calculation of the cup product structure of lens spaces in §3.3 or §3.4e we have $H^*(K(\mathbb{Z}_p, 1); \mathbb{Z}_p) \approx \Lambda_{\mathbb{Z}_p}[\alpha] \otimes \mathbb{Z}_p[\beta]$ with $|\alpha| = 1$ and $|\beta| = 2$, and we need to figure out how Γ acts on this cohomology ring.

Let $\gamma \in \Gamma$ be a generator, say $\gamma(x) = mx$. The induced action of γ on $\pi_1(K(\mathbb{Z}_p, 1))$ is also multiplication by m since we have taken $K(\mathbb{Z}_p, 1) = B\mathbb{Z}_p \times E\Gamma$ and γ takes an edge-loop $[g]$ in $B\mathbb{Z}_p$ to $[\gamma(g)] = [mg]$. Hence γ acts on $H_1(K(\mathbb{Z}_p, 1); \mathbb{Z})$ by multiplication by m. It follows that $\gamma(\alpha) = m\alpha$ and $\gamma(\beta) = m\beta$ since $H^1(K(\mathbb{Z}_p, 1); \mathbb{Z}p) \approx Hom(H_1(K(\mathbb{Z}_p, 1)), \mathbb{Z}_p)$ and $H^2(K(\mathbb{Z}_p, 1); \mathbb{Z}p) \approx Ext(H_1(K(\mathbb{Z}_p, 1)), \mathbb{Z}_p)$, and it is a general fact, following easily from the definitions, that multiplication by an integer m in an abelian group H induces multiplication by m in $Hom(H, G)$ and $Ext(H, G)$.

Thus $\gamma(\beta^k) = m^k\beta^k$ and $\gamma(\alpha\beta^k) = m^{k+1}\alpha\beta^k$. Since m was chosen to be a generator of the multiplicative group of invertible elements of $\mathbb{Z}_p$, it follows that the only elements of $H^*(K(\mathbb{Z}_p, 1); \mathbb{Z}_p)$ fixed by γ, hence by Γ, are the scalar multiples of $\beta^{i(p-1)}$ and $\alpha\beta^{i(p-1)-1}$. Thus $H^*(K(\mathbb{Z}_p, 1); \mathbb{Z}_p)^\Gamma = \Lambda_{\mathbb{Z}_p}[\alpha\beta^{p-2}] \otimes \mathbb{Z}_p[\beta^{p-1}]$, so we have produced

a space whose $\mathbb{Z}_p$ -cohomology ring is $\Lambda_{\mathbb{Z}_p}[x_{2p-3}] \otimes \mathbb{Z}_p[y_{2p-2}]$, subscripts indicating dimension.

Example 2. As an easy generalization of the preceding example, replace the group Γ there by a subgroup of $Aut(\mathbb{Z}_p)$ of order d, where d is any divisor of $p-1$. The new Γ is generated by the automorphism $x \mapsto m^{(p-1)/d}x$, and the same analysis shows that we obtain a space with $\mathbb{Z}_p$ -cohomology $\Lambda_{\mathbb{Z}_p}[x_{2d-1}] \otimes \mathbb{Z}_p[y_{2d}]$, subscripts again denoting dimension. For a given choice of d the condition that d divides $p-1$ says $p \equiv 1 \mod d$, which is satisfied by infinitely many p's, according to a classical theorem of Dirichlet.

Example 3. The two preceding examples can be modified so as to eliminate the exterior algebra factors. We shall do this by replacing $\mathbb{Z}_p$ by $\mathbb{Z}_{p^\infty}$, the union of the increasing sequence $\mathbb{Z}_p \subset \mathbb{Z}_{p^2} \subset \mathbb{Z}_{p^3} \subset \cdots$.

The first step is to show that $H^*(K(\mathbb{Z}_{p^\infty}, 1); \mathbb{Z}_p) \approx \mathbb{Z}_p[\beta]$ with $|\beta| = 2$. We know that $\widetilde{H}_*(K(\mathbb{Z}_{p^i}, 1); \mathbb{Z})$ consists of $\mathbb{Z}_{p^i}$'s in odd dimensions. The inclusion $\mathbb{Z}_{p^i} \hookrightarrow \mathbb{Z}_{p^{i+1}}$ induces a map $K(\mathbb{Z}_{p^i}, 1) \to K(\mathbb{Z}_{p^{i+1}}, 1)$ unique up to homotopy, and we can take this map to be a p-sheeted covering space since the covering space of a $K(\mathbb{Z}_{p^{i+1}}, 1)$ corresponding to the unique index p subgroup of its fundamental group is a $K(\mathbb{Z}_{p^i}, 1)$. The homology transfer formula $\pi_*\tau_* = p$ shows that the image of the induced map $H_n(K(\mathbb{Z}_{p^i}, 1); \mathbb{Z}) \to H_n(K(\mathbb{Z}_{p^{i+1}}, 1); \mathbb{Z})$ for n odd contains the multiples of p, hence this map is the inclusion $\mathbb{Z}_{p^i} \hookrightarrow \mathbb{Z}_{p^{i+1}}$. We can then use the universal coefficient theorem to compute the induced map $H^*(K(\mathbb{Z}_{p^{i+1}}, 1); \mathbb{Z}_p) \to H^*(K(\mathbb{Z}_{p^i}, 1); \mathbb{Z}_p)$. Namely, the inclusion $\mathbb{Z}_{p^i} \hookrightarrow \mathbb{Z}_{p^{i+1}}$ obviously induces the trivial map $Hom(\mathbb{Z}_{p^{i+1}}, \mathbb{Z}_p) \to Hom(\mathbb{Z}_{p^i}, \mathbb{Z}_p)$, so on odd-dimensional cohomology the induced map is trivial. On the other hand, from the diagram of free resolutions

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathbb{Z} & \xrightarrow{p^i} & \mathbb{Z} & \longrightarrow & \mathbb{Z}_{p^i} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \mathbb{1}} & & \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle p} & & \\
0 & \longrightarrow & \mathbb{Z} & \xrightarrow{p^{i+1}} & \mathbb{Z} & \longrightarrow & \mathbb{Z}_{p^{i+1}} & \longrightarrow & 0
\end{array}$$

it follows that the inclusion $\mathbb{Z}_{p^i} \hookrightarrow \mathbb{Z}_{p^{i+1}}$ induces an isomorphism $Ext(\mathbb{Z}_{p^{i+1}}, \mathbb{Z}_p) \to Ext(\mathbb{Z}_{p^i}, \mathbb{Z}_p)$, so $H^n(K(\mathbb{Z}_{p^{i+1}}, 1); \mathbb{Z}_p) \to H^n(K(\mathbb{Z}_{p^i}, 1); \mathbb{Z}_p)$ is an isomorphism for n even. Since $\mathbb{Z}_{p^\infty}$ is the union of the increasing sequence of subgroups $\mathbb{Z}_{p^i}$, $B\mathbb{Z}_{p^\infty}$ is the union of the increasing sequence of subcomplexes $B\mathbb{Z}_{p^i}$. Thus we can apply Proposition 3.49 to conclude that $H^*(K(\mathbb{Z}_{p^\infty}, 1); \mathbb{Z}_p)$ is zero in odd dimensions, while in even dimensions the map $H^*(K(\mathbb{Z}_{p^\infty}, 1); \mathbb{Z}_p) \to H^*(K(\mathbb{Z}_p, 1); \mathbb{Z}_p)$ induced by the inclusion $\mathbb{Z}_p \hookrightarrow \mathbb{Z}_{p^\infty}$ is an isomorphism. Thus $H^*(K(\mathbb{Z}_{p^\infty}, 1); \mathbb{Z}_p) \approx \mathbb{Z}_p[\beta]$ as claimed.

Next we show that the map $Aut(\mathbb{Z}_{p^\infty}) \to Aut(\mathbb{Z}_p)$ obtained by restriction to the subgroup $\mathbb{Z}_p \subset \mathbb{Z}_{p^\infty}$ is a split surjection. Automorphisms of $\mathbb{Z}_{p^i}$ are the maps $x \mapsto mx$ for $(m,p) = 1$, so the restriction map $Aut(\mathbb{Z}_{p^{i+1}}) \to Aut(\mathbb{Z}_{p^i})$ is surjective. Since $Aut(\mathbb{Z}_{p^\infty}) = \varprojlim Aut(\mathbb{Z}_{p^i})$, the restriction map $Aut(\mathbb{Z}_{p^\infty}) \to Aut(\mathbb{Z}_p)$ is also surjective. The order of $Aut(\mathbb{Z}_{p^i})$, the multiplicative group of invertible elements of $\mathbb{Z}_{p^i}$, is $p^i - p^{i-1} = p^{i-1}(p-1)$ and $p-1$ is relatively prime to p^{i-1}, so the abelian group $Aut(\mathbb{Z}_{p^i})$ contains a subgroup of order $p-1$. This subgroup maps onto the cyclic group $Aut(\mathbb{Z}_p)$ of the same order, so $Aut(\mathbb{Z}_{p^i}) \to Aut(\mathbb{Z}_p)$ is a split surjection, hence also $Aut(\mathbb{Z}_{p^\infty}) \to Aut(\mathbb{Z}_p)$.

Thus we have an action of $\Gamma = Aut(\mathbb{Z}_p)$ on $B\mathbb{Z}_{p^\infty}$ extending its natural action on $B\mathbb{Z}_p$. The Borel construction then gives an inclusion $B\mathbb{Z}_p \times_\Gamma E\Gamma \hookrightarrow B\mathbb{Z}_{p^\infty} \times_\Gamma E\Gamma$ inducing an isomorphism of $H^*(B\mathbb{Z}_{p^\infty} \times_\Gamma E\Gamma; \mathbb{Z}_p)$ onto the even-dimensional part of $H^*(B\mathbb{Z}_p \times_\Gamma E\Gamma; \mathbb{Z}_p)$, a polynomial algebra $\mathbb{Z}_p[y_{2p-2}]$. Similarly, if d is any divisor of $p-1$, then taking Γ to be the subgroup of $Aut(\mathbb{Z}_p)$ of order d yields a space with $\mathbb{Z}_p$-cohomology the polynomial ring $\mathbb{Z}_p[y_{2d}]$.

Example 4. Now we enlarge the preceding example by bringing in the permutation group, to produce a space with $\mathbb{Z}_p$-cohomology the polynomial ring $\mathbb{Z}_p[y_{2d}, y_{4d}, \cdots, y_{2nd}]$ where d is any divisor of $p-1$ and $p > n$. Let X be the product of n copies of $B\mathbb{Z}_{p^\infty}$ and let Γ be the group of homeomorphisms of X generated by permutations of the factors together with the actions of $\mathbb{Z}_d$ in each factor, constructed in the preceding example. We can view Γ as a group of $n \times n$ matrices with entries in $\mathbb{Z}_p$, the matrices obtained by replacing some of the 1's in a permutation matrix by elements of $\mathbb{Z}_p$ of multiplicative order a divisor of d. Thus there is a split short exact sequence $0 \to (\mathbb{Z}_d)^n \to \Gamma \to \Sigma_n \to 0$, and the order of Γ is $d^n n!$. The product space X has $H^*(X; \mathbb{Z}_p) \approx \mathbb{Z}_p[\beta_1, \cdots, \beta_n]$ with $|\beta_i| = 2$, so $H^*(X \times_\Gamma E\Gamma; Z_p) \approx \mathbb{Z}_p[\beta_1, \cdots, \beta_n]^\Gamma$ provided that p does not divide the order of Γ, i.e., $p > n$. For a polynomial to be invariant under the $\mathbb{Z}_d$ action in each factor it must be a polynomial in the powers β_i^d, and to be invariant under permutations of the variables it must be a symmetric polynomial in these powers. Since symmetric polynomials are exactly the polynomials in the elementary symmetric functions, the polynomials in the β_i's invariant under Γ form a polynomial ring $\mathbb{Z}_p[y_{2d}, y_{4d}, \cdots, y_{2nd}]$ with y_{2k} the sum of all products of k distinct powers β_i^d.

Example 5. As a further variant on the preceding example, choose a divisor q of d and replace Γ by its subgroup consisting of matrices for the which the product of the q^{th} powers of the nonzero entries is 1. This has the effect of enlarging the ring of polynomials invariant under the action, and in fact it can be shown that the invariant polynomials form a polynomial ring $\mathbb{Z}_p[y_{2d}, y_{4d}, \cdots, y_{2(n-1)d}, y_{2nq}]$, with the last generator y_{2nd} replaced by

$y_{2nq} = \prod_i \beta_i^q$. For example, if $n = 2$ and $q = 1$ we obtain $\mathbb{Z}_p[y_4, y_{2d}]$ with $y_4 = \beta_1\beta_2$ and $y_{2d} = \beta_1^d + \beta_2^d$. The group Γ in this case happens to be isomorphic to the dihedral group of order $2d$.

General Remarks

The problem of realizing graded polynomial rings $\mathbb{Z}_p[y]$ in one variable as cup product rings of spaces was discussed in §3.2, and Example 3 provides the remaining examples, showing that $|y|$ can be any even divisor of $2(p-1)$. In more variables the problem of realizing $\mathbb{Z}_p[y_1, \cdots, y_n]$ with specified dimensions $|y_i|$ is more difficult, but has been solved for odd primes p. Here is a sketch of the answer.

Assuming that p is odd, the dimensions $|y_i|$ are even. Call the number $d_i = |y_i|/2$ the *degree* of y_i. In the examples above this was in fact the degree of y_i as a polynomial in the 2-dimensional classes β_j invariant under the action of Γ. It was proved in [Dwyer-Miller-Wilkerson] that every realizable polynomial algebra $\mathbb{Z}_p[y_1, \cdots, y_n]$ is the ring of invariant polynomials $\mathbb{Z}_p[\beta_1, \cdots, \beta_n]^\Gamma$ for an action of a finite group Γ on $\mathbb{Z}_p[\beta_1, \cdots, \beta_n]$, where $|\beta_i| = 2$. The basic examples, whose products yield all realizable polynomial algebras, can be divided into two categories. First there are classifying spaces of Lie groups, each of which realizes a polynomial algebra for all but finitely many primes p. These are listed in the following table.

Lie group	degrees	primes
S^1	1	all
$SU(n)$	$2, 3, \cdots, n$	all
$Sp(n)$	$2, 4, \cdots, 2n$	all
$SO(2k)$	$2, 4, \cdots, 2k-2, k$	$p > 2$
G_2	$2, 6$	$p > 2$
F_4	$2, 6, 8, 12$	$p > 3$
E_6	$2, 5, 6, 8, 9, 12$	$p > 3$
E_7	$2, 6, 8, 10, 12, 14$	$p > 3$
E_8	$2, 8, 12, 14, 18, 20, 24, 30$	$p > 5$

The remaining examples have to be constructed by hand. They form two infinite families plus 30 sporadic exceptions, as shown in the next table.

degrees	primes
$d, 2d, \cdots, (n-1)d, nq$ with $q \mid d$	$p \equiv 1 \mod d$
$2, d$	$p \equiv -1 \mod d$
$4, 6$	$p \equiv 1 \mod 3$
$6, 12$	$p \equiv 1 \mod 3$
$4, 12$	$p \equiv 1 \mod 12$
$12, 12$	$p \equiv 1 \mod 12$
$8, 12$	$p \equiv 1 \mod 4$
$8, 24$	$p \equiv 1 \mod 8$
$12, 24$	$p \equiv 1 \mod 12$
$24, 24$	$p \equiv 1 \mod 24$
$6, 8$	$p \equiv 1, 3 \mod 8$
$8, 12$	$p \equiv 1 \mod 8$
$6, 24$	$p \equiv 1, 19 \mod 24$
$12, 24$	$p \equiv 1 \mod 24$
$20, 30$	$p \equiv 1 \mod 5$
$20, 60$	$p \equiv 1 \mod 20$
$30, 60$	$p \equiv 1 \mod 15$
$60, 60$	$p \equiv 1 \mod 60$
$12, 30$	$p \equiv 1, 4 \mod 15$
$12, 60$	$p \equiv 1, 49 \mod 60$
$12, 20$	$p \equiv 1, 9 \mod 20$
$2, 6, 10$	$p \equiv 1, 4 \mod 5$
$4, 6, 14$	$p \equiv 1, 2, 4 \mod 7$
$6, 9, 12$	$p \equiv 1 \mod 3$
$6, 12, 18$	$p \equiv 1 \mod 3$
$6, 12, 30$	$p \equiv 1, 4 \mod 15$
$4, 8, 12, 20$	$p \equiv 1 \mod 4$
$2, 12, 20, 30$	$p \equiv 1, 4 \mod 5$
$8, 12, 20, 24$	$p \equiv 1 \mod 4$
$12, 18, 24, 30$	$p \equiv 1 \mod 3$
$4, 6, 10, 12, 18$	$p \equiv 1 \mod 3$
$6, 12, 18, 24, 30, 42$	$p \equiv 1 \mod 3$

The first row of the table is the examples we have constructed, though our construction needed the extra condition that p not divide the order of the group Γ. For all entries in both tables the order of Γ, the group such that $\mathbb{Z}_p[y_1, \cdots, y_n] = \mathbb{Z}_p[\beta_1, \cdots, \beta_n]^\Gamma$, turns out to equal the product of the degrees. With the condition that p not divide this order, the method we used for the first row can also be applied to give examples for all the other rows. In some cases the congruence conditions on p, which are needed in order for Γ to be a subgroup of $GL_n(\mathbb{Z}_p)$, automatically imply that p does not divide the order of Γ. But when this is not the case a different construction of a space with the desired cohomology is needed. To find out more about this the reader can begin by consulting [Kane] and

[Notbohm].

For the prime 2 the realization problem is still not completely solved. The known examples are listed in the following table; again we give only the irreducible examples, which generate others by taking products. All but the last entry in the table arise from classifying spaces of Lie groups, as described in §4.4d. The construction for the last entry is in [Dwyer-Wilkerson].

Lie group	degrees
$O(1)$	1
$SO(n)$	$2, 3, \cdots, n$
$SU(n)$	$4, 6, \cdots, 2n$
$Sp(n)$	$4, 8, \cdots, 4n$
—	$8, 12, 14, 15$

i. Local Coefficients

Homology and cohomology with local coefficients are fancier versions of ordinary homology and cohomology which can be defined for nonsimply-connected spaces. In various situations these more refined homology and cohomology theories arise quite naturally, for example in extending Poincaré duality to non-orientable manifolds. We shall first give the quick algebraic definition of homology and cohomology with local coefficients, then we shall reinterpret this definition more geometrically in a way which looks more like ordinary homology and cohomology. The reinterpretation also allows the familiar properties of homology and cohomology to be extended to the local coefficient case with very little effort.

Let X be a path-connected space having a universal cover $\widetilde{X}$ and fundamental group π, so that X is the quotient of $\widetilde{X}$ by the action of π by covering transformations $\widetilde{x} \mapsto \gamma \cdot \widetilde{x}$ for $\gamma \in \pi$ and $\widetilde{x} \in \widetilde{X}$. The action of π on $\widetilde{X}$ induces an action of π on the group $C_n(\widetilde{X})$ of singular n-chains in $\widetilde{X}$, by sending a singular n-simplex $\sigma : \Delta^n \to \widetilde{X}$ to the composition $\Delta^n \xrightarrow{\sigma} \widetilde{X} \xrightarrow{\gamma} \widetilde{X}$. The action of π on $C_n(\widetilde{X})$ makes $C_n(\widetilde{X})$ a module over the group ring $\mathbb{Z}[\pi]$, which consists of the finite formal sums $\sum_i m_i\gamma_i$ with $m_i \in \mathbb{Z}$ and $\gamma_i \in \pi$, with the natural addition $\sum_i m_i\gamma_i + \sum_i n_i\gamma_i = \sum_i (m_i + n_i)\gamma_i$ and multiplication $\big(\sum_i m_i\gamma_i\big)\big(\sum_j n_j\gamma_j\big) = \sum_{i,j} m_i n_j \gamma_i\gamma_j$. The boundary maps $\partial : C_n(\widetilde{X}) \to C_{n-1}(\widetilde{X})$ are $\mathbb{Z}[\pi]$-module homomorphisms since the action of π on these groups comes from an action on $\widetilde{X}$.

If M is an arbitrary module over $\mathbb{Z}[\pi]$, we would like to define $C_n(X; M)$ to be $C_n(\widetilde{X}) \otimes_{\mathbb{Z}[\pi]} M$, but for tensor products over a noncommutative ring one has to be a little careful with left and right module structures. In general, if R is a ring, possibly noncommutative, one defines the tensor product $A \otimes_R B$ of a right R-module A and a left R-module B to be the abelian group with generators $a \otimes b$ for $a \in A$ and $b \in B$, subject to the relations

$$(a_1 + a_2) \otimes b = a_1 \otimes b + a_2 \otimes b, \qquad a \otimes (b_1 + b_2) = a \otimes b_1 + a \otimes b_2, \qquad ar \otimes b = a \otimes rb$$

In case $R = \mathbb{Z}[\pi]$, a left $\mathbb{Z}[\pi]$-module A can be regarded as a right $\mathbb{Z}[\pi]$-module by setting $a\gamma = \gamma^{-1}a$ for $\gamma \in \pi$. So the tensor product of two left $\mathbb{Z}[\pi]$-modules A and B is defined, and the relation $a\gamma \otimes b = a \otimes \gamma b$ becomes $\gamma^{-1}a \otimes b = a \otimes \gamma b$, or equivalently $a' \otimes b = \gamma a' \otimes \gamma b$ where $a' = \gamma^{-1}a$. To simplify notation we shall write $A \otimes_{\mathbb{Z}[\pi]} B$ as $A \otimes_\pi B$, consistent with the fact that the essential part of the $\mathbb{Z}[\pi]$-module structures is the action of π.

In particular, $C_n(\widetilde{X}) \otimes_\pi M$ is defined if M is a left $\mathbb{Z}[\pi]$-module. These chain groups $C_n(X; M) = C_n(\widetilde{X}) \otimes_\pi M$ form a chain complex with the boundary maps $\partial \otimes \mathbb{1}$, and the

homology groups of this chain complex are by definition the **homology groups with local coefficients**, denoted $H_n(X;M)$.

For cohomology one can set $C^n(X;M) = Hom_{\mathbb{Z}[\pi]}(C_n(\widetilde{X}),M)$, the $\mathbb{Z}[\pi]$-module homomorphisms $C_n(\widetilde{X}) \to M$, and these groups $C^n(X;M)$ form a cochain complex whose cohomology groups $H^n(X;M)$ are the **cohomology groups with local coefficients**.

Let us check that when M is a trivial $\mathbb{Z}[\pi]$-module, with $\gamma m = m$ for all $\gamma \in \pi$ and $m \in M$, then $H_n(X;M)$ is just ordinary homology with coefficients in the abelian group M. A singular n-simplex $\sigma : \Delta^n \to X$ has various lifts $\widetilde{\sigma} : \Delta^n \to \widetilde{X}$ which form an orbit of the action of π on $C_n(\widetilde{X})$. In $C_n(\widetilde{X}) \otimes_\pi M$ all these lifts are identified via the relation $\widetilde{\sigma} \otimes m = \gamma\widetilde{\sigma} \otimes \gamma m = \gamma\widetilde{\sigma} \otimes m$. Thus we can identify $C_n(\widetilde{X}) \otimes_\pi M$ with $C_n(X) \otimes M = C_n(X;M)$, so $H_n(X;M)$ reduces to ordinary homology with coefficients in M. The analogous statement for cohomology is also true since elements of $Hom_{\mathbb{Z}[\pi]}(C_n(\widetilde{X}),M)$ are functions from singular n-simplices $\widetilde{\sigma} : \Delta^n \to \widetilde{X}$ to M taking the same value on all elements of a π-orbit since the action of π on M is trivial, so $Hom_{\mathbb{Z}[\pi]}(C_n(\widetilde{X}),M)$ is identifiable with $Hom(C_n(X),M)$, ordinary cochains with coefficients in M.

Suppose we take $M = \mathbb{Z}[\pi]$, viewed as a module over itself via its ring structure. For a ring R with identity element, $A \otimes_R R$ is naturally isomorphic to A via the correspondence $a \otimes r \mapsto ar$. So we have a natural identification $C_n(\widetilde{X}) \otimes_\pi \mathbb{Z}[\pi] = C_n(\widetilde{X})$, hence $H_n(X;\mathbb{Z}[\pi]) \approx H_n(\widetilde{X})$. Generalizing this, let $X' \to X$ be the cover corresponding to a subgroup $\pi' \subset \pi$. Then the free abelian group $\mathbb{Z}[\pi/\pi']$ with basis the cosets $\gamma\pi'$ is a $\mathbb{Z}[\pi]$-module and $C_n(\widetilde{X}) \otimes_{\mathbb{Z}[\pi]} \mathbb{Z}[\pi/\pi'] \approx C_n(X')$, so $H_n(X;\mathbb{Z}[\pi/\pi']) \approx H_n(X')$. More generally, if A is an abelian group then $A[\pi/\pi']$ is a $\mathbb{Z}[\pi]$-module and $H_n(X;A[\pi/\pi']) \approx H_n(X';A)$. So homology of covering spaces is a special case of homology with local coefficients. The corresponding assertions for cohomology however are not true, as we shall see later in the section.

For a $\mathbb{Z}[\pi]$-module M, let π' be the kernel of the homomorphism $\rho : \pi \to Aut(M)$ defining the module structure, given by $\rho(\gamma)(m) = \gamma m$. If $X' \to X$ is the cover corresponding to the normal subgroup π' of π, then $C_n(\widetilde{X}) \otimes_\pi M \approx C_n(X') \otimes_\pi M \approx C_n(X') \otimes_{\mathbb{Z}[\pi/\pi']} M$. This gives a more efficient description of $H_n(X;G)$. As a special case, suppose that we take $M = \mathbb{Z}$, so $Aut(\mathbb{Z}) \approx \mathbb{Z}_2 = \{\pm 1\}$. For a nontrivial $\mathbb{Z}[\pi]$-module structure on M, π' is a subgroup of index 2 and $X' \to X$ is a 2-sheeted covering space. If τ is the nontrivial deck transformation of X', let $C_n^+(X') = \{\alpha \in C_n(X') \mid \tau_\sharp(\alpha) = \alpha\}$ and $C_n^-(X') = \{\alpha \in C_n(X') \mid \tau_\sharp(\alpha) = -\alpha\}$. Thus $C_n^\pm(X')$ has basis the chains $\sigma \pm \tau\sigma$ for $\sigma : \Delta^n \to X'$, and we have a short exact sequence

$$0 \longrightarrow C_n^-(X') \hookrightarrow C_n(X') \xrightarrow{\Sigma} C_n^+(X') \longrightarrow 0$$

where $\Sigma(\alpha) = \alpha + \tau_\sharp(\alpha)$. The transfer homomorphism $C_n(X) \to C_n^+(X')$ sending a singular simplex in X to the sum of its two lifts to X' is an isomorphism. Also, elements of $C_n^-(X')$ are exactly the cosets in the quotient $C_n(X') \otimes_\pi \mathbb{Z}$ of $C_n(X')$, so $C_n^-(X) \approx C_n(X') \otimes_\pi \mathbb{Z}$. These isomorphisms are isomorphisms of chain complexes and the previous displayed short exact sequence is a short exact sequence of chain complexes, so we get a long exact sequence of homology groups

$$\cdots \longrightarrow H_n(X;\mathbb{Z}) \longrightarrow H_n(X') \xrightarrow{p_*} H_n(X) \longrightarrow H_{n-1}(X;\mathbb{Z}) \longrightarrow \cdots$$

where the first and fourth terms are homology with local coefficients in the module $\mathbb{Z}$ and p_* is induced by the covering projection $p : X' \to X$.

Let us apply this exact sequence when X is a non-orientable n-manifold which is closed and connected, and $\mathbb{Z}$ is viewed as a $\mathbb{Z}[\pi]$-module by letting a loop γ in X act on $\mathbb{Z}$ by multiplication by $+1$ or -1 according to whether γ preserves or reverses local orientations of X; see §3.3. The double cover X' of X can be identified with the elements '± 1' in the cover $X_\mathbb{Z}$. Since X is non-orientable, X' is a connected component of $X_\mathbb{Z}$, and X' is orientable since $X_\mathbb{Z}$ is always orientable. The non-orientability of X implies that $H_n(X) = 0$, and of course $H_{n+1}(X) = 0$, so the exact sequence above gives $H_n(X;\mathbb{Z}) \approx H_n(X') \approx \mathbb{Z}$. Thus by taking homology with local coefficients we obtain a fundamental class for a non-orientable manifold.

Local Coefficients via Bundles of Groups

Now we wish to reinterpret homology and cohomology with local coefficients in more geometric terms, making it look more like ordinary homology and cohomology.

A **bundle of groups** is a covering space $p : E \to X$ together with a group structure on each subset $p^{-1}(x)$, such that all these groups $p^{-1}(x)$ are isomorphic to a fixed group G in the following special way: Each point of X has a neighborhood U for which there exists a homeomorphism $h_U : p^{-1}(U) \to U \times G$ taking each $p^{-1}(x)$ to $\{x\} \times G$ by a group isomorphism. This condition relates the topological and algebraic structures, which otherwise would have nothing to do with each other. Borrowing terminology from the theory of fiber bundles, the subsets $p^{-1}(x)$ are called the **fibers** of $p : E \to X$, and one speaks of E as a bundle of groups with fiber G.

For example, associated to an n-dimensional manifold M without boundary is a bundle of groups $p : M_\mathbb{Z} \to M$ with fiber $\mathbb{Z}$, defined in the following way. As a set, $M_\mathbb{Z}$ is the disjoint union of all the groups $H_n(M, M - \{x\};\mathbb{Z})$ as x ranges over M, with p sending $H_n(M, M - \{x\};\mathbb{Z})$ to x. A basis for the topology on $M_\mathbb{Z}$ consists of the sets

$B(\alpha_U)$ of images α_x of an element $\alpha_U \in H_n(M, M-U;\mathbb{Z})$ under the natural isomorphisms $\varphi_{U,x} : H_n(M, M-U;\mathbb{Z}) \to H_n(M, M-\{x\};\mathbb{Z})$, where U ranges over open balls in neighborhoods $\mathbb{R}^n \subset M$. The projection $p : B(\alpha_U) \to U$ is then a homeomorphism, so the map $h_U : p^{-1}(U) \to U \times \mathbb{Z}$, defined by choosing a point $x_0 \in U$ and an isomorphism $H_n(M, M-\{x_0\};\mathbb{Z}) \approx \mathbb{Z}$ and then setting $h_U(\alpha_x) = (p(\alpha_x), \varphi_{U,x_0}\varphi_{U,x}^{-1}(\alpha_x))$, is a homeomorphism taking the fiber $p^{-1}(x)$ isomorphically to $\{x\} \times \mathbb{Z}$. Replacing the coefficient group $\mathbb{Z}$ by any coefficient group G, one obtains in exactly the same way a bundle of groups $M_G \to M$ with fiber G.

Specializing to the case that M is the Möbius band, with its boundary circle deleted to get a manifold without boundary, $M_{\mathbb{Z}}$ is the quotient of $I \times (-1,1) \times \mathbb{Z}$ under the identifications $(0,t,n) \sim (1,-t,-n)$. Disregarding the last coordinate, this is a description of M itself, and the sign change in the $\mathbb{Z}$ coordinate is a result of the non-orientability of M. Note that $M_{\mathbb{Z}}$ contains a copy of M as the elements $0 \in H_2(M, M-\{x\})$. The rest of $M_{\mathbb{Z}}$ consists of disjoint annuli $M_{\pm n}$, the elements $\pm n \in H_2(M, M-\{x\};\mathbb{Z})$. Each projection $M_{\pm n} \to M$ is a 2-sheeted covering space isomorphic to the orientable double cover $\widetilde{M} \to M$ defined for all manifolds in §3.3.

In this Möbius band example, $M_{\mathbb{Z}}$ could also be described as the quotient of the product $\mathbb{R} \times (-1,1) \times \mathbb{Z}$ under the action of $\mathbb{Z} = \pi_1(M)$ defined by $m \cdot (x,t,n) = (x+m, (-1)^m t, (-1)^m n)$. More generally, suppose X is path-connected and has universal cover $\widetilde{X}$ and fundamental group π, so that X is the quotient of $\widetilde{X}$ by the action of π by covering transformations $\widetilde{x} \mapsto \gamma \cdot \widetilde{x}$ for $\gamma \in \pi$ and $\widetilde{x} \in \widetilde{X}$. Suppose also that for some group G a homomorphism $\rho : \pi \to Aut(G)$ is given, where $Aut(G)$ is the group of automorphisms of G. We can think of ρ as defining an action of π on G, $\gamma \cdot g = \rho(\gamma)(g)$. Factoring out the diagonal action of π on $\widetilde{X} \times G$ defined by $\gamma \cdot (\widetilde{x}, g) = (\gamma \cdot \widetilde{x}, \gamma \cdot g)$ gives a quotient space E_ρ, and the projection $\widetilde{X} \times G \to \widetilde{X}$ induces a map $p : E_\rho \to X$. This map $E_\rho \to X$ is a bundle of groups with fiber G since for a neighborhood $U \subset X$ over which $\widetilde{X}$ is a product $U \times \pi$, the part of $E_\rho \to X$ over U is the quotient of $U \times \pi \times G \to U \times \pi$ by the action of π, $\gamma(x, \eta, g) = (x, \gamma\eta, \gamma \cdot g)$, which simply identifies to $U \times G \to U$ all the disjoint product bundles $U \times \{\eta\} \times G \to U \times \{\eta\}$ for $\eta \in \pi$.

An alternative notation for E_ρ is $\widetilde{X} \times_\pi G$, suggestive of the parallel algebraic object $A \otimes_\pi B$.

There is an inverse construction, associating to a bundle of groups $p : E \to X$ with fiber G an action of $\pi_1(X)$ on G, defined as follows. Choose a basepoint $x_0 \in X$. Applying unique path lifting to a path γ from x_0 to $x_1 \in X$ yields a map $\rho(\gamma) : p^{-1}(x_0) \to p^{-1}(x_1)$ which is easily seen to be a group isomorphism depending only on the homotopy class of γ. Restricting attention to loops γ at x_0, the resulting map $\gamma \mapsto \rho(\gamma)$ is clearly a

homomorphism $\pi_1(X) \to Aut(p^{-1}(x_0))$. So we have associated an action $\rho(E)$ of $\pi_1(X)$ on G to the bundle of groups $E \to X$. For example, for the bundle $M_{\mathbb{Z}} \to M$ the action of a loop γ on the fiber $\mathbb{Z}$ is multiplication by ± 1 according to whether γ preserves or reverses orientation in M, i.e., whether γ lifts to a closed loop in the orientable double cover $\widetilde{M} \to M$ or not.

We leave it as an exercise for the reader to prove that these two constructions involving bundles of groups with fiber G and base X are inverses, namely:

(i) For an action ρ of $\pi_1(X)$ on G, $\rho(E_\rho) = \rho$.

(ii) For a bundle of groups E, $E_{\rho(E)}$ is isomorphic to E.

Here an **isomorphism** of bundles of groups is an isomorphism of covering spaces which is a group isomorphism in each fiber.

As another example, the action of $\pi_1(X)$ on itself by inner automorphisms corresponds to a bundle of groups $p : E \to X$ with fibers $p^{-1}(x) = \pi_1(X, x)$. This example is rather similar in spirit to the examples $M_{\mathbb{Z}} \to M$. In both cases one has a functor which associates a group to each point of a space, and all the groups at different points are isomorphic, but not canonically so. Different choices of isomorphisms are obtained by choosing different paths between two points, and loops give rise to an action of π_1 on the fibers.

For the rest of this section we restrict attention to the case of bundles of groups $p : E \to X$ whose fiber G is abelian. Homology groups of X with coefficients in E may then be defined as follows. Consider finite sums $\sum_i n_i\sigma_i$ where each $\sigma_i : \Delta^n \to X$ is a singular n-simplex in X and $n_i : \Delta^n \to E$ is a lifting of σ_i. The sum of two lifts n_i and m_i of the same σ_i is defined by $(n_i+m_i)(s) = n_i(s)+m_i(s)$, and is also a lift of σ_i. In this way the finite sums $\sum_i n_i\sigma_i$ form an abelian group $C_n(X; E)$, provided we allow the deletion of terms $n_i\sigma_i$ when n_i is the zero-valued lift. A boundary homomorphism $\partial : C_n(X; E) \to C_{n-1}(X; E)$ is defined by $\partial\big(\sum_i n_i\sigma_i\big) = \sum_{i,j}(-1)^j n_i\sigma_i\big|[v_0, \cdots, \widehat{v}_j, \cdots, v_n]$ where 'n_i' in the right side of the equation means the restricted lifting $n_i\big|[v_0, \cdots, \widehat{v}_j, \cdots, v_n]$. The proof that the usual boundary homomorphism ∂ satisfies $\partial^2 = 0$ applies equally well in the present context, so the groups $C_n(X; E)$ form a chain complex, whose homology groups are denoted $H_n(X; E)$.

In case E is the product bundle $X \times G$, lifts n_i are simply elements of G, so $H_n(X; E) = H_n(X; G)$, ordinary homology. In the general case, lifts $n_i : \Delta^n \to E$ are uniquely determined by their value at one point $s \in \Delta^n$, and these values can be specified arbitrarily since Δ^n is simply-connected, so the n_i's can be thought of as elements of $p^{-1}(\sigma_i(s))$, a group isomorphic to G. However if E is not a product, there is no canonical isomorphism between different fibers $p^{-1}(x)$ so one cannot identify $H_n(X; E)$ with

ordinary homology.

An alternative approach would be to take the coefficients n_i to be elements of the fiber group over a specific point of $\sigma_i(\Delta^n)$, say $\sigma_i(v_0)$. However, with such a definition the formula for the boundary operator ∂ becomes more complicated since there is no point of Δ^n which lies in all the faces.

Proposition 3.54. *If X is a path-connected space having a universal covering space, then the groups $H_n(X;E)$ are naturally isomorphic to the homology groups $H_n(X;G)$ with local coefficients in the $\mathbb{Z}[\pi]$-module G associated to E, where $\pi = \pi_1(X)$.*

Proof: By our earlier remarks, a bundle of groups $E \to X$ with fiber G is equivalent to an action of π on G, i.e., a $\mathbb{Z}[\pi]$-module structure on G. Thus E is identifiable with the quotient of $\widetilde{X} \times G$ by the diagonal action of π. For a chain $\sum n_i\sigma_i \in C_n(X;E)$, the coefficient n_i gives a lift of σ_i to E, and n_i in turn has various lifts to $\widetilde{X} \times G$. Thus we have natural surjections $C_n(\widetilde{X} \times G) \to C_n(E) \to C_n(X;E)$ expressing each of these groups as a quotient of the preceding one. More precisely, identifying $C_n(\widetilde{X} \times G)$ with $C_n(\widetilde{X}) \otimes \mathbb{Z}[G]$ in the obvious way, then $C_n(E)$ is the quotient of $C_n(\widetilde{X}) \otimes \mathbb{Z}[G]$ under the identifications $\widetilde{\sigma} \otimes g \sim \gamma \cdot \widetilde{\sigma} \otimes \gamma \cdot g$. This quotient is the tensor product $C_n(\widetilde{X}) \otimes_\pi \mathbb{Z}[G]$. To pass to the quotient $C_n(X;E)$ of $C_n(E) = C_n(\widetilde{X}) \otimes_\pi \mathbb{Z}[G]$ we need to take into account the sum operation in $C_n(X;E)$, addition of lifts $n_i : \Delta^n \to E$. This means that in sums $\widetilde{\sigma} \otimes g_1 + \widetilde{\sigma} \otimes g_2 = \widetilde{\sigma} \otimes (g_1 + g_2)$, the term $g_1 + g_2$ should be interpreted not in $\mathbb{Z}[G]$ but in the natural quotient G of $\mathbb{Z}[G]$. Hence $C_n(X;E)$ is identified with the quotient $C_n(\widetilde{X}) \otimes_\pi G$ of $C_n(\widetilde{X}) \otimes_\pi \mathbb{Z}[G]$. This natural identification commutes with the boundary homomorphisms, so the homology groups are also identified. □

More generally, if X has a number of path components X_α with universal covers $\widetilde{X}_\alpha$, then $C_n(X;E) = \bigoplus_\alpha \big(C_n(\widetilde{X}_\alpha) \otimes_{\mathbb{Z}[\pi_1(X_\alpha)]} G\big)$, so $H_n(X;E)$ splits accordingly as a direct sum of the local coefficient homology groups for the path components X_α.

We turn now to the question of whether homology with local coefficients satisfies axioms similar to those for ordinary homology. The main novelty is with the behavior of induced homomorphisms. In order for a map $f : X \to X'$ to induce a map on homology with local coefficients we must have bundles of groups $E \to X$ and $E' \to X'$ which are related in some way. The natural assumption to make is that there is a commutative diagram

$$\begin{array}{ccc} E & \xrightarrow{\widetilde{f}} & E' \\ \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle p'} \\ X & \xrightarrow{f} & X' \end{array}$$

such that $\widetilde{f}$ restricts to a homomorphism in each fiber. Then there is a chain homomorphism $f_\sharp : C_n(X;E) \to C_n(X';E')$ obtained by composing singular simplices with f and their lifts with $\widetilde{f}$, hence there is an induced homomorphism $f_* H_n(X;E) \to H_n(X';E')$. The fibers of E and E' need not be isomorphic groups, so in the case of trivial bundles this construction specializes to Bockstein homomorphisms. To avoid this extra complication we shall consider only the case that $\widetilde{f}$ restricts to an isomorphism on each fiber. With this condition, a commutative diagram as above will be called a **bundle map**.

Here is a method for constructing bundle maps. Starting with a map $f : X \to X'$ and a bundle of groups $p' : E' \to X'$, let

$$E = \{\, (x, e') \in X \times E' \mid f(x) = p'(e') \,\}.$$

This fits into a commutative diagram as above if we define $p(x,e') = x$ and $\widetilde{f}(x,e') = e'$. In particular, the fiber $p^{-1}(x)$ consists of pairs (x,e') with $p'(e') = f(x)$, so $\widetilde{f}$ is a bijection of this fiber with the fiber of $E' \to X'$ over $f(x)$. We use this bijection to give $p^{-1}(x)$ a group structure. To check that $p : E \to X$ is a bundle of groups, let $h' : (p')^{-1}(U') \to U' \times G$ be an isomorphism as in the definition of a bundle of groups. Define $h : p^{-1}(U) \to U \times G$ over $U = f^{-1}(U')$ by $h(x,e') = (x, h'_2(e'))$ where h'_2 is the second coordinate of h'. An inverse for h is $(x,g) \in (x, (h')^{-1}(f(x), g))$, and h is clearly an isomorphism on each fiber. Thus $p : E \to X$ is a bundle of groups, called the **pullback** of $E' \to X'$ via f, or the **induced bundle**, and one often writes $E = f^*(E')$.

Given any bundle map $E \to E'$ as in the diagram, it is routine to check that the map $E \to f^*(E')$, $e \mapsto (p(e), \widetilde{f}(e))$, is an isomorphism of bundles over X, so the pullback construction produces all bundle maps.

Viewing bundles of groups over X as $\mathbb{Z}[\pi_1(X)]$-modules, the pullback construction corresponds to making a $\mathbb{Z}[\pi_1(X')]$-module into a $\mathbb{Z}[\pi_1(X)]$-module by defining $\gamma \cdot g = f_*(\gamma) \cdot g$ for $f_* : \pi_1(X) \to \pi_1(X')$. This follows easily from the definitions. In particular, this implies that homotopic maps $f_0, f_1 : X \to X'$ induce isomorphic pullback bundles $f_0^*(E'), f_1^*(E')$. Hence the map $f_* H_n(X;E) \to H_n(X';E')$ induced by a bundle map depends only on the homotopy class of f.

Generalizing the definition of $H_n(X;E)$ to pairs (X,A) is straightforward, starting with the definition of $H_n(X,A;E)$ as the n^{th} homology group of the quotient complex $C_n(X;E)/C_n(A;E)$ where $p : E \to X$ becomes a bundle of groups over A by restriction to $p^{-1}(A)$. Associated to the pair (X,A) there is then a long exact sequence of homology groups with local coefficients in the bundle E. The excision property is proved just as for ordinary homology, via iterated barycentric subdivision. The final axiom for homology, involving disjoint unions, extends trivially to homology with local coefficients. Simplicial

and cellular homology also extend without difficulty to the case of local coefficients, as do the proofs that these forms of homology agree with singular homology for Δ-complexes and CW complexes, respectively. We leave the verifications of all these statements to the energetic reader.

Now we turn to cohomology. One might try defining $H^n(X;E)$ by simply dualizing, taking $Hom(C_n(X),E)$, but this makes no sense since E is not a group. Instead, the cochain group $C^n(X;E)$ is defined to consist of all functions φ assigning to each singular simplex $\sigma:\Delta^n\to X$ a lift $\varphi(\sigma):\Delta^n\to E$. In case E is the product $X\times G$, this amounts to assigning an element of G to each σ, so this definition generalizes ordinary cohomology. Coboundary maps $\delta:C^n(X;E)\to C^{n+1}(X;E)$ are defined just as with ordinary cohomology, and satisfy $\delta^2=0$, so we have cohomology groups $H^n(X;E)$, and in the relative case, $H^n(X,A;E)$, defined via relative cochains $C^n(X,A;E)=Ker(C^n(X;E)\to C^n(A;E))$.

For a path-connected X with universal cover $\widetilde{X}$ and fundamental group π, we can identify $H^n(X;E)$ with $H^n(X;G)$, cohomology with local coefficients in the $\mathbb{Z}[\pi]$-module G corresponding to E, by identifying $C^n(X;E)$ with $Hom_{\mathbb{Z}[\pi]}(C_n(\widetilde{X}),G)$ in the following way. An element $\varphi\in C^nX;E)$ assigns to each $\sigma:\Delta^n\to X$ a lift to E. Regarding E as the quotient of $\widetilde{X}\times G$ under the diagonal action of π, a lift of σ to E is the same as an orbit of a lift to $\widetilde{X}\times G$. Such an orbit is a function f assigning to each lift $\widetilde{\sigma}:\Delta^n\to\widetilde{X}$ an element $f(\widetilde{\sigma})\in G$ such that $f(\gamma\widetilde{\sigma})=\gamma f(\widetilde{\sigma})$ for all $\gamma\in G$, i.e., an element of $Hom_{\mathbb{Z}[\pi]}(C_n(\widetilde{X}),G)$.

The basic properties of ordinary cohomology in §3.1 extend without great difficulty to cohomology groups with local coefficients. To define the map $f^*:H^n(X';E')\to H^n(X;E)$ induced by a bundle map

$$\begin{array}{ccc} E=f^*(E') & \xrightarrow{\widetilde{f}} & E' \\ \downarrow p & & \downarrow p' \\ X & \xrightarrow{\quad f\quad} & X' \end{array}$$

it suffices to observe that a singular simplex $\sigma:\Delta^n\to X$ and a lift $\widetilde{\sigma}':\Delta^n\to E'$ of $f\sigma$ define a lift $\widetilde{\sigma}=(\sigma,\widetilde{\sigma}'):\Delta^n\to f^*(E)$ of σ. To show that $f\simeq g$ implies $f^*=g^*$ requires some modification of the proof of the corresponding result for ordinary cohomology in §3.1, which proceeded by dualizing the proof for homology. In the local coefficient case one constructs a chain homotopy P^* satisfying $g^\sharp-f^\sharp=P^*\delta+\delta P^*$ directly from the subdivision of $\Delta^n\times I$ used in the proof of the homology result. Similar remarks apply to proving excision and Mayer-Vietoris sequences for cohomology with local coefficients. To prove the equivalence of simplicial and cellular cohomology with singular cohomology, in the local coefficient context, one should use the axiomatic limit argument from the previous section to show that $H^n(X^k;E)\approx H^n(X;E)$ for $k>n$, rather than the argument given

in §3.1. Once again details will be left to the reader.

The difference between homology with local coefficients and cohomology with local coefficients is illuminated by comparing the following proposition with our earlier identification of $H_*(X;\mathbb{Z}[\pi_1(X)])$ with the ordinary homology of the universal cover of X.

Proposition 3.55. *If X is a finite CW complex with universal cover $\widetilde{X}$ and fundamental group π, then for all n, $H^n(X;\mathbb{Z}[\pi])$ is isomorphic to $H^n_c(\widetilde{X};\mathbb{Z})$, cohomology of $\widetilde{X}$ with compact supports and ordinary integer coefficients.*

For example, consider the the n-dimensional torus T^n, the product of n circles, with fundamental group $\pi = \mathbb{Z}^n$ and universal cover $\mathbb{R}^n$. We have $H_i(T^n;\mathbb{Z}[\pi]) \approx H_i(\mathbb{R}^n)$, which is zero except for a $\mathbb{Z}$ in dimension 0, but $H^i(T^n;\mathbb{Z}[\pi]) \approx H^i_c(\mathbb{R}^n)$ vanishes except for a $\mathbb{Z}$ in dimension n, as we saw in the previous section.

To prove the proposition we shall use a few general facts about cohomology with compact supports. One significant difference between ordinary cohomology and cohomology with compact supports is in induced maps. A map $f\colon X \to Y$ induces $f^\sharp\colon C^n_c(Y;G) \to C^n_c(X;G)$ and hence $f^*\colon H^n_c(Y;G) \to H^n_c(X;G)$ provided that f is **proper**: the preimage $f^{-1}(K)$ of each compact set K in Y is compact in X. Thus if $\varphi \in C^n(Y;G)$ vanishes on chains in $Y - K$ then $f^\sharp(\varphi) \in C^n(X;G)$ vanishes on chains in $X - f^{-1}(K)$. Further, to guarantee that $f \simeq g$ implies $f^* = g^*$ we should restrict attention to homotopies which are proper as maps $X \times I \to Y$. Relative groups $H^n_c(X,A;G)$ are defined when A is a closed subset of X, which guarantees that the inclusion $A \hookrightarrow X$ is a proper map. With these constraints the basic theory of §3.1 translates without difficulty to cohomology with compact supports.

In particular, for a locally compact CW complex X one can compute $H^*_c(X;G)$ using *finite* cellular cochains, the cellular cochains vanishing on all but finitely many cells. Namely, to compute $H^n_c(X^n, X^{n-1};G)$ using excision one first has to identify this group with $H^n_c(X^n, N(X^{n-1});G)$ where $N(X^{n-1})$ is a closed neighborhood of X^{n-1} in X^n obtained by deleting an open n-disk from the interior of each n-cell. If X is locally compact, the obvious deformation retraction of $N(X^{n-1})$ onto X^{n-1} is a proper homotopy equivalence, hence via long exact sequences and the five-lemma we obtain isomorphisms $H^n_c(X^n, X^{n-1};G) \approx H^n_c(X^n, N(X^{n-1});G)$, and by excision the latter group can be identified with the finite cochains.

Proof of 3.55: As noted above, we can compute $H^*_c(\widetilde{X};\mathbb{Z})$ using the groups $C^n_f(\widetilde{X};\mathbb{Z})$ of finite cellular cochains $\varphi\colon C_n \to \mathbb{Z}$, where $C_n = H_n(\widetilde{X}^n, \widetilde{X}^{n-1})$. Giving $\widetilde{X}$ the CW structure lifting that on X, then since X is compact, finite cellular cochains are exactly homomorphisms $\varphi\colon C_n \to \mathbb{Z}$ such that for each cell e^n of $\widetilde{X}$, $\varphi(\gamma e^n)$ is nonzero for only

finitely many covering transformations $\gamma \in \pi$. Such a φ determines a $\mathbb{Z}[\pi]$-homomorphism $\widehat{\varphi}: C_n \to \mathbb{Z}[\pi]$ defined by $\widehat{\varphi}(e^n) = \sum_\gamma \varphi(\gamma^{-1}e^n)\gamma$; this is a $\mathbb{Z}[\pi]$-homomorphism since if we replace the summation index γ in the right side of $\varphi(\eta e^n) = \sum_\gamma \varphi(\gamma^{-1}\eta e^n)\gamma$ by $\eta\gamma$, we get $\sum_\gamma \varphi(\gamma^{-1}e^n)\eta\gamma$. The function $\varphi \mapsto \widehat{\varphi}$ defines a homomorphism $C^n_f(\widetilde{X};\mathbb{Z}) \to Hom_{\mathbb{Z}[\pi]}(C_n, \mathbb{Z}[\pi])$ which is injective since φ is recoverable from $\widehat{\varphi}$ as the coefficient of $\gamma = 1$, and is surjective since a $\mathbb{Z}[\pi]$-homomorphism $\psi: M \to \mathbb{Z}[\pi]$ has the form $\psi(x) = \sum_\gamma \psi_\gamma(x)\gamma$ with $\psi_\gamma \in Hom_{\mathbb{Z}}(M, \mathbb{Z})$ satisfying $\psi_\gamma(x) = \psi_1(\gamma^{-1}x)$, so ψ_1 determines ψ. The isomorphisms $C^n_f(\widetilde{X};\mathbb{Z}) \approx Hom_{\mathbb{Z}[\pi]}(C_n, \mathbb{Z}[\pi])$ are isomorphisms of cochain complexes, so the respective cohomology groups $H^n_c(\widetilde{X};\mathbb{Z})$ and $H^n(X;\mathbb{Z}[\pi])$ are isomorphic. □

Cup and cap product work easily with local coefficients in a bundle of rings, the latter concept being defined in the obvious way. The cap product can be used to extend Poincaré duality to non-orientable manifolds M, using local coefficients in $M_{\mathbb{Z}}$ or more generally M_R for a ring R:

Theorem 3.56. *For an arbitrary closed n-manifold M there is a fundamental class $[M] \in H_n(M; M_R)$ such that $[M]\frown : H^k(M; M_R) \to H_{n-k}(M; M_R)$ is an isomorphism for all k.*

With the definitions we have given this is all perfectly straightforward, so we shall not stop to give details here.

Exercises

1. Compute $H_*(S^1; E)$ and $H^*(S^1; E)$ where $E \to S^1$ is the nontrivial bundle with fiber $\mathbb{Z}$.

2. Show the manifold M is R-orientable iff the bundle M_R is isomorphic to the product $M \times R$.

3. Compute the homology groups with local coefficients $H_n(M; M_{\mathbb{Z}})$ for a closed non-orientable surface M.

4. Let $\mathcal{B}(X; G)$ be the set of isomorphism classes of bundles of groups $E \to X$ with fiber G, and let $E_0 \to BAut(G)$ be the bundle corresponding to the 'identity' action $\rho: Aut(G) \to Aut(G)$. Show that for a CW complex X the map $\Phi: [X, BAut(G)] \to \mathcal{B}(X, G)$, $\Phi([f]) = f^*(E_0)$, is a bijection, where $[X, Y]$ denotes the set of homotopy classes of maps $X \to Y$.

5. Show that if finite connected CW complexes X and Y are homotopy equivalent, then their universal covers $\widetilde{X}$ and $\widetilde{Y}$ are proper homotopy equivalent.

6. If X is a finite nonsimply-connected graph, show that $H^n(X;\mathbb{Z}[\pi_1(X)])$ is zero unless $n = 1$, when it is the direct sum of a countably infinite number of $\mathbb{Z}$'s. [Use Proposition 3.55 and compute $H_c^n(\widetilde{X})$ as $\varinjlim H^n(\widetilde{X}, \widetilde{X} - T_i)$ for a suitable sequence of finite subtrees $T_1 \subset T_2 \subset \cdots$ of $\widetilde{X}$ with $\bigcup_i T_i = \widetilde{X}$.]

7. Show that homology groups $H_n^{lf}(X;G)$ can be defined using **locally finite** chains, which are formal sums $\sum_\sigma g_\sigma \sigma$ of singular simplices $\sigma : \Delta^n \to X$ with coefficients $g_\sigma \in G$, such that each $x \in X$ has a neighborhood meeting the images of only finitely many σ's with $g_\sigma \neq 0$. Develop this homology theory far enough to show that for a locally compact CW complex X, $H_n^{lf}(X;G)$ can be computed using infinite cellular chains $\sum_\alpha g_\alpha e_\alpha^n$.

Chapter 3
Cohomology

Cohomology is an algebraic variant of homology, the result of a simple dualization in the definition. Not surprisingly, cohomology groups satisfy axioms much like the axioms for homology, except that induced homomorphisms go in the opposite direction as a result of the dualization. The basic difference between homology and cohomology is thus that cohomology groups are contravariant functors while homology groups are covariant. In terms of intrinsic information, however, there is not a big difference between homology groups and cohomology groups. The homology groups of a space determine its cohomology groups, and the converse holds at least when the homology groups are finitely generated.

What is a little surprising is that contravariance leads to extra structure in cohomology. This first appears in a natural product, called cup product, which makes the cohomology groups of a space into a ring. This is an extremely useful piece of additional structure, and much of this chapter is devoted to studying cup products, which are considerably more subtle than the additive structure of cohomology.

How does contravariance lead to a product in cohomology which is not present in homology? In homology one does have a natural 'cross product' $H_i(X) \times H_j(Y) \longrightarrow H_{i+j}(X \times Y)$, defined in §3.4b, so if we take $Y = X$ we have the first half of a hypothetical product

$$H_i(X) \times H_j(X) \longrightarrow H_{i+j}(X \times X) \longrightarrow H_{i+j}(X)$$

The difficulty is in defining the second map. The natural thing would be for this to be induced by a map $X \times X \to X$. The multiplication map in a topological group or more generally an H-space is such a map, and the resulting product, called Pontryagin product, can be quite useful; see §3.4c. But for general X, the only natural maps $X \times X \to X$ are the projections onto one of the factors, and since these projections collapse the other factor to a point, the resulting product in homology is rather trivial.

With cohomology, however, the situation is better. One still has a cross product $H^i(X) \times H^j(Y) \longrightarrow H^{i+j}(X \times Y)$ which can be constructed in the same way as in homology, so one can again take $X = Y$ and get the first half of a product

$$H^i(X) \times H^j(X) \longrightarrow H^{i+j}(X \times X) \longrightarrow H^{i+j}(X)$$

But now the second map would be induced by a map $X \to X \times X$, and there is an obvious candidate for this map, the diagonal map $\Delta(x) = (x, x)$. This turns out to work nicely, giving a very well-behaved product in cohomology, the cup product.

Another sort of extra structure in cohomology whose existence is traceable to contravariance is provided by cohomology operations. These make the cohomology groups of a space into a module over a certain ring, whose algebraic structure is rather complicated. Cohomology operations lie at a depth somewhat greater than the cup product structure, so we defer their study to Chapter 4.4k.

The extra layer of algebra in cohomology arising from the dualization in its definition may seem at first to be separating it further from topology, but there are a quite a few topological situations where cohomology arises quite naturally. One of these is Poincaré duality, the topic of the third section of this chapter. Another is obstruction theory, in §4.3. Characteristic classes in vector bundle theory (see [Milnor-Stasheff] or [VBKT]) provide yet another instance.

From the viewpoint of homotopy theory, cohomology is in some ways more basic than homology. As we shall see at the end of §4.1, cohomology has a description as homotopy classes of maps which is in some respects dual to the definition of homotopy groups. There is an analog of this for homology, but it is not nearly as nice. This homotopical viewpoint is explored in some generality in §4.4e and §4.4f.

1. Cohomology Groups

Homology groups $H_n(X)$ are the result of a two-stage process: First one forms a chain complex $\cdots \longrightarrow C_n \xrightarrow{\partial} C_{n-1} \longrightarrow \cdots$ of singular, simplicial, or cellular chains, then one takes the homology groups of this chain complex, $Ker(\partial)/Im(\partial)$. To obtain homology groups $H_n(X;G)$ with coefficients in an abelian group G, there is an intermediate step, which can be described as tensoring the chain complex with G. The idea of cohomology is to insert a different middle step: Replace each group C_n by its dual **cochain group** $C_n^* = Hom(C_n, G)$, the group of homomorphisms $C_n \to G$, and replace each boundary map $\partial : C_n \to C_{n-1}$ by its dual **coboundary map** $\delta = \partial^* : C_{n-1}^* \to C_n^*$. Notice that δ goes in the opposite direction from ∂, increasing rather than decreasing dimension. This is because, in general, a homomorphism $\alpha : A \to B$ induces a dual homomorphism $\alpha^* : Hom(B, G) \to Hom(A, G)$ defined by $\alpha^*(\varphi) = \varphi\alpha$. Dual homomorphisms obviously satisfy the properties $(\alpha\beta)^* = \beta^*\alpha^*$, $\mathbb{1}^* = \mathbb{1}$, $0^* = 0$. In particular, since $\partial\partial = 0$ it follows that $\delta\delta = 0$, so **cohomology groups** $Ker(\delta)/Im(\delta)$ can be defined.

The plan for this section is first to sort out the algebra and show that the cohomology groups of a chain complex are determined algebraically by the homology groups. Then, after this algebraic excursion, we will show that cohomology groups of spaces satisfy axioms like those for homology. The payoff for all this formal work will begin to be apparent in

subsequent sections.

The Universal Coefficient Theorem

Let us begin with a simple example. Consider the chain complex

$$\begin{array}{ccccccccccc} 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{0} & \mathbb{Z} & \xrightarrow{2} & \mathbb{Z} & \xrightarrow{0} & \mathbb{Z} & \longrightarrow & 0 \\ & & \| & & \| & & \| & & \| & & \\ & & C_3 & & C_2 & & C_1 & & C_0 & & \end{array}$$

where $\mathbb{Z} \xrightarrow{2} \mathbb{Z}$ is the map $x \mapsto 2x$. If we dualize by taking $Hom(-, G)$ with $G = \mathbb{Z}$, we obtain the cochain complex

$$\begin{array}{ccccccccccc} 0 & \longleftarrow & \mathbb{Z} & \xleftarrow{0} & \mathbb{Z} & \xleftarrow{2} & \mathbb{Z} & \xleftarrow{0} & \mathbb{Z} & \longleftarrow & 0 \\ & & \| & & \| & & \| & & \| & & \\ & & C_3^* & & C_2^* & & C_1^* & & C_0^* & & \end{array}$$

In the original chain complex the homology groups are $\mathbb{Z}$'s in dimensions 0 and 3, together with a $\mathbb{Z}_2$ in dimension 1. In the cohomology of the cochain complex we still get $\mathbb{Z}$'s in dimensions 0 and 3, but the $\mathbb{Z}_2$ has shifted up a dimension, to dimension 2.

More generally, consider any chain complex of finitely generated free abelian groups. It was an exercise in §3.1 to show that such a chain complex splits as the direct sum of elementary complexes $0 \longrightarrow \mathbb{Z} \longrightarrow 0$ and $0 \longrightarrow \mathbb{Z} \xrightarrow{m} \mathbb{Z} \longrightarrow 0$. Applying $Hom(-, \mathbb{Z})$, we obtain the direct sum of the corresponding dual complexes $0 \longleftarrow \mathbb{Z} \longleftarrow 0$ and $0 \longleftarrow \mathbb{Z} \xleftarrow{m} \mathbb{Z} \longleftarrow 0$. Thus the cohomology groups are the same as the homology groups except that torsion is shifted up one dimension. We will see later in this section that the same relation between homology and cohomology holds whenever the homology groups are finitely generated, even when the chain groups are not finitely generated. It would also be quite easy to see in this example what happens if $Hom(-, \mathbb{Z})$ is replaced by $Hom(-, G)$, since the dual elementary cochain complexes would then be $0 \longleftarrow G \longleftarrow 0$ and $0 \longleftarrow G \xleftarrow{m} G \longleftarrow 0$.

Consider now a completely general chain complex C of free abelian groups, of the form $\cdots \longrightarrow C_n \xrightarrow{\partial} C_{n-1} \longrightarrow \cdots$. Dualizing this by taking $Hom(-, G)$ gives a cochain complex $\cdots \longleftarrow C_n^* \xleftarrow{\delta} C_{n-1}^* \longleftarrow \cdots$. Our goal is to show how the cohomology groups $H^n(C; G) = Ker(\delta)/Im(\delta)$ of this cochain complex are determined solely by G and the homology groups $H_n(C) = Ker(\partial)/Im(\partial)$.

An optimistic first guess might be that $H^n(C; G)$ is isomorphic to $Hom(H_n(C), G)$, but the example above, where H_2 was zero while H^2 was nonzero, shows this cannot be true in general. Nevertheless, there is a natural map $h : H^n(C; G) \to Hom(H_n(C), G)$,

defined as follows. Denote the cycles and boundaries by $Z_n = Ker(\partial) \subset C_n$ and $B_n = Im(\partial) \subset C_n$. A class in $H^n(C;G)$ is represented by a homomorphism $\varphi : C_n \to G$ such that $\delta\varphi = 0$, i.e., $\varphi\partial = 0$, or in other words, φ vanishes on B_n. The restriction of φ to Z_n then induces a quotient homomorphism $h(\varphi) : Z_n/B_n \to G$, an element of $Hom(H_n(C), G)$. If φ is in $Im(\delta)$, say $\varphi = \delta\psi = \psi\partial$, then φ is zero on Z_n, so $h(\varphi) = 0$. Thus there is a well-defined quotient map $h : H^n(C;G) \to Hom(H_n(C), G)$, which is obviously a homomorphism.

It is not hard to see that h is surjective. The short exact sequence $0 \to Z_n \to C_n \to B_{n-1} \to 0$ splits since B_{n-1} is free, being a subgroup of the free abelian group C_{n-1}. Thus there is a projection $p : C_n \to Z_n$ which restricts to the identity on Z_n. Composing with p gives a way of extending homomorphisms $Z_n \to G$ to homomorphisms $C_n \to G$, In particular, this extends homomorphisms $Z_n \to G$ which vanish on B_n, i.e., homomorphisms $H_n(C) \to G$, to homomorphisms $C_n \to G$ which vanish on B_n, i.e., elements of $Ker(\delta)$. Thus we have a homomorphism $Hom(H_n(C), G) \to Ker(\delta)$. Composing this with the quotient map $Ker(\delta) \to H^n(C;G)$ gives a map $Hom(H_n(C), G) \to H^n(C;G)$ whose composition with h is the identity map on $Hom(H_n(C), G)$ since the effect of composing with h is simply to undo the effect of extending homomorphisms via p. This shows h is surjective, and in fact it shows that the short exact sequence

$$0 \longrightarrow Ker(h) \longrightarrow H^n(C;G) \xrightarrow{h} Hom(H_n(C), G) \longrightarrow 0$$

splits.

The remaining task then is to analyze $Ker(h)$. A convenient way to start the process is to consider not just the chain complex C, but also its subcomplexes consisting of the cycles and the boundaries. Thus we consider the commutative diagram of short exact sequences

(i)
$$\begin{array}{ccccccccc} 0 & \longrightarrow & Z_{n+1} & \longrightarrow & C_{n+1} & \xrightarrow{\partial} & B_n & \longrightarrow & 0 \\ & & \downarrow 0 & & \downarrow \partial & & \downarrow 0 & & \\ 0 & \longrightarrow & Z_n & \longrightarrow & C_n & \xrightarrow{\partial} & B_{n-1} & \longrightarrow & 0 \end{array}$$

where the vertical boundary maps on Z_{n+1} and B_n are the restrictions of the boundary map in the complex C, hence are zero. Dualizing (i) gives a commutative diagram

(ii)
$$\begin{array}{ccccccccc} 0 & \longleftarrow & Z^*_{n+1} & \longleftarrow & C^*_{n+1} & \longleftarrow & B^*_n & \longleftarrow & 0 \\ & & \uparrow 0 & & \uparrow \delta & & \uparrow 0 & & \\ 0 & \longleftarrow & Z^*_n & \longleftarrow & C^*_n & \longleftarrow & B^*_{n-1} & \longleftarrow & 0 \end{array}$$

The rows here are exact since, as we have already remarked, the rows of (i) split, and the dual of a split short exact sequence is a split short exact sequence because of the natural isomorphism $Hom(A \oplus B, G) \approx Hom(A, G) \oplus Hom(B, G)$.

We may view (ii), like (i), as part of a short exact sequence of chain complexes, with the coboundary maps in the Z_n^* and B_n^* complexes zero since they are the duals of those in (i). Hence the associated long exact sequence of homology groups has the form

$$(iii) \qquad \cdots \longleftarrow B_n^* \longleftarrow Z_n^* \longleftarrow H^n(C;G) \longleftarrow B_{n-1}^* \longleftarrow Z_{n-1}^* \longleftarrow \cdots$$

The 'boundary maps' $Z_n^* \to B_n^*$ in this long exact sequence are in fact the dual maps i_n^* of the inclusions $i_n : B_n \to Z_n$, as one sees by recalling how these boundary maps are defined: In (ii) one takes an element of Z_n^*, pulls this back to C_n^*, applies δ to get an element of C_{n+1}^*, then pulls this back to B_n^*. The first of these steps extends a homomorphism $\varphi : Z_n \to G$ to C_n, the second step composes this extension with ∂, and the third step undoes this composition and restricts to B_n. The net effect is just to restrict φ from Z_n to B_n.

A long exact sequence can always be broken up into short exact sequences, and doing this for the sequence (iii) yields short exact sequences

$$(iv) \qquad 0 \longleftarrow Ker(i_n^*) \longleftarrow H^n(C;G) \longleftarrow Coker(i_{n-1}^*) \longleftarrow 0$$

The group $Ker(i_n^*)$ can be identified naturally with $Hom(H_n(C), G)$ since elements of $Ker(i_n^*)$ are homomorphisms $Z_n \to G$ which vanish on the subgroup B_n, and such homomorphisms are the same as homomorphisms $Z_n/B_n \to G$. Under this identification of $Ker(i_n^*)$ with $Hom(H_n(C), G)$, the map $H^n(C;G) \to Ker(i_n^*)$ in (iv) becomes the map h considered earlier.

Thus we can rewrite (iv) as a short exact sequence

$$(v) \qquad 0 \longrightarrow Coker(i_{n-1}^*) \longrightarrow H^n(C;G) \xrightarrow{h} Hom(H_n(C), G) \longrightarrow 0$$

which splits. Our objective is to show that the more mysterious term $Coker(i_{n-1}^*)$ depends only on $H_{n-1}(C)$ and G, in a natural, functorial way.

First let us observe that $Coker(i_{n-1}^*)$ would be zero if it were always true that the dual of a short exact sequence is exact, since the dual of the short exact sequence

$$(vi) \qquad 0 \longrightarrow B_{n-1} \xrightarrow{i_{n-1}} Z_{n-1} \longrightarrow H_{n-1}(C) \longrightarrow 0$$

is the sequence

$$(vii) \qquad 0 \longleftarrow B_{n-1}^* \xleftarrow{i_{n-1}^*} Z_{n-1}^* \longleftarrow H_{n-1}(C)^* \longleftarrow 0$$

so if this were exact at B^*_{n-1}, i^*_{n-1} would be surjective, hence $Coker(i^*_{n-1})$ would be zero. This argument does apply if $H_{n-1}(C)$ happens to be free, since (vi) splits in this case, which implies that (vii) is also split exact. So in this case the map h in (v) is an isomorphism. However, in the general case it is easy to produce examples of short exact sequences whose duals are not exact, e.g., if we dualize $0 \longrightarrow \mathbb{Z} \xrightarrow{n} \mathbb{Z} \longrightarrow \mathbb{Z}_n \longrightarrow 0$ by applying $Hom(-,\mathbb{Z})$ we get $0 \longleftarrow \mathbb{Z} \xleftarrow{n} \mathbb{Z} \longleftarrow 0 \longleftarrow 0$ which fails to be exact at the left-hand $\mathbb{Z}$, precisely the place we are interested in for $Coker(i^*_{n-1})$.

We might mention in passing that the loss of exactness at the left end of a short exact sequence after dualization is in fact all that goes wrong, in view of the following:

Exercise. If $A \to B \to C \to 0$ is exact, then dualizing by applying $Hom(-,G)$ yields an exact sequence $A^* \leftarrow B^* \leftarrow C^* \leftarrow 0$.

However, we will not need this fact in what follows.

The sequence (vi) has the special feature that both B_{n-1} and Z_{n-1} are free, so (vi) can be regarded as a free resolution of $H_{n-1}(C)$, where a **free resolution** of an abelian group H is an exact sequence

$$\cdots \longrightarrow F_2 \xrightarrow{f_2} F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} H \longrightarrow 0$$

with each F_n free. If we dualize this free resolution by applying $Hom(-,G)$, we may lose exactness, but at least we get a chain complex (or perhaps we should say 'cochain complex,' but algebraically there is no difference):

$$\cdots \longleftarrow F_1^* \xleftarrow{f_1^*} F_0^* \xleftarrow{f_0^*} H^* \longleftarrow 0$$

Let us use the temporary notation $H^n(F;G)$ for the homology group $Ker(f^*_{n+1})/Im(f^*_n)$ of this dual complex. Note that the group $Coker(i^*_{n-1})$ which we are interested in is $H^1(F;G)$ where F is the free resolution in (vi). Part (b) of the following lemma therefore shows that $Coker(i^*_{n-1})$ depends only on H_{n-1} and G.

Lemma 3.1.

(a) *Given free resolutions F and F' of abelian groups H and H', then every homomorphism $\alpha : H \to H'$ can be extended to a chain map between these two free resolutions:*

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & F_2 & \xrightarrow{f_2} & F_1 & \xrightarrow{f_1} & F_0 & \xrightarrow{f_0} & H & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \alpha_2} & & \downarrow{\scriptstyle \alpha_1} & & \downarrow{\scriptstyle \alpha_0} & & \downarrow{\scriptstyle \alpha} & & \\
\cdots & \longrightarrow & F'_2 & \xrightarrow{f'_2} & F'_1 & \xrightarrow{f'_1} & F'_0 & \xrightarrow{f'_0} & H & \longrightarrow & 0
\end{array}$$

Furthermore, any two such chain maps extending α are chain homotopic.

(b) *Given any two free resolutions F and F' of H, there are canonical isomorphisms $H^n(F;G) \approx H^n(F';G)$ for all n.*

Proof: The α_i's will be constructed inductively. Since the F_i's are free, it suffices to define each α_i on a basis for F_i. To define α_0, observe that surjectivity of f_0' implies that for each basis element x of F_0 there exists $x' \in F_0'$ such that $f_0'(x') = \alpha f_0(x)$, so we define $\alpha_0(x) = x'$. We would like to define α_1 in the same way, sending a basis element $x \in F_1$ to an element $x' \in F_1'$ such that $f_1'(x') = \alpha_0 f_1(x)$. Such an x' will exist if $Im(\alpha_0 f_1) \subset Im(f_1')$. Since $Im(f_1') = Ker(f_0')$, this condition becomes $Im(\alpha_0 f_1) \subset Ker(f_0')$, or $f_0'\alpha_0 f_1 = 0$, which holds since $f_0'\alpha_0 = \alpha f_0$ and $f_0 f_1 = 0$. The same procedure defines all the subsequent α_i's.

If we have another chain map given by maps $\alpha_i' : F_i \to F_i'$, then the differences $\beta_i = \alpha_i - \alpha_i'$ define a chain map extending the zero homomorphism $H \to H'$. It will suffice to construct maps $\lambda_i : F_i \to F_{i+1}'$ defining a chain homotopy from β_i to 0, i.e., with $\beta_i = f_{i+1}'\lambda_i + \lambda_{i-1} f_i$, where $\lambda_{-1} : H \to F_0'$ is defined to be zero. The λ_i's are constructed inductively by a procedure much like the construction of the α_i's. We wish to define λ_i to take a basis element $x \in F_i$ to an element $x' \in F_{i+1}'$ such that $f_{i+1}'(x') = \beta_i(x) - \lambda_{i-1} f_i(x)$. This will be possible if $Im(\beta_i - \lambda_{i-1} f_i) \subset Im(f_{i+1}') = Ker(f_i')$, in other words if $f_i'(\beta_i - \lambda_{i-1} f_i) = 0$. Using the relations $f_i'\beta_i = \beta_{i-1} f_i$ and $\beta_{i-1} = f_i'\lambda_{i-1} + \lambda_{i-2} f_{i-1}$, which holds by induction, we have

$$\begin{aligned} f_i'(\beta_i - \lambda_{i-1} f_i) &= f_i'\beta_i - (f_i'\lambda_{i-1}) f_i \\ &= \beta_{i-1} f_i - (\beta_{i-1} - \lambda_{i-2} f_{i-1}) f_i = 0 \end{aligned}$$

as desired. This finishes the proof of (a).

In the situation of (a), dualization gives maps $\alpha_n^* : F_n'^* \to F_n^*$ forming a chain map between the dual complexes. Therefore we have induced homomorphisms on cohomology $\alpha^* : H^n(F';G) \to H^n(F;G)$. These do not depend on the choice of α_n's since any other choices α_n' are chain homotopic, say via chain homotopies λ_n, and then α_n^* and $\alpha_n'^*$ are chain homotopic via λ_n^* since the dual of the relation $\alpha_i - \alpha_i' = f_{i+1}'\lambda_i + \lambda_{i-1} f_i$ is $\alpha_i^* - \alpha_i'^* = \lambda_i^* f_{i+1}'^* + f_i^* \lambda_{i-1}^*$.

The induced homomorphisms $\alpha^* : H^n(F';G) \to H^n(F;G)$ satisfy $(\beta\alpha)^* = \alpha^*\beta^*$ for a composition $H \xrightarrow{\alpha} H' \xrightarrow{\beta} H''$ with a free resolution F'' of H'' also given, since one can choose the compositions $\beta_n\alpha_n$ of extensions α_n of α and β_n of β as an extension of $\beta\alpha$. In particular, if we take α to be an isomorphism and β to be its inverse, with $F'' = F$, then $\alpha^*\beta^* = (\beta\alpha)^* = \mathbb{1}$, the latter equality coming from the obvious extension of $\mathbb{1} : H \to H$

by the identity map of F. The same reasoning shows $\beta^*\alpha^* = \mathbb{1}$, so α^* is an isomorphism. Finally, if we specialize further, taking α to be the identity but with two different free resolutions F and F', we get a canonical isomorphism $\mathbb{1}^* : H^n(F;G) \to H^n(F';G)$. □

The group $H^n(F;G)$ is denoted $Ext^n(H,G)$. It is defined for arbitrary abelian groups H since a free resolution of H always exists. Namely, choose a set of generators for H and let F_0 be a free abelian group with basis in one-to-one correspondence with these generators, so we have a surjective homomorphism $f_0 : F_0 \to H$ sending the basis elements to the chosen generators. Now repeat the process with $Kerf_0$ in place of H, constructing a homomorphism $f_1 : F_1 \to F_0$ sending a basis for a free abelian group F_1 onto generators for $Kerf_0$. And similarly for $Kerf_1$, etc. This general construction is not the most efficient one possible since in fact $Kerf_0$ is free, being a subgroup of a free abelian group, so there is always a free resolution $0 \to F_1 \to F_0 \to H \to 0$ with $F_1 = Kerf_0$.

Since there is a free resolution $0 \to F_1 \to F_0 \to H \to 0$, it follows that $Ext^n(H,G) = 0$ for $n > 1$, and then $Ext^1(H,G)$ is abbreviated to $Ext(H,G)$. The 'Ext' notation arises from the fact that $Ext(H,G)$ has an interpretation as the set of isomorphism classes of extensions of G by H, i.e., short exact sequences $0 \to G \to J \to H \to 0$, with a natural definition of isomorphism between such exact sequences. This is explained in books on homological algebra, for example [K.Brown], [Hilton-Stammbach], or [MacLane 63]. This interpretation of $Ext(H,G)$ is seldom needed in algebraic topology, however.

By the Exercise stated earlier, exactness of $F_1 \to F_0 \to H \to 0$ implies exactness of the dual $F_1^* \leftarrow F_0^* \leftarrow H^* \leftarrow 0$. This means that $H^0(F;G)$ as defined above is zero. Rather than having $Ext^0(H,G)$ be automatically zero, it is better to define $H^n(F;G)$ as the n^{th} homology group of the complex $\cdots \leftarrow F_1^* \leftarrow F_0^* \leftarrow 0$ with the term H^* omitted. This makes $Ext^0(H,G) = H^0(F;G) = H^* = Hom(H,G)$, by the exactness of $F_1^* \leftarrow F_0^* \leftarrow H^* \leftarrow 0$.

Summarizing, we have established the following algebraic result:

Theorem 3.2. *If a chain complex C of free abelian groups has homology groups $H_n(C)$, then the cohomology groups $H^n(C;G)$ of the cochain complex $Hom(C_n,G)$ are determined by split exact sequences*

$$0 \longrightarrow Ext(H_{n-1}(C),G) \longrightarrow H^n(C;G) \xrightarrow{h} Hom(H_n(C),G) \longrightarrow 0$$ □

This is known as the **universal coefficient theorem for cohomology** because it is formally analogous to the universal coefficient theorem for homology in §3.4a, which expresses homology with arbitrary coefficients in terms of homology with $\mathbb{Z}$ coefficients.

Computing $Ext(H,G)$ for finitely generated H is not difficult using the following three properties:

$Ext(H\oplus H',G) \approx Ext(H,G)\oplus Ext(H',G)$

$Ext(H,G) = 0$ if H is free.

$Ext(\mathbb{Z}_n,G) \approx G/nG$

The first of these comes from using the direct sum of free resolutions of H and H' as a free resolution for $H\oplus H'$. If H is free, the free resolution $0 \to H \to H \to 0$ yields the second property, while the third comes from dualizing the free resolution

$$0 \longrightarrow \mathbb{Z} \xrightarrow{n} \mathbb{Z} \longrightarrow \mathbb{Z}_n \longrightarrow 0$$

to produce an exact sequence

$$\begin{array}{ccccccccc} 0 & \longleftarrow & Ext(\mathbb{Z}_n,G) & \longleftarrow & Hom(\mathbb{Z},G) & \xleftarrow{n} & Hom(\mathbb{Z},G) & \longleftarrow & Hom(\mathbb{Z}_n,G) \longleftarrow 0 \\ & & \| & & \| & & \| & & \\ & & G/nG & \longleftarrow & G & \xleftarrow{n} & G & & \end{array}$$

In particular, these three properties imply that $Ext(H,\mathbb{Z})$ is isomorphic to the torsion subgroup of H if H is finitely generated. Since $Hom(H,\mathbb{Z})$ is isomorphic to the free part of H if H is finitely generated, we have:

Corollary 3.3. *If the homology groups H_n and H_{n-1} of a chain complex C of free abelian groups are finitely generated, with torsion subgroups $T_n \subset H_n$ and $T_{n-1} \subset H_{n-1}$, then $H^n(C;\mathbb{Z}) \approx (H_n/T_n)\oplus T_{n-1}$.* □

It is useful in many situations to know that the short exact sequences in the universal coefficient theorem are natural, i.e., a chain map α between chain complexes C and C' of free abelian groups induces a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & Ext(H_{n-1}(C),G) & \longrightarrow & H^n(C;G) & \xrightarrow{h} & Hom(H_n(C),G) & \longrightarrow & 0 \\ & & \uparrow{\scriptstyle (\alpha_*)^*} & & \uparrow{\scriptstyle \alpha^*} & & \uparrow{\scriptstyle (\alpha_*)^*} & & \\ 0 & \longrightarrow & Ext(H_{n-1}(C'),G) & \longrightarrow & H^n(C';G) & \xrightarrow{h} & Hom(H_n(C'),G) & \longrightarrow & 0 \end{array}$$

This is apparent if one just thinks about the construction; one obviously obtains a map between the short exact sequences (iv) containing $Ker(i_n^*)$ and $Coker(i_{n-1}^*)$, the identification $Ker(i_n^*) = Hom(H_n(C),G)$ is certainly natural, and the proof of Lemma 3.1 shows that $Ext(H,G)$ depends naturally on H.

However, the splitting in the universal coefficient theorem is not natural since it depends on the choice of the projections $p\colon C_n \to Z_n$. An exercise at the end of the section gives a topological example showing that the splitting in fact cannot be natural.

The naturality property together with the five-lemma proves:

Corollary 3.4. *If a chain map between chain complexes of free abelian groups induces an isomorphism on homology groups, then it induces an isomorphism on cohomology groups with any coefficient group* G. □

One could attempt to generalize the algebraic machinery of the universal coefficient theorem by replacing abelian groups with modules over a chosen ring R and Hom by Hom_R, the R-module homomorphisms. The key fact about abelian groups that was needed was that subgroups of free abelian groups are free. Submodules of free R-modules are free if R is a principal ideal domain, so in this case the generalization is automatic. One obtains natural split short exact sequences

$$0 \longrightarrow Ext_R(H_{n-1}(C), G) \longrightarrow H^n(C;G) \xrightarrow{h} Hom_R(H_n(C), G) \longrightarrow 0$$

where C is a chain complex of free R-modules, with boundary maps R-module homomorphisms, and the coefficient group G is also an R-module. If R is a field, for example, then R-modules are always free and so the Ext_R term is always zero.

Even without any hypotheses on R the construction of Ext^n works, and one obtains groups $Ext^n_R(H, G)$ associated to R-modules G and H. For sufficiently complicated rings R these groups $Ext^n_R(H, G)$ can be nonzero for $n > 1$. In certain more advanced topics in algebraic topology these Ext^n_R groups play an essential role.

Cohomology of Spaces

Now we return to topology. Given a space X and an abelian group G, we define the group $C^n(X;G)$ of **singular n-cochains with coefficients in G** to be the dual group $Hom(C_n(X), G)$ of the singular chain group $C_n(X)$. Thus an n-cochain $\varphi \in C^n(X;G)$ assigns to each singular n-simplex $\sigma : \Delta^n \to X$ a value $\varphi(\sigma) \in G$. Since the singular n-simplices form a basis for $C_n(X)$, these values can be chosen arbitrarily, so n-cochains are the same as functions from singular n-simplices to G.

The **coboundary map** $\delta : C^n(X;G) \to C^{n+1}(X;G)$ is the dual ∂^*, so for a cochain $\varphi \in C^n(X;G)$, its coboundary $\delta(\varphi)$ is the composition $C_{n+1}(X) \xrightarrow{\partial} C_n(X) \xrightarrow{\varphi} G$. This means that for a singular $(n+1)$-simplex $\sigma : \Delta^{n+1} \to X$ we have

$$\delta\varphi(\sigma) = \sum_i (-1)^i \varphi(\sigma \,|\, [v_0, \cdots, \widehat{v}_i, \cdots, v_{n+1}])$$

It is automatic that $\delta^2 = 0$ since δ^2 is the dual of $\partial^2 = 0$. Therefore we can define the **cohomology group $H^n(X;G)$ with coefficients in G** to be the quotient $Ker(\delta)/Im(\delta)$ at $C^n(X;G)$ in the cochain complex

$$\cdots \longleftarrow C^{n+1}(X;G) \xleftarrow{\delta} C^n(X;G) \xleftarrow{\delta} C^{n-1}(X;G) \longleftarrow \cdots \longleftarrow C^0(X;G) \longleftarrow 0$$

Elements of $Ker(\delta)$ are **cocycles**, and elements of $Im(\delta)$ are **coboundaries**. For a cochain φ to be a cocycle means that $\delta(\varphi) = \varphi\partial = 0$, i.e., φ vanishes on boundaries.

Since the chain groups $C_n(X)$ are free, the algebraic universal coefficient theorem takes on the topological guise of split short exact sequences

$$0 \longrightarrow Ext(H_{n-1}(X), G) \longrightarrow H^n(X; G) \longrightarrow Hom(H_n(X), G) \longrightarrow 0$$

which describe how cohomology groups with arbitrary coefficients are determined purely algebraically by homology groups with $\mathbb{Z}$ coefficients.

There is a simpler statement when G is a field F, using the version of the algebraic universal coefficient theorem for modules over a principal ideal domain. The chain complex C consisting of the free F-modules $C_n = C_n(X) \otimes F$ has homology groups $H_n(X; F)$, and its dual complex $Hom_F(C_n, F)$ is the same as $Hom(C_n(X), F)$, so its cohomology groups are $H^n(C; F) = H^n(X; F)$. Since all F-modules are free, $Ext_F(H, G)$ is always zero and the universal coefficient theorem becomes an isomorphism:

$$H^n(X; F) \approx Hom_F(H_n(X; F), F)$$

Thus, with field coefficients, cohomology is the exact dual of homology. Note that when $F = \mathbb{Z}_p$ or $\mathbb{Q}$ we have $Hom_F(H, G) = Hom(H, G)$, the group homomorphisms.

For the remainder of this section we will go through the main features of singular homology and check that they extend without much difficulty to cohomology.

(1) Reduced Groups. Reduced groups $\widetilde{H}^n(X; G)$ can be defined by dualizing the augmented chain complex $\cdots \to C_0(X) \to \mathbb{Z} \to 0$, so $\widetilde{H}^n(X; G) = H^n(X; G)$ for $n > 0$ and $\widetilde{H}^0(X; G) \approx Hom(\widetilde{H}_0(X), G)$ by the universal coefficient theorem.

(2) Relative Groups and the Long Exact Sequence of a Pair. To define relative groups $H^n(X, A; G)$ for a pair (X, A) we first dualize the short exact sequence

$$0 \longrightarrow C_n(A) \xrightarrow{\ i\ } C_n(X) \xrightarrow{\ j\ } C_n(X, A) \longrightarrow 0$$

by applying $Hom(-, G)$ to get

$$0 \longleftarrow C^n(A; G) \xleftarrow{\ i^*\ } C^n(X; G) \xleftarrow{\ j^*\ } C^n(X, A; G) \longleftarrow 0$$

This is also exact since $C_n(X, A)$ is free, so the first sequence splits. In fact, $C_n(X, A)$ can be viewed as the subgroup of $C_n(X)$ with basis the singular n-simplices in X which

do not lie in A, so $C^n(X,A;G)$ is all functions from such simplices to G. This is identified naturally with the kernel of i^*, all the functions on n-simplices in X which vanish on n-simplices in A. Thus we can view the relative cochains as a subgroup of the absolute cochains, the ones which vanish on chains in A. Relative coboundary maps $\delta: C^n(X,A;G) \to C^{n+1}(X,A;G)$ are obtained as restrictions of the absolute δ's, so relative cohomology groups $H^n(X,A;G)$ are defined. The maps i^* and j^* commute with δ since i and j commute with ∂, so the preceding displayed short exact sequence of cochain groups is part of a short exact sequence of cochain complexes, giving rise to an associated long exact sequence of cohomology groups

$$\cdots \longrightarrow H^n(X,A;G) \xrightarrow{j^*} H^n(X;G) \xrightarrow{i^*} H^n(A;G) \xrightarrow{\delta} H^{n+1}(X,A;G) \longrightarrow \cdots$$

By similar reasoning one obtains the long exact sequence for a triple (X,A,B) and the long exact sequence for reduced cohomology groups of a pair.

(3) Induced Homomorphisms. A map $f: X \to Y$ induces chain maps $f_\sharp: C_n(X) \to C_n(Y)$ with dual cochain maps $f^\sharp: C^n(Y;G) \to C^n(X;G)$. The relation $f_\sharp\partial = \partial f_\sharp$ dualizes to $\delta f^\sharp = f^\sharp\delta$, so $f^\sharp$ induces homomorphisms $f^*: H^n(Y;G) \to H^n(X;G)$. In the relative case a map $f: (X,A) \to (Y,B)$ induces $f^*: H^n(Y,B;G) \to H^n(X,A;G)$ by the same reasoning, and in fact f induces a map between short exact sequences of cochain complexes, hence a map between long exact sequences of cohomology groups, with commuting squares. The properties $(fg)^\sharp = g^\sharp f^\sharp$ and $\mathbb{1}^\sharp = \mathbb{1}$ imply $(fg)^* = g^* f^*$ and $\mathbb{1}^* = \mathbb{1}$, so $X \mapsto H^n(X;G)$ and $(X,A) \mapsto H^n(X,A;G)$ are contravariant functors, the 'contra' indicating that induced maps go in the reverse direction.

The algebraic universal coefficient theorem applies also to relative cohomology, and there is a naturality statement: A map $f: (X,A) \to (Y,B)$ induces a commutative diagram

$$\begin{array}{ccccccc}
0 \longrightarrow & Ext(H_{n-1}(X,A),G) & \longrightarrow & H^n(X,A;G) & \xrightarrow{h} & Hom(H_n(X,A),G) & \longrightarrow 0 \\
& \uparrow{\scriptstyle (f_*)^*} & & \uparrow{\scriptstyle f^*} & & \uparrow{\scriptstyle (f_*)^*} & \\
0 \longrightarrow & Ext(H_{n-1}(Y,B),G) & \longrightarrow & H^n(Y,B;G) & \xrightarrow{h} & Hom(H_n(Y,B),G) & \longrightarrow 0
\end{array}$$

This follows from the naturality of the algebraic universal coefficient sequences since the vertical maps are induced by the chain maps $f_\sharp: C_n(X,A) \to C_n(Y,B)$. When the subspaces A and B are empty we obtain the absolute forms of these results.

(4) Homotopy Invariance. The statement is that if $f \simeq g: (X,A) \to (Y,B)$, then $f^* = g^*: H^n(Y,B) \to H^n(X,A)$. This is proved by direct dualization of the proof for homology. From the proof of Theorem 2.7 we have a chain homotopy P satisfying $g_\sharp - f_\sharp = \partial P + P\partial$.

This relation dualizes to $g^\sharp - f^\sharp = P^*\delta + \delta P^*$, so P^* is a chain homotopy between the maps $f^\sharp, g^\sharp : C^n(Y;G) \to C^n(X;G)$. This restricts also to a chain homotopy between $f^\sharp$ and $g^\sharp$ on relative cochains, the cochains which are zero on singular simplices in the subspaces B and A. Since $f^\sharp$ and $g^\sharp$ are chain homotopic, they induce the same homomorphism $f^* = g^*$ on cohomology.

(5) Excision. For cohomology this says that for subspaces $Z \subset A \subset X$ with the closure of Z contained in the interior of A, the inclusion $i : (X - Z, A - Z) \hookrightarrow (X, A)$ induces isomorphisms $i^* : H^n(X, A; G) \to H^n(X - Z, A - Z; G)$ for all n. This follows from the corresponding result for homology by the naturality of the universal coefficient theorem and the five-lemma. Alternatively, if one wishes to avoid appealing to the universal coefficient theorem, the proof of excision for homology dualizes easily to cohomology by the following argument. In the proof for homology there were chain maps $\rho : C_n(X) \to C_n(A + B)$ and $\iota : C_n(A + B) \to C_n(X)$ such that $\rho\iota = \mathbb{1}$ and $\mathbb{1} - \iota\rho = \partial D + D\partial$ for a chain homotopy D. Dualize all this by taking $Hom(-, G)$, and note that since ρ, ι, and D take chains in A to chains in A, the dual maps take homomorphisms vanishing on chains in A to homomorphisms vanishing on chains in A. This means we have maps ρ^* and ι^* between $C^n(A + B, A; G)$ and $C^n(X, A; G)$ which induce isomorphisms on cohomology since $\iota^*\rho^* = \mathbb{1}$ and $\mathbb{1} - \rho^*\iota^* = D^*\delta + \delta D^*$. Furthermore, there is an obvious identification of $C^n(A + B, A; G)$ with $C^n(B, A \cap B; G)$. Thus we get isomorphisms $H^n(B, A \cap B; G) \approx H^n(X, A)$ induced by i.

(6) Axioms for Cohomology. These are exactly dual to the axioms for homology. Restricting attention to CW complexes again, a (reduced) **cohomology theory** is a sequence of contravariant functors $\widetilde{h}^n$ from CW complexes to abelian groups, together with natural coboundary homomorphisms $\delta : \widetilde{h}^n(A) \to \widetilde{h}^{n+1}(X/A)$ for CW pairs (X, A), satisfying the following axioms:

(1) If $f \simeq g : X \to Y$, then $f^* = g^* : \widetilde{h}^n(Y) \to \widetilde{h}^n(X)$.

(2) For each CW pair (X, A) there is a long exact sequence

$$\cdots \xrightarrow{\delta} \widetilde{h}^n(X/A) \xrightarrow{q^*} \widetilde{h}^n(X) \xrightarrow{i^*} \widetilde{h}^n(A) \xrightarrow{\delta} \widetilde{h}^{n+1}(X/A) \xrightarrow{q^*} \cdots$$

where i is the inclusion and q is the quotient map.

(3) For a wedge sum $X = \bigvee_\alpha X_\alpha$ with inclusions $i_\alpha : X_\alpha \hookrightarrow X$, the product map $\prod_\alpha i_\alpha^* : \widetilde{h}^n(X) \to \prod_\alpha \widetilde{h}^n(X_\alpha)$ is an isomorphism for each n.

We have already seen that the first two axioms hold for singular cohomology, and the third is easy to check by the same argument as for homology. Note that for cohomology this

axiom involves direct product, rather than the direct sum which appeared in the homology version.

Similar remarks concerning the relation between reduced and unreduced theories can be made for cohomology as were made for homology.

(7) Simplicial Cohomology. If X is a Δ-complex and $A \subset X$ is a subcomplex, then the simplicial chain groups $\Delta_n(X, A)$ dualize to the simplicial cochain groups $\Delta^n(X, A; G) = Hom(\Delta_n(X, A), G)$, and the resulting cohomology groups are by definition the simplicial cohomology groups $H^n_\Delta(X, A; G)$. Since the inclusions $\Delta_n(X, A) \subset C_n(X, A)$ induce isomorphisms $H^\Delta_n(X, A) \approx H_n(X, A)$, Corollary 3.4 implies that the dual maps $C^n(X, A; G) \to \Delta^n(X, A; G)$ also induce isomorphisms $H^n(X, A; G) \approx H^n_\Delta(X, A; G)$.

(8) Cellular Cohomology. For a CW complex X this is defined via the cellular cochain complex formed by the horizontal sequence in the following diagram, where coefficients in a given group G are understood and the cellular coboundary maps d_n are the compositions $\delta_n j_n$, making the triangles commute. Note that $d_n d_{n-1} = 0$ since $j_n \delta_{n-1} = 0$.

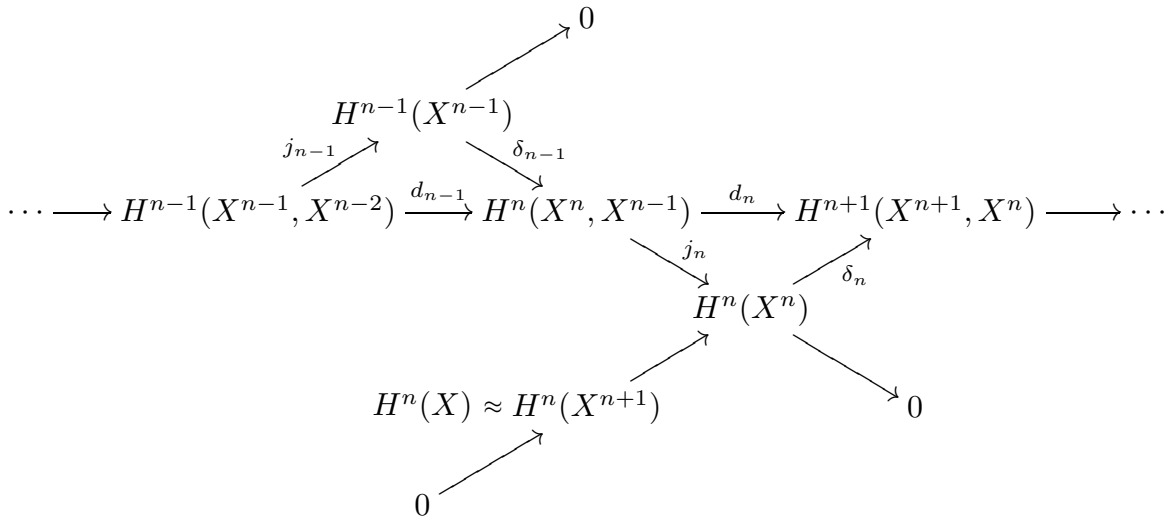

Theorem 3.5. *$H^n(X; G) \approx Ker(d_n)/Im(d_{n-1})$. Furthermore, the cellular cochain complex $\{H^n(X^n, X^{n-1}; G), d_n\}$ is isomorphic to the dual of the cellular chain complex, obtained by applying $Hom(-, G)$.*

Proof: The universal coefficient theorem implies that $H^k(X^n, X^{n-1}; G) = 0$ for $k \neq n$. The long exact sequence of the pair (X^n, X^{n-1}) then gives isomorphisms $H^k(X^n; G) \approx H^k(X^{n-1}; G)$ for $k \neq n,\ n-1$. Hence by induction on n we obtain $H^k(X^n; G) = 0$ if $k > n$. Thus the diagonal sequences in the preceding diagram are exact. The universal coefficient theorem also gives $H^k(X, X^{n+1}; G) = 0$ for $k \leq n+1$, so $H^n(X; G) \approx$

$H^n(X^{n+1};G)$. The diagram then yields isomorphisms

$$H^n(X;G) \approx H^n(X^{n+1};G) \approx Ker\delta_n \approx Ker(d_n)/Im(\delta_{n-1}) \approx Ker(d_n)/Im(d_{n-1})$$

For the second statement in the theorem we have the diagram

$$\begin{array}{ccccc} H^k(X^k,X^{k-1};G) & \longrightarrow & H^k(X^k;G) & \xrightarrow{\delta} & H^{k+1}(X^{k+1},X^k;G) \\ \downarrow h & & \downarrow h & & \downarrow h \\ Hom(H_k(X^k,X^{k-1},G) & \longrightarrow & Hom(H_k(X^k),G) & \xrightarrow{\partial^*} & Hom(H_{k+1}(X^{k+1},X^k),G) \end{array}$$

The first square commutes by naturality of h and the second square commutes because coboundary maps are defined as duals of boundary maps. The first and third vertical maps are isomorphisms by the universal coefficient theorem, and the result follows. □

(9) Mayer-Vietoris Sequences. In the absolute case these take the form

$$\cdots \longrightarrow H^n(X;G) \xrightarrow{\Psi} H^n(A;G) \oplus H^n(B;G) \xrightarrow{\Phi} H^n(A\cap B;G) \longrightarrow H^{n+1}(X;G) \longrightarrow \cdots$$

where X is the union of the interiors of A and B. This is the long exact sequence associated to the short exact sequence of cochain complexes

$$0 \longrightarrow C^n(A+B,G) \xrightarrow{\psi} C^n(A;G) \oplus C^n(B;G) \xrightarrow{\varphi} C^n(A\cap B;G) \longrightarrow 0$$

Here $C^n(A+B;G)$ is the dual of the subgroup $C_n(A+B) \subset C_n(X)$ consisting of sums of singular n-simplices lying in A or in B. The inclusion $C_n(A+B) \subset C_n(X)$ is a chain homotopy equivalence by Proposition 2.14, so the dual restriction map $C^n(X;G) \to C^n(A+B;G)$ is also a chain homotopy equivalence, hence induces an isomorphism on cohomology. The map ψ has coordinates the two restrictions to A and B, and φ take the difference of the restrictions to $A\cap B$, so it is obvious that φ is onto with kernel the image of ψ.

There is a relative Mayer-Vietoris sequence

$$\cdots \longrightarrow H^n(X,Y;G) \longrightarrow H^n(A,C;G) \oplus H^n(B,D;G) \longrightarrow H^n(A\cap B, C\cap D;G) \longrightarrow \cdots$$

for a pair $(X,Y) = (A\cup B, C\cup D)$ with $C \subset A$ and $D \subset B$ such that X is the union of the interiors of A and B while Y is the union of the interiors of C and D. To derive this, consider first the map of short exact sequences of cochain complexes:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C^n(X,Y;G) & \longrightarrow & C^n(X;G) & \longrightarrow & C^n(Y;G) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & C^n(A+B,C+D;G) & \longrightarrow & C^n(A+B;G) & \longrightarrow & C^n(C+D;G) & \longrightarrow & 0 \end{array}$$

Here $C^n(A+B, C+D; G)$ is defined as the kernel of the restriction map $C^n(A+B; G) \longrightarrow C^n(C+D; G)$, so the second sequence is exact. The vertical maps are restriction. The second and third of these induce isomorphisms on cohomology, as we have seen, so by the five-lemma the first vertical map also induces isomorphisms on cohomology. The relative Mayer-Vietoris sequence is then the long exact sequence associated to the short exact sequence of cochain complexes

$$0 \longrightarrow C^n(A+B, C+D; G) \xrightarrow{\psi} C^n(A, C; G) \oplus C^n(B, D; G) \xrightarrow{\varphi} C^n(A \cap B, C \cap D; G) \longrightarrow 0$$

This is exact since it is the dual of the short exact sequence

$$0 \longrightarrow C_n(A \cap B, C \cap D) \longrightarrow C_n(A, C) \oplus C_n(B, D) \longrightarrow C_n(A + B, C + D) \longrightarrow 0$$

from §2.2, which splits since $C_n(A + B, C + D)$ is free with basis the singular n-simplices in A or B which do not lie in C or in D.

Exercises

1. Show that $Ext(H, G)$ is a contravariant functor of H for fixed G, and a covariant functor of G for fixed H.

2. Show that the maps $G \to G$ and $H \to H$ which multiply each element by the integer n each induce multiplication by n in $Ext(H, G)$.

3. Show that $H^0(X; G) \approx Hom(H_0(X), G)$ directly from the definitions.

4. Directly from the definitions, compute the simplicial cohomology groups of $S^1 \times S^1$, $\mathbb{R}\mathrm{P}^2$, and the Klein bottle, with $\mathbb{Z}$ and $\mathbb{Z}_2$ coefficients, using the Δ-complex structures given in §2.1.

5. Show that if A is a retract of X then $H^n(X; G) \approx H^n(A; G) \oplus H^n(X, A; G)$.

6. Show that if $f : S^n \to S^n$ has degree d then the induced map $f^* : H^n(S^n; G) \to H^n(S^n; G)$ is multiplication by d.

7. Let X be a Moore space $M(\mathbb{Z}_m, n)$ obtained from S^n by attaching a cell e^{n+1} by a map of degree m. Show the quotient map $X \to X/S^n = S^{n+1}$ induces the trivial map on $\widetilde{H}_i(-; \mathbb{Z})$ for all i, but not on $H^{n+1}(-; \mathbb{Z})$, while the inclusion $S^n \hookrightarrow X$ induces the trivial map on $\widetilde{H}^i(-; \mathbb{Z})$ for all i, but not on $H_n(-; \mathbb{Z})$. Deduce that the splitting in the universal coefficient theorem for cohomology cannot be natural.

8. Show $H^k(X, X^n; G) = 0$ if X is a CW complex and $k \leq n$, by using the cohomology version of the second proof of the corresponding result for homology in Lemma 2.25.

2. Cup Product

In the introduction to this chapter we sketched a definition of cup product in terms of another product called cross product. However, since it takes a fair bit of effort to set up the cross product rigorously, this will not actually be the definition of cup product we give first. Instead we start with a more elementary definition involving an explicit formula with simplices. But one should not be misled into thinking the explicit formula is essential to the existence of cup products, or that the basic properties of cup product depend crucially on the formula. The formula is merely a convenient starting point, and useful for working out a few examples without much machinery.

To define the cup product in cohomology we take coefficients in a ring R, the most common choices being $\mathbb{Z}$, $\mathbb{Z}_p$, and $\mathbb{Q}$. Given cochains $\varphi \in C^k(X;R)$ and $\psi \in C^\ell(X;R)$, their **cup product** is the cochain $\varphi \smile \psi \in C^{k+\ell}(X;R)$ whose value on a singular simplex $\sigma : \Delta^{k+\ell} \to X$ is

$$(\varphi \smile \psi)(\sigma) = \varphi(\sigma\big|[v_0, \cdots, v_k])\psi(\sigma\big|[v_k, \cdots, v_{k+\ell}])$$

where the right-hand side is the product in R. To see that this cup product of cochains induces a cup product of cohomology classes we need a formula relating it to the coboundary map:

Lemma 3.6. *$\delta(\varphi \smile \psi) = \delta\varphi \smile \psi + (-1)^k \varphi \smile \delta\psi$ for $\varphi \in C^k(X;R)$ and $\psi \in C^\ell(X;R)$.*

Proof: For $\sigma : \Delta^{k+\ell+1} \to X$ we have

$$(\delta\varphi \smile \psi)(\sigma) = \sum_{i=0}^{k+1} (-1)^i \varphi(\sigma\big|[v_0, \cdots, \widehat{v}_i, \cdots, v_{k+1}])\psi(\sigma\big|[v_{k+1}, \cdots, v_{k+\ell+1}])$$

$$(-1)^k(\varphi \smile \delta\psi)(\sigma) = \sum_{i=k}^{k+\ell+1} (-1)^i \varphi(\sigma\big|[v_0, \cdots, v_k])\psi(\sigma\big|[v_k, \cdots, \widehat{v}_i, \cdots, v_{k+\ell+1}])$$

Adding these two expressions, the last term of the first sum cancels the first term of the second sum, and the remaining terms are exactly $\delta(\varphi \smile \psi)(\sigma) = (\varphi \smile \psi)(\partial\sigma)$ since $\partial\sigma = \sum_{i=0}^{k+\ell+1} (-1)^i \sigma\big|[v_0, \cdots, \widehat{v}_i, \cdots, v_{k+\ell+1}]$. □

From the formula $\delta(\varphi \smile \psi) = \delta\varphi \smile \psi \pm \varphi \smile \delta\psi$ it is apparent that the cup product of two cocycles is again a cocycle. Also, the cup product of a cocycle and a coboundary, in either order, is a coboundary since $\varphi \smile \delta\psi = \pm\delta(\varphi \smile \psi)$ if $\delta\varphi = 0$, and $\delta\varphi \smile \psi = \delta(\varphi \smile \psi)$ if $\delta\psi = 0$. It follows that there is an induced cup product

$$H^k(X;R) \times H^\ell(X;R) \xrightarrow{\smile} H^{k+\ell}(X;R)$$

This is associative and distributive since cup product at the cochain level is. If R has an identity element 1, then the 0-cocycle taking the value 1 on each singular 0-simplex defines an element $1 \in H^0(X;R)$ which is an identity for $\smile$.

A cup product for simplicial cohomology can be defined by the same formula as for singular cohomology, so the canonical isomorphism between simplicial and singular cohomology respects cup products. Here are three examples of direct calculations of cup products using simplicial cohomology.

Example 1. Let M be the closed orientable surface of genus g with the Δ-complex structure shown in the first of the two figures below.

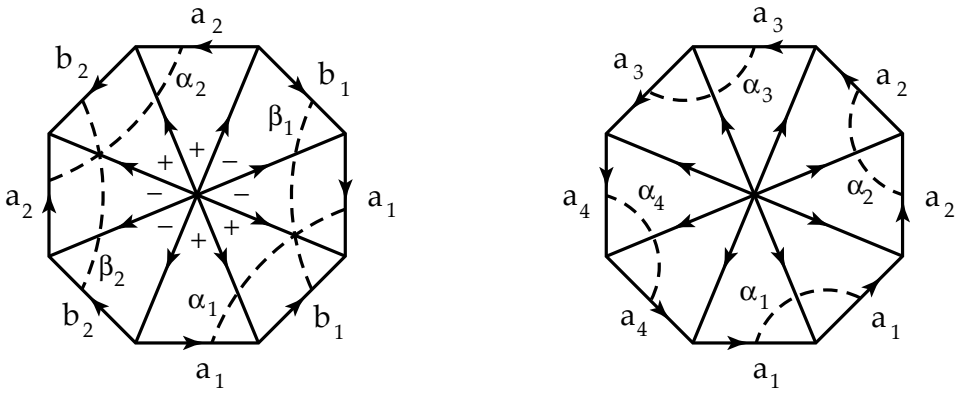

Taking $\mathbb{Z}$ coefficients, the cup product of interest is $H^1(M) \times H^1(M) \to H^2(M)$. A basis for $H_1(M)$ is formed by the edges a_i and b_i, as we showed in §2.2 when we computed the homology of M using cellular homology. We have $H^1(M) \approx Hom(H_1(M), \mathbb{Z})$ by cellular cohomology or the universal coefficient theorem, and the specified basis for $H_1(M)$ gives an isomorphism $Hom(H_1(M), \mathbb{Z}) \approx H_1(M)$ under which the cohomology class α_i dual to a_i assigns the value 1 to a_i and 0 to the other basis elements, and similarly for cohomology classes β_i dual to b_i. To represent α_i by a simplicial cocycle φ_i we need to choose values for φ_i on the edges radiating out from the central vertex in such a way that $\delta\varphi_i = 0$. A nice way to do this is to consider the arc labelled α_i in the figure, which represents a loop in M meeting a_i in one point and disjoint from all the other basis elements a_j and b_j. We define φ_i to have the value 1 on edges meeting the arc α_i and the value 0 on all other edges. Thus φ_i counts the number of intersections of each edge with the arc α_i. Then one calculates that $\delta\varphi_i = 0$ from the definition of δ. In similar fashion we obtain a cocycle ψ_i counting intersections with the arc β_i, and ψ_i represents the cohomology class β_i dual to b_i.

Now we can compute cup products by applying the definition. Keeping in mind the fact that the ordering of the vertices of each 2-simplex is compatible with the indicated ori-

entations of its edges, we see for example that $\varphi_1 \smile \psi_1$ takes the value 0 on all 2-simplices except the one with outer edge b_1 in the lower right part of the figure, where it takes the value 1. Thus $\varphi_1 \smile \psi_1$ takes the value 1 on the 2-chain c formed by the sum of all the 2-simplices with the signs indicated in the figure. It is an easy calculation that $\partial c = 0$, so c represents an element of $H_2(M)$. The fact that $(\varphi_1 \smile \psi_1)(c)$ is a generator of $\mathbb{Z}$ implies both that c represents a generator of $H_2(M) \approx \mathbb{Z}$ and that $\varphi_1 \smile \psi_1$ represents the dual generator γ of $H^2(M) \approx Hom(H_2(M), \mathbb{Z}) \approx \mathbb{Z}$. Thus $\alpha_1 \smile \beta_1 = \gamma$. In similar fashion one computes:

$$\alpha_i \smile \beta_j = \left\{ \begin{array}{ll} \gamma, & i = j \\ 0, & i \neq j \end{array} \right\} = -(\beta_i \smile \alpha_j), \qquad \alpha_i \smile \alpha_j = 0, \qquad \beta_i \smile \beta_j = 0$$

These relations determine the cup product $H^1(M) \times H^1(M) \to H^2(M)$ completely since cup product is distributive. Notice that cup product is not commutative in this example since $\alpha_i \smile \beta_i = -(\beta_i \smile \alpha_i)$. We will show in Theorem 3.7 below that this is the worst that can happen: Cup product is commutative up to a sign depending only on dimension.

One can see in this example that nonzero cup products of distinct classes α_i and β_j occur precisely when the corresponding loops α_i and β_j intersect. This is also true for the cup product of α_i or β_i with itself if we allow ourselves to take two copies of the corresponding loop and deform one of them to be disjoint from the other. This intersection behavior could hardly be mere accident, and in fact it is an aspect of Poincaré duality.

Example 2. The closed non-orientable surface N of genus g can be treated in similar fashion if we use $\mathbb{Z}_2$ coefficients. Using the Δ-complex structure pictured in the second figure above, the edges a_i give a basis for $H_1(N; \mathbb{Z}_2)$, and the dual basis elements $\alpha_i \in H^1(N; \mathbb{Z}_2)$ can be represented by cocycles with values given by counting intersections with the arcs labelled α_i in the figure. Then one computes that $\alpha_i \smile \alpha_i$ is the nonzero element of $H^2(N; \mathbb{Z}_2) \approx \mathbb{Z}_2$ and $\alpha_i \smile \alpha_j = 0$ for $i \neq j$. In particular, when $g = 1$ we have $N = \mathbb{R}\mathrm{P}^2$, and the cup product of a generator of $H^1(\mathbb{R}\mathrm{P}^2; \mathbb{Z}_2)$ with itself is a generator of $H^2(\mathbb{R}\mathrm{P}^2; \mathbb{Z}_2)$.

The remark following Example 1 concerning Poincaré duality applies here also, but with the following difference: When one tries to deform a second copy of the loop α_i in the present example to be disjoint from the original copy, the best one can do is make it intersect the original in one point. This reflects the fact that $\alpha_i \smile \alpha_i$ is now nonzero.

Example 3. Let X be the 2-dimensional CW complex obtained by attaching a 2-cell to S^1 by the degree m map $S^1 \to S^1$, $z \mapsto z^m$. Using cellular cohomology, or cellular homology and the universal coefficient theorem, we see that $H^n(X; \mathbb{Z})$ consists of a $\mathbb{Z}$ for $n = 0$

and a $\mathbb{Z}_m$ for $n = 2$, so the cup product structure with $\mathbb{Z}$ coefficients is uninteresting. However with $\mathbb{Z}_m$ coefficients we have $H^i(X;\mathbb{Z}_m) \approx \mathbb{Z}_m$ for $i = 0, 1, 2$ so there is the possibility that the cup product of two 1-dimensional classes can be nontrivial.

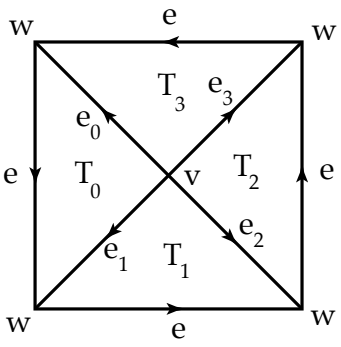

To obtain a Δ-complex structure on X, take a regular m-gon subdivided into m triangles T_i around a central vertex v, as shown in the figure for the case $m = 4$, then identify all the outer edges by rotations of the m-gon. This gives X a Δ-complex structure with 2 vertices, $m+1$ edges, and m 2-simplices. A generator α of $H^1(X;\mathbb{Z}_m)$ is represented by a cocycle φ assigning the value 1 to the edge e, which generates $H_1(X)$. The condition that φ be a cocycle means that $\varphi(e_i) + \varphi(e) = \varphi(e_{i+1})$ for all i, subscripts being taken $\bmod m$. So we may take $\varphi(e_i) = i \in \mathbb{Z}_m$. Hence $(\varphi \smile \varphi)(T_i) = \varphi(e_i)\varphi(e) = i$. The map $h : H^2(X;\mathbb{Z}_m) \to Hom(H_2(X;\mathbb{Z}_m), \mathbb{Z}_m)$ is an isomorphism since $\sum_i T_i$ is a generator of $H_2(X;\mathbb{Z}_m)$ and there are 2-cocycles taking the value 1 on $\sum_i T_i$, for example the cocycle taking the value 1 on one T_i and 0 on all the others. The cocycle $\varphi \smile \varphi$ takes the value $0 + 1 + \cdots + (m-1)$ on $\sum_i T_i$, hence represents $0 + 1 + \cdots + (m-1)$ times a generator β of $H^2(X;\mathbb{Z}_m)$. In $\mathbb{Z}_m$ the sum $0+1+\cdots+(m-1)$ is 0 if m is odd and k if $m = 2k$ since the terms 1 and $m-1$ cancel, 2 and $m-2$ cancel, etc. Thus, writing α^2 for $\alpha \smile \alpha$, we have $\alpha^2 = 0$ if m is odd and $\alpha^2 = k\beta$ if $m = 2k$. For example, if $m = 2$, X is $\mathbb{R}\mathrm{P}^2$ and $\alpha^2 = \beta$ in $H^2(\mathbb{R}\mathrm{P}^2;\mathbb{Z}_2)$, reproving the case $g = 1$ in Example 2.

The same formula which defines the cup product $H^k(X;R) \times H^\ell(X;R) \to H^{k+\ell}(X;R)$ also gives relative products

$$H^k(X;R) \times H^\ell(X,A;R) \xrightarrow{\smile} H^{k+\ell}(X,A;R)$$
$$H^k(X,A;R) \times H^\ell(X;R) \xrightarrow{\smile} H^{k+\ell}(X,A;R)$$
$$H^k(X,A;R) \times H^\ell(X,A;R) \xrightarrow{\smile} H^{k+\ell}(X,A;R)$$

since if φ or ψ vanishes on chains in A then so does $\varphi \smile \psi$. When A and B are open subsets of X, or when A and B are subcomplexes of the CW complex X, there is a more general relative cup product

$$H^k(X,A;R) \times H^\ell(X,B;R) \xrightarrow{\smile} H^{k+\ell}(X, A \cup B; R)$$

obtained as follows. The absolute cup product restricts to a cup product $C^k(X,A;R) \times C^\ell(X,B;R) \to C^{k+\ell}(X, A+B; R)$ where $C^n(X, A+B; R)$ is the subgroup of $C^n(X;R)$

consisting of cochains vanishing on chains in A and on chains in B. The inclusions $C^n(X, A\cup B; R) \subset C^n(X, A+B; R)$ induce isomorphisms on cohomology, as we saw in the derivation of Mayer-Vietoris sequences in §3.1. Therefore the cup product $C^k(X, A; R) \times C^\ell(X, B; R) \to C^{k+\ell}(X, A + B; R)$ induces a cup product $H^k(X, A; R) \times H^\ell(X, B; R) \to H^{k+\ell}(X, A \cup B; R)$.

The Cohomology Ring

Since cup product is associative and distributive, it is natural to try to make it the multiplication in a ring structure on the cohomology groups of a space X. This is easy to do if we simply define $H^*(X; R)$ to be the direct sum of the groups $H^n(X; R)$, and then cup product clearly makes $H^*(X; R)$ into a ring, with identity if R has an identity. Similarly, $H^*(X, A; R)$ is a ring via the relative cup product. Taking scalar multiplication by elements of R into account, these rings can also be regarded as R-algebras.

For example, the calculations in either Example 2 or Example 3 above show that $H^*(\mathbb{R}\mathrm{P}^2; \mathbb{Z}_2)$ consists of the polynomials $a_0 + a_1\alpha + a_2\alpha^2$ with coefficients $a_i \in \mathbb{Z}_2$, so $H^*(\mathbb{R}\mathrm{P}^2; \mathbb{Z}_2)$ is the quotient $\mathbb{Z}_2[\alpha]/(\alpha^3)$ of the polynomial ring $\mathbb{Z}_2[\alpha]$ by the ideal generated by α^3. More generally, we will see later in this section that $H^*(\mathbb{R}\mathrm{P}^n; \mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha]/(\alpha^{n+1})$ and $H^*(\mathbb{R}\mathrm{P}^\infty; \mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha]$. Similar statements hold also for $\mathbb{C}\mathrm{P}^n$ and $\mathbb{C}\mathrm{P}^\infty$ using $\mathbb{Z}$ coefficients.

Adding cohomology classes of different dimensions to form $H^*(X; R)$ is mainly a matter of formal convenience and has little topological significance. One always regards the cohomology ring as a **graded ring**: a ring with a decomposition as a sum $\bigoplus_{k\ge 0} A_k$ of additive subgroups A_k such that the multiplication takes $A_k \times A_\ell$ to $A_{k+\ell}$.

The natural question of whether the cohomology ring is commutative is answered by the following:

Theorem 3.7. *With a commutative coefficient ring R we have $\alpha \smile \beta = (-1)^{k\ell}\beta \smile \alpha$ for all $\alpha \in H^k(X, A; R)$ and $\beta \in H^\ell(X, A; R)$.*

Taking $\alpha = \beta$, this implies in particular that if α is an element of $H^k(X, A; R)$ with k odd, then $2\alpha^2 = 0$ in $H^{2k}(X, A; R)$. Hence if $H^{2k}(X, A; R)$ has no elements of order two, then $\alpha^2 = 0$. For example, if X is the 2-complex obtained by attaching a disk to S^1 by a map of degree m, as in Example 3 above, then we can deduce that the square of a generator of $H^1(X; \mathbb{Z}_m)$ is zero if m is odd, and is either zero or the unique element of $H^2(X; \mathbb{Z}_m) \approx \mathbb{Z}_m$ of order two if m is even. As we saw there, the square is in fact nonzero when m is even.

A graded ring satisfying the commutativity property of the theorem is usually called simply **commutative** in the context of algebraic topology, in spite of the potential for misunderstanding. In the older literature one finds the more accurate but slightly awkward terms graded-commutative, anti-commutative, and skew-commutative.

Proof of 3.7: Let us first treat the case $A = \varnothing$, for simplicity. Given cochains $\varphi \in C^k(X;R)$ and $\psi \in C^\ell(X,R)$, one can see from the definition that the cup products $\varphi \smile \psi$ and $\psi \smile \varphi$ differ only by a permutation of the vertices of $\Delta^{k+\ell}$. The idea of the proof is to study a particularly nice permutation of vertices, namely totally reversing their order, replacing $[v_0,\cdots,v_n]$ by $[v_n,\cdots,v_0]$. This has the nice feature of also reversing the ordering of vertices in any face.

For a singular n-simplex $\sigma\colon[v_0,\cdots,v_n] \to X$, let $\overline{\sigma}$ be the singular n-simplex obtained by preceding σ by the linear homeomorphism of $[v_0,\cdots,v_n]$ reversing the order of the vertices. Thus $\overline{\sigma}(v_i) = \sigma(v_{n-i})$. This reversal of vertices is the product of $n+(n-1)+\cdots+1 = n(n+1)/2$ transpositions of adjacent vertices, each of which reverses orientation of the n-simplex, being a reflection across an $(n-1)$-dimensional hyperplane. So to take orientations into account we would expect that a sign $\varepsilon_n = (-1)^{n(n+1)/2}$ ought to be inserted. Hence we define a homomorphism $\rho\colon C_n(X) \to C_n(X)$ by $\rho(\sigma) = \varepsilon_n\overline{\sigma}$. We will show that ρ is a chain map chain homotopic to the identity, so ρ induces the identity on cohomology.

From this the theorem quickly follows. Namely, the formulas

$$\begin{aligned}(\rho^*\varphi \smile \rho^*\psi)(\sigma) &= \varphi(\varepsilon_k\sigma\big|[v_k,\cdots,v_0])\psi(\varepsilon_\ell\sigma\big|[v_{k+\ell},\cdots,v_k])\\ \rho^*(\psi\smile\varphi)(\sigma) &= \varepsilon_{k+\ell}\psi(\sigma\big|[v_{k+\ell},\cdots,v_k])\varphi(\sigma\big|[v_k,\cdots,v_0])\end{aligned}$$

show that $\varepsilon_k\varepsilon_\ell(\rho^*\varphi\smile\rho^*\psi) = \varepsilon_{k+\ell}\rho^*(\psi\smile\varphi)$, since we assume R is commutative. A trivial calculation gives $\varepsilon_{k+\ell} = (-1)^{k\ell}\varepsilon_k\varepsilon_\ell$, hence $(\rho^*\varphi\smile\rho^*\psi) = (-1)^{k\ell}\rho^*(\psi\smile\varphi)$. Since ρ is chain homotopic to the identity, the ρ^*'s disappear when we pass to cohomology classes, and so we obtain the desired formula $\alpha\smile\beta = (-1)^{k\ell}\beta\smile\alpha$.

The chain map property $\partial\rho = \rho\partial$ can be verified by calculating, for a singular n-simplex σ,

$$\begin{aligned}\partial\rho(\sigma) &= \varepsilon_n\sum_i(-1)^i\sigma\big|[v_n,\cdots,\widehat{v}_{n-i},\cdots,v_0]\\ \rho\partial(\sigma) &= \rho(\sum_i(-1)^i\sigma\big|[v_0,\cdots,\widehat{v}_i,\cdots,v_n])\\ &= \varepsilon_{n-1}\sum_i(-1)^{n-i}\sigma\big|[v_n,\cdots,\widehat{v}_{n-i},\cdots,v_0]\end{aligned}$$

which reduces us to the easily-checked identity $\varepsilon_n = (-1)^n\varepsilon_{n-1}$.

To define a chain homotopy between ρ and the identity we are motivated by the construction of the prism operator P in the proof that homotopic maps induce the same homomorphism on homology, in §2.1. This used a subdivision of $\Delta^n \times I$ into $(n+1)$-simplices with vertices v_i in $\Delta^n \times 0$ and w_i in $\Delta^n \times 1$, the vertex w_i lying directly above v_i. Using the same subdivision, and letting $\pi : \Delta^n \times I \to \Delta^n$ be the projection, we now define $P : C_n(X) \to C_{n+1}(X)$ by

$$P(\sigma) = \sum_i (-1)^i \varepsilon_{n-i}(\sigma\pi)\big|[v_0, \cdots, v_i, w_n, \cdots, w_i].$$

Thus the w-vertices are written in reverse order, and there is a compensating sign ε_{n-i}. One can view this formula as arising from the Δ-complex structure on $\Delta^n \times I$ in which the vertices are ordered $v_0, \cdots, v_n, w_n, \cdots, w_0$ rather than $v_0, \cdots, v_n, w_0, \cdots, w_n$.

To show $\partial P + P\partial = \rho - \mathbb{1}$ we first calculate ∂P, leaving out σ's and $\sigma\pi$'s for notational simplicity:

$$\begin{aligned}\partial P = &\sum_{j\le i} (-1)^i (-1)^j \varepsilon_{n-i}[v_0, \cdots, \widehat{v}_j, \cdots, v_i, w_n, \cdots, w_i] \\ &+ \sum_{j\ge i} (-1)^i (-1)^{i+1+n-j} \varepsilon_{n-i}[v_0, \cdots, v_i, w_n, \cdots, \widehat{w}_j, \cdots, w_i].\end{aligned}$$

The $j = i$ terms in these two sums give

$$\begin{aligned}&\varepsilon_n[w_n, \cdots, w_0] \;+\; \sum_{i>0} \varepsilon_{n-i}[v_0, \cdots, v_{i-1}, w_n, \cdots, w_i] \\ &\qquad + \sum_{i<n} (-1)^{n+i+1} \varepsilon_{n-i}[v_0, \cdots, v_i, w_n, \cdots, w_{i+1}] \;-\; [v_0, \cdots, v_n].\end{aligned}$$

In this expression the two summation terms cancel since replacing i by $i-1$ in the second sum produces a new sign $(-1)^{n+i}\varepsilon_{n-i+1} = -\varepsilon_{n-i}$. The remaining two terms $\varepsilon_n[w_n, \cdots, w_0]$ and $-[v_0, \cdots, v_n]$ represent $\rho(\sigma) - \sigma$. So in order to show that $\partial P + P\partial = \rho - \mathbb{1}$, it remains to check that in the formula for ∂P above, the terms with $j \ne i$ give $-P\partial$. Calculating $P\partial$ from the definitions, we have

$$\begin{aligned}P\partial = &\sum_{i<j} (-1)^i (-1)^j \varepsilon_{n-i-1}[v_0, \cdots, v_i, w_n, \cdots, \widehat{w}_j, \cdots, w_i] \\ &+ \sum_{i>j} (-1)^{i-1} (-1)^j \varepsilon_{n-i}[v_0, \cdots, \widehat{v}_j, \cdots, v_i, w_n, \cdots, w_i].\end{aligned}$$

Since $\varepsilon_{n-i} = (-1)^{n-i}\varepsilon_{n-i-1}$, this finishes the verification that $\partial P + P\partial = \rho - \mathbb{1}$, and so the theorem is proved when $A = \varnothing$. The proof also applies when $A \ne \varnothing$ since the maps ρ and P take chains in A to chains in A, so the dual homomorphisms ρ^* and P^* act on relative cochains. □

Induced homomorphisms are ring homomorphisms:

Proposition 3.8. *For a map $f: X \to Y$, the induced maps $f^*: H^n(Y;R) \to H^n(X;R)$ satisfy $f^*(\alpha \smile \beta) = f^*(\alpha) \smile f^*(\beta)$, and similarly in the relative case.*

Proof: This comes from the corresponding cochain formula $f^\sharp(\varphi) \smile f^\sharp(\psi) = f^\sharp(\varphi \smile \psi)$:

$$\begin{aligned}(f^\sharp\varphi \smile f^\sharp\psi)(\sigma) &= f^\sharp\varphi(\sigma\big|[v_0,\cdots,v_k])f^\sharp\psi(\sigma\big|[v_k,\cdots,v_{k+\ell}]) \\ &= \varphi(f\sigma\big|[v_0,\cdots,v_k])\psi(f\sigma\big|[v_k,\cdots,v_{k+\ell}]) \\ &= (\varphi\smile\psi)(f\sigma) = f^\sharp(\varphi\smile\psi)\end{aligned}$$

□

External Cup Product

One might guess that there should be some connection between cup product and product spaces, and indeed this is the case. We can define an **external cup product**

$$\mu: H^*(X;R) \otimes_R H^*(Y;R) \to H^*(X\times Y;R)$$

by $\mu(a\otimes b) = p_1^*(a) \smile p_2^*(b)$ where p_1 and p_2 are the projections of $X\times Y$ onto X and Y. (As we point out in the general discussion of tensor products near the end of §3.4a, if R is $\mathbb{Z}_n$ or a subring of $\mathbb{Q}$, then $\otimes_R$ is the same as $\otimes$, the ordinary tensor product of abelian groups.) The map μ is a homomorphism of additive structures and becomes a ring homomorphism if we define the multiplication in a tensor product of graded rings by $(a\otimes b)(c\otimes d) = (-1)^{|b||c|}ac\otimes bd$ where $|x|$ denotes the dimension of x, since

$$\begin{aligned}\mu\big((a\otimes b)(c\otimes d)\big) &= \mu\big((-1)^{|b||c|}ac\otimes bd\big) \\ &= (-1)^{|b||c|}p_1^*(a\smile c)\smile p_2^*(b\smile d) \\ &= (-1)^{|b||c|}p_1^*(a)\smile p_1^*(c)\smile p_2^*(b)\smile p_2^*(d) \\ &= p_1^*(a)\smile p_2^*(b)\smile p_1^*(c)\smile p_2^*(d) \\ &= \mu(a\otimes b)\mu(c\otimes d).\end{aligned}$$

In favorable cases the external cup product map μ is an isomorphism, for example for the torus $S^1\times S^1$, with $R=\mathbb{Z}$, according to the calculation in Example 1 earlier in this section. Here is a general result along these lines:

Theorem 3.9. *The external product $H^*(X;R)\otimes_R H^*(Y;R) \to H^*(X\times Y;R)$ is an isomorphism of rings if X and Y are CW complexes and $H^k(Y;R)$ is a finitely generated free R-module for all k.*

Results of this type, computing homology or cohomology of a product space, are known as **Künneth formulas**. The hypothesis that X and Y are CW complexes will be

shown to be unnecessary in §4.2 when we consider CW approximations to arbitrary spaces. On the other hand, the freeness hypothesis cannot always be dispensed with, as we shall see in §3.4b when we obtain a completely general Künneth formula for the homology of a product space.

Note that when the conclusion of the theorem holds, the ring structure in $H^*(X \times Y; R)$ is determined by the ring structures in $H^*(X; R)$ and $H^*(Y; R)$. Example 3 in §3.4e shows that some hypotheses are necessary in order for this to be true.

Before proving the theorem, let us look at a family of examples.

Example: Exterior Algebras. Generalizing the case of $S^1 \times S^1$, the theorem implies that, with $\mathbb{Z}$ coefficients, $H^*(S^k \times S^\ell)$ is free with additive basis $\{1, \alpha, \beta, \alpha \smile \beta\}$ where $\alpha \in H^k(S^k \times S^\ell)$ and $\beta \in H^\ell(S^k \times S^\ell)$ are pullbacks of generators of $H^k(S^k)$ and $H^\ell(S^\ell)$, hence satisfy $\alpha^2 = \beta^2 = 0$. The rest of the ring structure of $H^*(S^k \times S^\ell)$ is determined by the commutativity relation $\alpha\beta = (-1)^{k\ell}\beta\alpha$. If both k and ℓ are odd, $H^*(S^k \times S^\ell)$ is thus the exterior algebra $\Lambda_{\mathbb{Z}}[\alpha, \beta]$. In general, the **exterior algebra** $\Lambda_R[\alpha_1, \alpha_2, \cdots]$ over a commutative ring R with identity is the free R-module with basis the finite products $\alpha_{i_1} \cdots \alpha_{i_k}$ with $i_1 < \cdots < i_k$ and $k \geq 0$, with multiplication defined by the rules $\alpha_i\alpha_j = -\alpha_j\alpha_i$ and $\alpha_i^2 = 0$, the latter relation being a special case of the former if $2 \neq 0$ in R. Thus $\Lambda_R[\alpha_1, \alpha_2, \cdots]$ is the graded tensor product over R of the one-variable exterior algebras $\Lambda_R[\alpha_i]$ if the α_i's have odd dimension. In particular, $H^*(S^{k_1} \times \cdots \times S^{k_n}) \approx \Lambda_{\mathbb{Z}}[\alpha_1, \cdots, \alpha_n]$ if the dimensions k_i are all odd. With some k_i's even, one would have the tensor product of an exterior algebra for the odd-dimensional spheres and truncated polynomial rings $\mathbb{Z}[\alpha]/(\alpha^2)$ for the even-dimensional spheres. Of course, $\Lambda_{\mathbb{Z}}[\alpha]$ and $\mathbb{Z}[\alpha]/(\alpha^2)$ are isomorphic as rings, but when one takes tensor products in the graded sense it becomes important to distinguish them as graded rings, with α odd-dimensional in $\Lambda_{\mathbb{Z}}[\alpha]$ and even-dimensional in $\mathbb{Z}[\alpha]/(\alpha^2)$. These remarks apply more generally with any coefficient ring R in place of $\mathbb{Z}$, though when $R = \mathbb{Z}_2$ there is no need to distinguish between the odd and even dimensional cases since signs become irrelevant.

Proof of 3.9: We prove the theorem first for finite-dimensional X by induction on the dimension, using the diagram shown below. (The coefficient ring R will be fixed throughout the proof, so we omit it from the notation.) The upper triangle is obtained from the long exact sequence of the pair (X, A) by tensoring with $H^*(Y)$. By assumption $H^*(Y)$ is free, a direct sum of copies of the coefficient ring R, so tensoring with $H^*(Y)$ preserves exactness. The lower triangle is the long exact sequence of the pair $(X \times Y, A \times Y)$. The left-hand vertical map μ is defined for relative cohomology by the same formula as in the absolute case.

(∗)

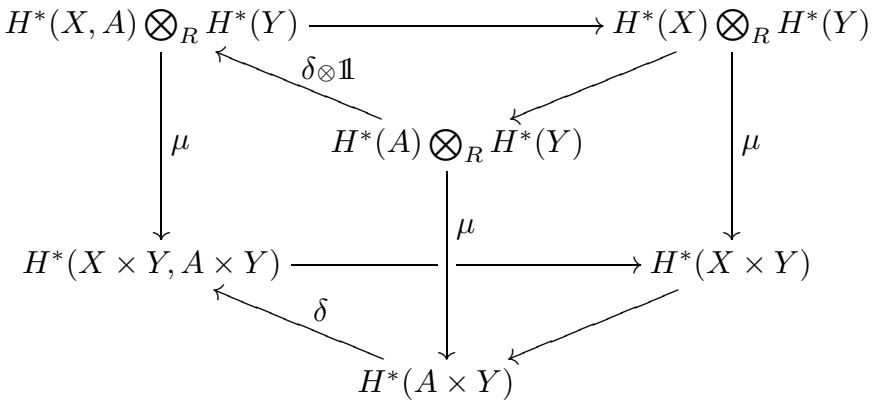

The two squares in the diagram not involving coboundary maps commute by the obvious naturality property of external cup product. Commmutativity of the remaining square can be checked as follows. Represent elements $\alpha \in H^k(A)$ and $\beta \in H^\ell(Y)$ by cocycles φ and ψ. Choose a cochain $\overline{\varphi} \in C^k(X)$ extending φ. Then $\delta\overline{\varphi}$ represents $\delta\alpha \in H^{k+1}(X,A)$, so $\mu(\delta\otimes\mathbb{1})(\alpha\otimes\beta)$ is represented by $p_1^\sharp(\delta\overline{\varphi}) \smile p_2^\sharp(\psi)$. On the other hand, $\mu(\alpha\otimes\beta)$ is represented by $q_1^\sharp\varphi \smile q_2^\sharp\psi$ where q_1 and q_2 are the projections of $A\times Y$ onto its two factors, so to apply δ to this we first extend it to the cochain $p_1^\sharp\overline{\varphi} \smile p_2^\sharp\psi$ on $X\times Y$, then apply δ. Finally, $p_1^\sharp(\delta\overline{\varphi}) \smile p_2^\sharp(\psi) = \delta p_1^\sharp(\overline{\varphi}) \smile p_2^\sharp(\psi) = \delta(p_1^\sharp\overline{\varphi} \smile p_2^\sharp\psi)$ since $\delta p_2^\sharp\psi = 0$.

The induction starts with the fairly trivial case that X is zero-dimensional. In this case $H^*(X\times Y)$ is a direct product of copies of $H^*(Y)$, one copy for each point of X. We assume $H^i(Y)$ is finitely generated and free, so let $\{\beta_j\}$ be a finite basis for this group. Then $H^i(X\times Y)$ consists of all sums $\sum_j r_j\beta_j$ for $r_j : X \to R$. Such a sum is exactly $\mu(\sum_j r_j\otimes\beta_j)$, so the theorem holds in this case.

For the inductive step, consider the diagram $(*)$ with $(X,A) = (X^n, X^{n-1})$. Induction on n and the five-lemma reduce us to showing the left-hand μ is an isomorphism. Let $f : (\coprod_\alpha D_\alpha^n, \coprod_\alpha \partial D_\alpha^n) \to (X^n, X^{n-1})$ consist of characteristic maps for the n-cells of X. Then we have a commutative square

$$\begin{array}{ccc} H^*(X^n,X^{n-1})\otimes_R H^*(Y) & \xrightarrow{f^*\otimes\mathbb{1}} & H^*(\coprod_\alpha D_\alpha^n, \coprod_\alpha \partial D_\alpha^n)\otimes_R H^*(Y) \\ \downarrow\mu & & \downarrow\mu \\ H^*(X^n\times Y, X^{n-1}\times Y) & \xrightarrow{(f\times\mathbb{1})^*} & H^*(\coprod_\alpha D_\alpha^n\times Y, \coprod_\alpha \partial D_\alpha^n\times Y) \end{array}$$

The two horizontal maps are isomorphisms since the quotient spaces are homeomorphic in both cases. So it suffices to show the right-hand μ is an isomorphism. But this follows by taking $(X,A) = (\coprod_\alpha D_\alpha^n, \coprod_\alpha \partial D_\alpha^n)$ in $(*)$ and applying the five-lemma and the induction

hypothesis, since the result for $X = \coprod_\alpha D^n_\alpha$ follows from the zero-dimensional case, as $\coprod_\alpha D^n_\alpha$ deformation retracts to a discrete set. This finishes the induction step.

The case that X is infinite-dimensional follows directly from the finite-dimensional case since the restriction map $H^k(X) \to H^k(X^n)$ is an isomorphism for large n, by cellular cohomology for example, and similarly for $H^k(X \times Y, X \times B) \to H^k(X^n \times Y, X^n \times B)$. □

For some applications it is useful to have a relative version of the preceding theorem. The relative external product is

$$\mu : H^*(X, A; R) \otimes_R H^*(Y, B; R) \to H^*(X \times Y, A \times Y \cup X \times B; R)$$

for CW pairs (X, A) and (Y, B), defined as in the absolute case by $\mu(a \otimes b) = p_1^*(a) \smile p_2^*(b)$ where $p_1^*(a) \in H^*(X \times Y, A \times Y; R)$ and $p_2^*(b) \in H^*(X \times Y, X \times B; R)$.

Theorem 3.10. *If (X, A) and (Y, B) are CW pairs and $H^k(Y, B; R)$ is a finitely generated free R-module for each k, then the external product $H^*(X, A; R) \otimes_R H^*(Y, B; R) \to H^*(X \times Y, A \times Y \cup X \times B; R)$ is an isomorphism of rings.*

Proof: First note that the case $B = \varnothing$ follows from the absolute case by applying the five-lemma to the large diagram in the proof of the absolute case. So it remains to deduce the case $B \neq \varnothing$ from the case $B = \varnothing$.

The following commutative diagram shows that collapsing B to a point reduces us to the case that B is a point.

$$\begin{array}{ccc} H^*(X, A) \bigotimes_R H^*(Y, B) & \xleftarrow{\approx} & H^*(X, A) \bigotimes_R H^*(Y/B, B/B) \\ \downarrow \mu & & \downarrow \mu \\ H^*(X \times Y, A \times Y \cup X \times B) & \xleftarrow{\approx} & H^*(X \times (Y/B), A \times (Y/B) \cup X \times (B/B)) \end{array}$$

Here the lower map is an isomorphism since the quotient spaces $(X \times Y)/(A \times Y \cup X \times B)$ and $\big(X \times (Y/B)\big)/\big(A \times (Y/B) \cup X \times (B/B)\big)$ are the same.

Now when B is a point $y_0 \in Y$, consider the commutative diagram

$$\begin{array}{ccccc} H^*(X, A) \bigotimes_R H^*(Y, y_0) & \longrightarrow & H^*(X, A) \bigotimes_R H^*(Y) & \longrightarrow & H^*(X, A) \bigotimes_R H^*(y_0) \\ \downarrow \mu & & \downarrow \mu & & \downarrow \mu \\ & & & \nearrow & H^*(X \times y_0, A \times y_0) \\ & & & & \uparrow \approx \\ H^*(X \times Y, X \times y_0 \cup A \times Y) & \longrightarrow & H^*(X \times Y, A \times Y) & \longrightarrow & H^*(X \times y_0 \cup A \times Y, A \times Y) \end{array}$$

Since y_0 is a retract of Y, the upper row of this diagram is a split short exact sequence. The lower row is the long exact sequence of a triple, and it too is a split short exact sequence since $(X\times y_0, A\times y_0)$ is a retract of $(X\times Y, A\times Y)$. The middle and right vertical maps μ are isomorphisms by the case $B=\varnothing$ since $H^k(Y;R)$ is a finitely generated free R-module if $H^k(Y,y_0;R)$ is. The five-lemma then implies that the μ on the left is an isomorphism as well. □

The relative external product for pairs (X,x_0) and (Y,y_0) gives a **reduced external cup product**

$$\widetilde H^*(X;R)\otimes_R\widetilde H^*(Y;R)\to\widetilde H^*(X\wedge Y;R)$$

where $X\wedge Y$ is the smash product $X\times Y/(X\times\{y_0\}\cup\{x_0\}\times Y)$. The preceding theorem implies that this reduced external product is an isomorphism if $\widetilde H^*(X;R)$ or $\widetilde H^*(Y;R)$ is free and finitely generated in each dimension. For example, $\widetilde H^n(X;R)\approx\widetilde H^{n+k}(X\wedge S^k;R)$ via external cup product with a generator of $H^k(S^1;R)\approx R$. The space $X\wedge S^k$ is the k-fold reduced suspension Σ^kX of X, so we see that the **suspension isomorphisms** $\widetilde H^n(X;R)\approx\widetilde H^{n+k}(\Sigma^kX;R)$ which exist by elementary exact sequence arguments can also be obtained via external cup product with a generator of $\widetilde H^*(S^k;R)$.

One reason why cup products are harder to compute than the additive structure of cohomology is that there is no direct definition of cup products in cellular cohomology. However, there is a nice definition of external cup product in cellular cohomology. To distinguish this from the previously defined external cup product, let us call it the **cross product**

$$H^k(X;R)\times H^\ell(Y;R)\xrightarrow{\ \times\ }H^{k+\ell}(X\times Y;R)$$

for CW complexes X and Y. This comes from a cross product of cellular cochains defined by $(\varphi\times\psi)(e^k_\alpha\times e^\ell_\beta)=\varphi(e^k_\alpha)\psi(e^\ell_\beta)$ where a cellular n-cochain is regarded as a function on cells taking the value 0 on cells of dimension other than n. It is an exercise to verify the cellular coboundary formula $\delta(\varphi\times\psi)=\delta\varphi\times\psi+(-1)^k\varphi\times\delta\psi$, from the cellular boundary formula $\partial(e^k_\alpha\times e^\ell_\beta)=\partial e^k_\alpha\times e^\ell_\beta+(-1)^ke^k_\alpha\times\partial e^\ell_\beta$ in Proposition 3.35. Hence there is an induced cross product in cellular cohomology, just as Lemma 3.6 gives a cup product in singular cohomology.

To show this cross product coincides with the external cup product, let e^k_α and e^ℓ_β be cells in X and Y, with characteristic maps Φ_α and Φ_β. Then we have a commutative

diagram of external cup products, with coefficients in R understood:

$$\begin{array}{ccccc}
H^k(D^k,\partial D^k) & \times & H^\ell(D^\ell,\partial D^\ell) & \longrightarrow & H^{k+\ell}(D^k\times D^\ell,\partial D^k\times D^\ell\cup D^k\times\partial D^\ell)\\
\uparrow{\scriptstyle \Phi_\alpha^*} & & \uparrow{\scriptstyle \Phi_\beta^*} & & \uparrow{\scriptstyle (\Phi_\alpha\times\Phi_\beta)^*}\\
H^k(X^k,X^{k-1}) & \times & H^\ell(Y^\ell,Y^{\ell-1}) & \longrightarrow & H^{k+\ell}(X^k\times Y^\ell,X^{k-1}\times Y^\ell\cup X^k\times Y^{\ell-1})\\
\downarrow & & \downarrow & & \downarrow\\
H^k(X^k) & \times & H^\ell(Y^\ell) & \longrightarrow & H^{k+\ell}(X^k\times Y^\ell)\\
\uparrow & & \uparrow & & \uparrow\\
H^k(X) & \times & H^\ell(Y) & \longrightarrow & H^{k+\ell}(X\times Y)
\end{array}$$

In the upper row the map is just the ring multiplication $R\times R\to R$. Therefore the external cup product in the second row is the same as the cross product of cellular cochains. Recalling how the isomorphism between cellular and singular homology is constructed, as in the diagram preceding Theorem 3.5, the rest of the diagram above then shows that cross product agrees with external cup product. Some authors in fact use the term 'cross product' for what we have called the external cup product.

One could go a step further and define cup product in terms of cross product or external cup product, using the diagonal map $\Delta: X\to X\times X$, $x\mapsto(x,x)$. Namely, the composition

$$H^k(X;R)\times H^\ell(X;R)\xrightarrow{\times} H^{k+\ell}(X\times X;R)\xrightarrow{\Delta^*} H^{k+\ell}(X;R)$$

is exactly the cup product since $\Delta^*\mu(a\otimes b)=\Delta^*(p_1^*(a)\smile p_2^*(b))=\Delta^*(p_1^*(a))\smile\Delta^*(p_2^*(b))=a\smile b$ since both compositions $p_1\Delta$ and $p_2\Delta$ are the identity map of X.

Unfortunately Δ is not a cellular map, taking k-skeleton to k-skeleton for each k, so Δ does not induce a map of cellular cochains. It is possible to homotope Δ to a cellular map, by appealing to 'cellular approximation,' Theorem 4.12, but this involves arbitrary choices. For example, the diagonal of a square can be pushed across either adjacent triangle. In particular cases one might hope to understand the geometry well enough to compute an explicit cellular approximation to the diagonal map, but usually other techniques for computing cup products are preferable.

Spaces with Polynomial Cohomology

Among the simplest graded rings are polynomial rings $R[\alpha]$ and their truncated versions $R[\alpha]/(\alpha^n)$, consisting of polynomials of degree less than n. The example we have seen is $H^*(\mathbb{R}\mathrm{P}^2;\mathbb{Z}_2)\approx\mathbb{Z}_2[\alpha]/(\alpha^3)$. Generalizing this, we have:

Theorem 3.11. *$H^*(\mathbb{R}\mathrm{P}^n;\mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha]/(\alpha^{n+1})$ and $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha]$, where α is a generator of $H^1(\mathbb{R}\mathrm{P}^n;\mathbb{Z}_2)$ or $H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$, respectively. For complex projective spaces, $H^*(\mathbb{C}\mathrm{P}^n;\mathbb{Z}) \approx \mathbb{Z}[\alpha]/(\alpha^{n+1})$ and $H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}) \approx \mathbb{Z}[\alpha]$ where α is a generator of $H^2(\mathbb{C}\mathrm{P}^n;\mathbb{Z})$ or $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$.*

This turns out to be an extremely important result, and it can be proved in a number of different ways. The proof we give here consists of a fairly direct reduction to a special case of Theorem 3.10 on the relative external cup product. Another proof, by induction and Poincaré duality, will be given in the next section. A third proof is contained in §4.4d as an application of the Gysin sequence.

Proof: The argument will be the same in both the real and complex cases, so let us consider the complex case because $\mathbb{Z}$ coefficients are slightly more subtle than $\mathbb{Z}_2$ coefficients. Let us abbreviate $\mathbb{C}\mathrm{P}^n$ to P^n. We assume the additive structure of $H^*(P^n;\mathbb{Z})$ is already known, for example from the standard CW structure on P^n. If we can show that for α a generator of $H^2(P^n;\mathbb{Z})$, the power α^n is a generator of $H^{2n}(P^n;\mathbb{Z})$, then the isomorphism $H^*(P^n;\mathbb{Z}) \approx \mathbb{Z}[\alpha]/(\alpha^{n+1})$ will follow. For if α^n generates $H^{2n}(P^n;\mathbb{Z})$, then α^i must generate $H^{2i}(P^n;\mathbb{Z})$ for each i between 1 and n, otherwise if $\alpha^i = k\beta$ for some $k > 1$ then $\alpha^n = k\beta\alpha^{n-i}$, which cannot generate $H^{2n}(P^n;\mathbb{Z})$.

Viewing P^n as the quotient of $\mathbb{C}^{n+1} - \{0\}$ under the identifications given by scalar multiplication, then the points of P^n represented by points $(z_0, \cdots, z_n) \in \mathbb{C}^{n+1}$ with the coordinate $z_i \neq 0$ form an open set U_i in P^n. In fact U_i is homeomorphic to $\mathbb{C}^n$ since points in U_i are uniquely representable as points $(z_0, \cdots, z_n) \in \mathbb{C}^{n+1}$ with $z_i = 1$. The complement $P^n - U_i$ is a copy of P^{n-1}, and we can regard P^n as being obtained from this P^{n-1} by attaching the $2n$-cell U_i.

Now let us work our way through the following commutative diagram, with coefficient in $\mathbb{Z}$ implicit.

$$\begin{array}{ccccccc}
H^2(P^n) & \times \cdots \times & H^2(P^n) & \longrightarrow & H^{2n}(P^n) \\
\uparrow\approx & & \uparrow\approx & & \uparrow\approx \\
H^2(P^n, U_1) & \times \cdots \times & H^2(P^n, U_n) & \longrightarrow & H^{2n}(P^n, P^n - P^0) \\
\downarrow & & \downarrow & & \downarrow\approx \\
H^2(U_0, V_1) & \times \cdots \times & H^2(U_0, V_n) & \longrightarrow & H^{2n}(U_0, U_0 - P^0) \\
\downarrow\approx & & \downarrow\approx & \nearrow & \\
H^2(\mathbb{C}, \mathbb{C} - \{0\}) & \times \cdots \times & H^2(\mathbb{C}, \mathbb{C} - \{0\}) & &
\end{array}$$

The top row is the n-fold cup product map. The second row is a relative version of this, using the fact that $U_1 \cup \cdots \cup U_n$ is the complement of the point $P^0 = (1, 0, \cdots, 0)$. The

vertical isomorphisms $H^2(P^n) \approx H^2(P^n, U_i)$ come from the fact that U_i is contractible, and the last vertical isomorphism $H^{2n}(P^n) \approx H^{2n}(P^n, P^n - P^0)$ holds since $P^n - P^0$ deformation retracts onto $P^n - U_0 = P^{n-1}$ and $H^{2n}(P^n) \approx H^{2n}(P^n, P^{n-1})$ by cellular cohomology. The third row of the diagram is obtained from the second by restricting to U_0 and setting $V_i = U_0 \cap U_i$, the points of the form $(1, z_1, \cdots, z_n)$ with $z_i \neq 0$. The right-hand vertical map between second and third rows is an isomorphism by excision. The i^{th} term in the bottom row is obtained from the term above it by restricting to the coordinate z_i. Note that the pair (U_0, V_i) is homeomorphic to the product of $(\mathbb{C}, \mathbb{C} - \{0\})$ with $\mathbb{C}^{n-1}$, which accounts for the isomorphisms between the last two rows. Under these isomorphisms the upward diagonal map in the diagram, which is the relative external cup product, corresponds to the horizontal map in the third row, which is the ordinary relative cup product.

The generator α of the upper-left term $H^2(P^n)$ in the diagram maps to a generator $\overline{\alpha}$ of the lower-left term $H^2(\mathbb{C}, \mathbb{C} - \{0\})$. This can be seen by restricting to the coordinates z_0 and z_1, which give a 2-sphere $P^1 \subset P^n$ with the subspace '$\mathbb{C}$' identified with the complement of the point $(z_0, z_1) = (1, 0)$. Similarly, the generator α of the i^{th} factor $H^2(P^n)$ in the top row maps to the generator $\overline{\alpha}$ of the i^{th} factor of the bottom row. By Theorem 3.10 the n-fold external product of $\overline{\alpha}$ with itself generates $H^{2n}(U_0, U_0 - P^0)$, since we can repace $(\mathbb{C}, \mathbb{C} - \{0\})$ by $(D^2, \partial D^2)$ and $(U_0, U_0 - P^0)$ by $(D^{2n}, \partial D^{2n})$. The commutativity of the diagram then implies that α^n generates $H^{2n}(P^n)$, finishing the proof for $\mathbb{C}\mathrm{P}^n$. The result for $\mathbb{C}\mathrm{P}^\infty$ follows from the result for CP^n since the restriction map $H^i(\mathbb{C}\mathrm{P}^\infty) \to H^i(\mathbb{C}\mathrm{P}^n)$ is an isomorphism for $i \leq 2n$.

These arguments work just as well for $\mathbb{R}\mathrm{P}^n$ and $\mathbb{R}\mathrm{P}^\infty$ if we use $\mathbb{Z}_2$ coefficients and change all superscripts '2' to '1.' □

There are also quaternionic projective spaces $\mathbb{H}\mathrm{P}^n$ and $\mathbb{H}\mathrm{P}^\infty$, defined exactly as was done in the complex case in §0.2, with CW structures of the form $e^0 \cup e^4 \cup e^8 \cup \cdots$. Associativity of quaternion multiplication is needed for the identification $v \sim \lambda v$ to be an equivalence relation, so the definition does not extend to octonionic projective spaces, though there is an octonionic projective plane $\mathbb{O}\mathrm{P}^2$ which will be defined in §4.3. The cup product structure in quaternionic projectives spaces is just like that in complex projective spaces:

$$H^*(\mathbb{H}\mathrm{P}^\infty; \mathbb{Z}) \approx \mathbb{Z}[\alpha] \quad \text{and} \quad H^*(\mathbb{H}\mathrm{P}^n; \mathbb{Z}) \approx \mathbb{Z}[\alpha]/(\alpha^{n+1}), \quad \text{with } \alpha \in H^4$$

The same proof as in the complex case works here too.

The following theorem of Hopf is a nice algebraic application of the cup product structure in $H^*(\mathbb{R}\mathrm{P}^n;\mathbb{Z}_2)$.

Theorem 3.12. *If $\mathbb{R}^n$ has the structure of a division algebra over the scalar field $\mathbb{R}$, then n must be a power of 2.*

Proof: If $\mathbb{R}^n$ is a division algebra, define $g: S^{n-1}\times S^{n-1}\to S^{n-1}$ by $g(x,y)=xy/|xy|$. This is well-defined since there are no zero-divisors, and continuous by the bilinearity of the multiplication. From the relations $(-x)y=-(xy)=x(-y)$ it follows that $g(-x,y)=-g(x,y)=g(x,-y)$. This implies that g induces a quotient map $h:\mathbb{R}\mathrm{P}^{n-1}\times\mathbb{R}\mathrm{P}^{n-1}\to\mathbb{R}\mathrm{P}^{n-1}$.

We claim that $h^*:H^1(\mathbb{R}\mathrm{P}^{n-1};\mathbb{Z}_2)\to H^1(\mathbb{R}\mathrm{P}^{n-1}\times\mathbb{R}\mathrm{P}^{n-1};\mathbb{Z}_2)$ is the map $h^*(\gamma)=\alpha+\beta$ where γ generates $H^1(\mathbb{R}\mathrm{P}^{n-1};\mathbb{Z}_2)$ and α and β are the pullbacks of γ under the projections of $\mathbb{R}\mathrm{P}^{n-1}\times\mathbb{R}\mathrm{P}^{n-1}$ onto its two factors. This claim can be proved as follows. We may assume $n>2$, so $\pi_1(\mathbb{R}\mathrm{P}^{n-1})\approx\mathbb{Z}_2$. Let $\lambda:I\to S^{n-1}$ be a path joining a point x to the antipodal point $-x$. Then for fixed y, the path $s\mapsto g(\lambda(s),y)$ joins $g(x,y)$ to $g(-x,y)=-g(x,y)$. Hence, identifying antipodal points, h takes a nontrivial loop in the first $\mathbb{R}\mathrm{P}^{n-1}$ factor of $\mathbb{R}\mathrm{P}^{n-1}\times\mathbb{R}\mathrm{P}^{n-1}$ to a nontrivial loop in $\mathbb{R}\mathrm{P}^{n-1}$. The same argument works for the second factor. Hence the restriction of h to the 1-skeleton $S^1\vee S^1$ is homotopic to the map which includes each S^1 summand of $S^1\vee S^1$ into $\mathbb{R}\mathrm{P}^{n-1}$ as the 1-skeleton. Since restriction to the 1-skeleton is an isomorphism on $H^1(-;\mathbb{Z}_2)$ for both $\mathbb{R}\mathrm{P}^{n-1}$ and $\mathbb{R}\mathrm{P}^{n-1}\times\mathbb{R}\mathrm{P}^{n-1}$, it follows that $h^*(\gamma)=\alpha+\beta$.

Since $\gamma^n=0$ we have $0=h^*(\gamma^n)=(\alpha+\beta)^n=\sum_k\binom{n}{k}\alpha^k\beta^{n-k}$. This is an equation in the ring $H^*(\mathbb{R}\mathrm{P}^{n-1}\times\mathbb{R}\mathrm{P}^{n-1};\mathbb{Z}_2)\approx\mathbb{Z}_2[\alpha,\beta]/(\alpha^n,\beta^n)$, so the coefficient $\binom{n}{k}$ must be zero in $\mathbb{Z}_2$ for all k in the range $0<k<n$. It is a rather easy number theory fact that this happens only when n is a power of 2. Namely, an equivalent statement is that $(1+x)^n=1+x^n$ in the polynomial ring $\mathbb{Z}_2[x]$ only when n is a power of 2. To prove this, write n as a sum of powers of 2, $n=n_1+\cdots+n_k$ with $n_1<\cdots<n_k$. Then $(1+x)^n=(1+x)^{n_1}\cdots(1+x)^{n_k}=(1+x^{n_1})\cdots(1+x^{n_k})$ since squaring is an additive homomorphism with $\mathbb{Z}_2$ coefficients. If one multiplies the product $(1+x^{n_1})\cdots(1+x^{n_k})$ out, no terms combine or cancel since $n_i\geq 2n_{i-1}$ for each i, and so the resulting polynomial has 2^k terms. Thus if this polynomial equals $1+x^n$ we must have $k=1$, which means that n is a power of 2. □

The cup product structure for $\mathbb{R}\mathrm{P}^\infty$ with $\mathbb{Z}$ coefficients can easily be deduced from the cup product structure with $\mathbb{Z}_2$ coefficients, as follows. In general, a ring homomorphism $R\to S$ induces a ring homomorphism $H^*(X,A;R)\to H^*(X,A;S)$. In the case of the

projection $\mathbb{Z} \to \mathbb{Z}_2$ we get for $\mathbb{R}\mathrm{P}^\infty$ an induced chain map of cellular cochain complexes with $\mathbb{Z}$ and $\mathbb{Z}_2$ coefficients:

$$\begin{array}{ccccccccccccc}
\cdots & \longleftarrow & \mathbb{Z} & \xleftarrow{2} & \mathbb{Z} & \xleftarrow{0} & \mathbb{Z} & \xleftarrow{2} & \mathbb{Z} & \xleftarrow{0} & \mathbb{Z} & \longleftarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
\cdots & \longleftarrow & \mathbb{Z}_2 & \xleftarrow{0} & \mathbb{Z}_2 & \xleftarrow{0} & \mathbb{Z}_2 & \xleftarrow{0} & \mathbb{Z}_2 & \xleftarrow{0} & \mathbb{Z}_2 & \longleftarrow & 0
\end{array}$$

The vertical maps are surjective, so the ring homomorphism $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}) \to H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$ is injective in positive dimensions, with image the even-dimensional part of $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$. (This can also be deduced from the universal coefficient theorem.) Hence $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}) \approx \mathbb{Z}[\alpha]/(2\alpha)$ for α a generator of $H^2(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}) \approx \mathbb{Z}_2$.

The cup product structure in $H^*(\mathbb{R}\mathrm{P}^n;\mathbb{Z})$ can be computed in a similar fashion, though the description is a little cumbersome:

$$H^*(\mathbb{R}\mathrm{P}^{2k};\mathbb{Z}) \approx \mathbb{Z}[\alpha]/(2\alpha, \alpha^{k+1})$$
$$H^*(\mathbb{R}\mathrm{P}^{2k+1};\mathbb{Z}) \approx \mathbb{Z}[\alpha,\beta]/(2\alpha, \alpha^{k+1}, \beta^2, \alpha\beta)$$

where β is a generator of $H^{2k+1}(\mathbb{R}\mathrm{P}^{2k+1};\mathbb{Z}) \approx \mathbb{Z}$. Thus the rings $H^*(\mathbb{R}\mathrm{P}^{2k+1};\mathbb{Z})$ and $H^*(\mathbb{R}\mathrm{P}^{2k} \vee S^{2k+1};\mathbb{Z})$ are isomorphic, though with $\mathbb{Z}_2$ coefficients this is no longer true, since the generator $\alpha \in H^1(\mathbb{R}\mathrm{P}^{2k+1};\mathbb{Z}_2)$ has $\alpha^{2k+1} \neq 0$, while $\alpha^{2k+1} = 0$ for the generator $\alpha \in H^1(\mathbb{R}\mathrm{P}^{2k} \vee S^{2k+1};\mathbb{Z}_2)$.

We next describe a construction for enlarging S^{2n} to a space $J(S^{2n})$ whose cohomology ring $H^*(J(S^{2n});\mathbb{Z})$ is almost the polynomial ring $\mathbb{Z}[x]$ on a generator x of dimension $2n$. And with $\mathbb{Q}$ coefficients, $H^*(J(S^{2n});\mathbb{Q})$ is exactly the polynomial ring $\mathbb{Q}[x]$. This construction, known as the **James reduced product**, is also of interest because of its connections with loopspaces, described in §4.4i.

For a space X, let X^k be the product of k copies of X. From the disjoint union $\coprod_{k\geq 1} X^k$, let us form a quotient space $J(X)$ by identifying $(x_1,\cdots,x_i,\cdots,x_k)$ with $(x_1,\cdots,\widehat{x}_i,\cdots,x_k)$ if $x_i = e$, a chosen basepoint of X. Points of $J(X)$ can thus be thought of as k-tuples $(x_1,\cdots,x_k)$, $k \geq 0$, with no $x_i = e$. Inside $J(X)$ is the subspace $J_m(X)$ consisting of the points $(x_1,\cdots,x_k)$ with $k \leq m$. This can be viewed as a quotient space of X^m under the identifications $(x_1,\cdots,x_i,e,\cdots,x_m) \sim (x_1,\cdots,e,x_i,\cdots,x_m)$. For example, $J_1(X) = X$ and $J_2(X) = X \times X/(x,e) \sim (e,x)$. If X is a CW complex with e a 0-cell, the quotient map $X^m \to J_m(X)$ glues together the m subcomplexes of X^m where one coordinate is e. These glueings are by homeomorphisms taking cells onto cells, so $J_m(X)$ inherits a CW structure from X^m. There are natural inclusions $J_m(X) \subset J_{m+1}(X)$ as subcomplexes, and $J(X)$ is the union of these subcomplexes, hence is also a CW complex.

Proposition 3.13. *For $n > 0$, $H^*\big(J(S^n);\mathbb{Z}\big)$ consists of a $\mathbb{Z}$ in each dimension a multiple of n. If n is even, the i^{th} power of a generator of $H^n\big(J(S^n);\mathbb{Z}\big)$ is $i!$ times a generator of $H^{in}\big(J(S^n);\mathbb{Z}\big)$, for each $i \geq 1$. If n is odd $H^*\big(J(S^n);\mathbb{Z}\big)$ is isomorphic as a graded ring to $H^*(S^n;\mathbb{Z}) \otimes H^*\big(J(S^{2n});\mathbb{Z}\big)$.*

Thus for n even, $H^*\big(J(S^n);\mathbb{Z}\big)$ can be identified with the subring of the polynomial ring $\mathbb{Q}[x]$ additively generated by the monomials $x^i/i!$. This subring is called a **divided polynomial algebra**.

Proof: Giving S^n its usual CW structure, the resulting CW structure on $J(S^n)$ consists of exactly one cell in each dimension a multiple of n. Thus if $n > 1$ we deduce immediately from cellular cohomology that $H^*\big(J(S^n);\mathbb{Z}\big)$ consists of $\mathbb{Z}$'s in dimensions a multiple of n. To see that this conclusion is also true when $n = 1$, consider the quotient map $q : (S^n)^m \to J_m(S^n)$. This carries each cell of $(S^n)^m$ homeomorphically onto a cell of $J_m(S^n)$. so the cellular boundary maps in $J_m(S^n)$ will be trivial if the same is true in $(S^n)^m$; and the cellular boundary maps in $(S^n)^m$ must be trivial in order for $H^n\big((S^n)^m;\mathbb{Z}\big)$ to be what it is.

Since q maps each of the m n-cells of $(S^n)^m$ homeomorphically onto the n-cell of $J_m(S^n)$, we see from cellular cohomology that a generator $\alpha \in H^n\big(J_m(S^n);\mathbb{Z}\big)$ pulls back by q^* to the sum $\alpha_1 + \cdots + \alpha_m$ of the generators of $H^n\big((S^n)^m;\mathbb{Z}\big)$ corresponding to the n-cells of $(S^n)^m$. Assuming that n is even, the cup product structure in $H^*\big(J(S^n);\mathbb{Z}\big)$ is strictly commutative and $H^*\big((S^n)^m;\mathbb{Z}\big) \approx \mathbb{Z}[\alpha_1, \cdots, \alpha_m]/(\alpha_1^2, \cdots, \alpha_m^2)$. The power α^m then pulls back to $(\alpha_1 + \cdots + \alpha_m)^m = m!\alpha_1 \cdots \alpha_m$, where the product $\alpha_1 \cdots \alpha_m$ generates $H^{mn}\big((S^n)^m;\mathbb{Z}\big) \approx \mathbb{Z}$. The map q is a homeomorphism from the mn-cell of $(S^n)^m$ to the mn-cell of $J_m(S^n)$, so q^* is an isomorphism on H^{mn}. This implies that α^m is $m!$ times a generator of $H^{mn}\big(J_m(S^n);\mathbb{Z}\big)$. Since the cells of $J(S^n) - J_m(S^n)$ have dimension at least $(m+1)n$, the inclusion $J_m(S^n) \subset J(S^n)$ induces isomorphisms on H^i for $i \leq mn$. Thus if we let x_i denote a suitable generator of $H^{in}\big(J(S^n);\mathbb{Z}\big)$, we have $x_1^m = m!x_m$ for all m.

The case n odd is similar and will be left as an exercise for the reader. □

With coefficients in a field, the cup product structure for $J(S^n)$ becomes somewhat simpler. For example, $H^*\big(J(S^{2k});\mathbb{Q}\big) \approx \mathbb{Q}[x]$ since $x_1^i = i!x_i$ is nonzero for all m. Similarly, for p prime we have $H^*\big(J_{p-1}(S^{2k});\mathbb{Z}_p\big) \approx \mathbb{Z}_p[x]/(x^p)$, a truncated polynomial ring. More generally,

$$H^*\big(J(S^{2k});\mathbb{Z}_p\big) \approx \mathbb{Z}_p[x_1, x_p, x_{p^2}, \cdots]/(x_1^p, x_p^p, x_{p^2}^p, \cdots) = \bigotimes_{i\geq 0} \mathbb{Z}_p[x_{p^i}]/(x_{p^i}^p)$$

according to an exercise at the end of the section.

The examples of projective spaces lead naturally to the following question: Given a coefficient ring R and an integer $d > 0$, is there a space X having $H^*(X;R) \approx R[\alpha]$ with $\alpha \in H^d(X;R)$? Historically, it took major advances in the theory to answer this simple-looking question. Here is a table giving all the possible values of d for some of the most obvious and important choices of R, namely $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{Z}_2$, and $\mathbb{Z}_p$ with p an odd prime:

R	d
$\mathbb{Z}$	2, 4
$\mathbb{Q}$	any even number
$\mathbb{Z}_2$	1, 2, 4
$\mathbb{Z}_p$	any even divisor of $2(p-1)$

As we have seen, projective spaces give the examples for $\mathbb{Z}$ and $\mathbb{Z}_2$. Examples for $\mathbb{Q}$ are the spaces $J(S^d)$, and examples for $\mathbb{Z}_p$ are constructed in §3.4h. Showing that no other d's are possible takes considerably more work. The fact that d must be even when $R \neq \mathbb{Z}_2$ is a consequence of the commutativity property of cup product. In Theorem 4.89 and Corollary 4.90 we will settle the case $R = \mathbb{Z}$ and show that d must be a power of 2 for $R = \mathbb{Z}_2$ and an even divisor of $2(p-1)$ times a power of p for $R = \mathbb{Z}_p$, p odd. Ruling out the remaining cases is best done using K-theory, as in [VBKT]; the classical reference is [Adams-Atiyah]. However there is one slightly anomalous case, $R = \mathbb{Z}_2$, $k = 8$, which must be treated by special arguments; see [Toda63].

In fact, for each even d there exists a CW complex X_d consisting of a single cell in each dimension a multiple of d, such that X_d is simultaneously an example for all the admissible choices of coefficients R in the table. For example, we may take $X_2 = \mathbb{C}\mathrm{P}^\infty$ and $X_4 = \mathbb{H}\mathrm{P}^\infty$. The next space X_6 would have $H^*(X_6;\mathbb{Z}_p) \approx \mathbb{Z}_p[\alpha]$ for $p = 7, 13, 19, 31, \cdots$, primes of the form $3s+1$, the condition $6|2(p-1)$ being equivalent to $p = 3s+1$. (By a famous theorem of Dirichlet there are infinitely many primes in any such arithmetic progression.) Note that, in terms of $\mathbb{Z}$ coefficients, X_d must have the property that for a generator α of $H^d(X_k;\mathbb{Z})$, each power α^i is an integer a_i times a generator of $H^{di}(X_d;\mathbb{Z})$, with $a_i \neq 0$ if $H^*(X_d;\mathbb{Q}) \approx \mathbb{Q}[\alpha]$ and a_i relatively prime to p if $H^*(X_d;\mathbb{Z}_p) \approx \mathbb{Z}_p[\alpha]$. See [Hoffman-Porter] or Volume II for a construction of X_d.

One might also ask about realizing the truncated polynomial ring $R[\alpha]/(\alpha^{n+1})$, for which $\mathbb{R}\mathrm{P}^n$, $\mathbb{C}\mathrm{P}^n$, and $\mathbb{H}\mathrm{P}^n$ give the examples we have seen so far, not counting the trivial case $n = 1$ for which spheres provide examples. The analysis for polynomial rings also settles which truncated polynomial rings are realizable; there are just a few more than for

the full polynomial rings.

There is also the question of realizing polynomial rings $R[\alpha_1, \cdots, \alpha_k]$ with generators α_i in specified dimensions d_i. The tensor product $R[\alpha_1, \cdots, \alpha_m] \otimes_R R[\beta_1, \cdots, \beta_n]$ is equal to $R[\alpha_1, \cdots, \alpha_m, \beta_1, \cdots, \beta_n]$, so the product of two spaces with polynomial cohomology is again a space with polynomial cohomology, assuming the number of polynomial generators is finite in each dimension. For example, the k-fold product $(\mathbb{C}\mathrm{P}^\infty)^k$ has $H^*((\mathbb{C}\mathrm{P}^\infty)^k; \mathbb{Z}) \approx \mathbb{Z}[\alpha_1, \cdots, \alpha_k]$ with each α_i 2-dimensional. Similarly, products of the spaces $J(S^{d_i})$ realize all choices of even d_i's with $\mathbb{Q}$ coefficients.

However, with $\mathbb{Z}$ and $\mathbb{Z}_p$ coefficients, products of one-variable examples do not exhaust all the possibilities. As we show in §4.4d, there are three other basic examples with $\mathbb{Z}$ coefficients:

1. Generalizing the space $\mathbb{C}\mathrm{P}^\infty$ of complex lines through the origin in $\mathbb{C}^\infty$, there is the **Grassmannian** $G_k(\mathbb{C}^\infty)$ of k-dimensional vector subspaces of $\mathbb{C}^\infty$, and this has $H^*(G_k(\mathbb{C}^\infty); \mathbb{Z}) \approx \mathbb{Z}[\alpha_1, \cdots, \alpha_k]$ with α_i of dimension $2i$. This space is also known as $BU(n)$, 'the classifying space of the unitary group $U(n)$.' It is central to the study of vector bundles and K-theory.

2. Replacing $\mathbb{C}$ by $\mathbb{H}$, there is the quaternionic Grassmannian $G_k(\mathbb{H}^\infty)$, also known as $BSp(n)$, with $H^*(G_k(\mathbb{H}^\infty); \mathbb{Z}) \approx \mathbb{Z}[\alpha_1, \cdots, \alpha_k]$ with α_i of dimension $4i$.

3. The classifying space $BSU(n)$ of the special unitary group has cohomology the same as $BU(n)$ but with the generator α_1 omitted, so $H^*(BSU(n); \mathbb{Z}) \approx \mathbb{Z}[\alpha_2, \cdots, \alpha_k]$ with α_i $2i$-dimensional.

These examples and their products account for all the realizable polynomial cup product rings in several variables with $\mathbb{Z}$ coefficients, as was shown in [Adams-Wilkerson]. The situation for $\mathbb{Z}_p$ coefficients is quite a bit more complicated and will be discussed in §3.4h.

These questions are all special cases of the general

Realization Problem. Which graded commutative R-algebras occur as cup product algebras $H^*(X; R)$ of spaces X?

This is a difficult problem, with the degree of difficulty depending strongly on the coefficient ring R. The most accessible case is $R = \mathbb{Q}$, where essentially every graded commutative $\mathbb{Q}$-algebra is realizable [Quillen]. Next in order of difficulty is $R = \mathbb{Z}_p$ with p prime. This is much harder than the case of $\mathbb{Q}$, and only partial results, obtained with much labor, are known, mainly about realizing polynomial rings. Finally there is $R = \mathbb{Z}$, about which very little is known beyond what is implied by the $\mathbb{Z}_p$ cases.

Polynomial algebras are examples of *free* graded commutative algebras, having no unnecessary relations. In general, a free graded commutative algebra is a tensor product of single-generator free graded commutative algebras. The latter are either polynomial algebras $R[\alpha]$ on even-dimension generators α or quotients $R[\alpha]/(2\alpha^2)$ with α odd-dimensional. Note that if R is a field then $R[\alpha]/(2\alpha^2)$ is either the exterior algebra $\Lambda_R[\alpha]$ if the characteristic of R is not 2, or the polynomial algebra $R[\alpha]$ otherwise. Every graded commutative algebra is a quotient of a free one, clearly.

Example: Subcomplexes of the n-Torus. To give just a small hint of the endless variety of nonfree cup product algebras which can be realized, consider subcomplexes of the n-torus T^n, the product of n copies of S^1. Here we give S^1 its standard minimal cell structure and T^n the resulting product cell structure. We know that $H^*(T^n;\mathbb{Z})$ is the exterior algebra $\Lambda_{\mathbb{Z}}[\alpha_1,\cdots,\alpha_n]$, with the monomial $\alpha_{i_1}\cdots\alpha_{i_k}$ corresponding via cellular cohomology to the k-cell $e^1_{i_1}\cdots e^1_{i_k}$. So if we pass to a subcomplex $X\subset T^n$ by omitting certain cells, then $H^*(X;\mathbb{Z})$ is the quotient of $\Lambda_{\mathbb{Z}}[\alpha_1,\cdots,\alpha_n]$ obtained by setting the monomials corresponding to the omitted cells equal to zero. Since we are dealing with rings, we are factoring out by an ideal in $\Lambda_{\mathbb{Z}}[\alpha_1,\cdots,\alpha_n]$, the ideal generated by the monomials corresponding to the 'minimal' omitted cells, those whose boundary is entirely contained in X. For example, if we take X to be the subcomplex of T^3 obtained by deleting the cells $e^1_1e^1_2e^1_3$ and $e^1_2e^1_3$, then $H^*(X;\mathbb{Z})\approx\Lambda_{\mathbb{Z}}[\alpha_1,\alpha_2,\alpha_3]/(\alpha_2\alpha_3)$.

How many different subcomplexes of T^n are there? To each subcomplex $X\subset T^n$ we can associate a finite simplicial complex C_X by the following procedure. View T^n as the quotient of the n-cube I^n obtained by identifying opposite faces. If we intersect the cube $I^n\subset\mathbb{R}^n$ with the hyperplane $x_1+\cdots x_n=\varepsilon$ for small $\varepsilon>0$, we get a simplex Δ^{n-1}, and if $q\colon I^n\to T^n$ is the quotient map, then $\Delta^{n-1}\cap q^{-1}(X)$ is a subcomplex C_X of Δ^{n-1}. The k-cells of X correspond exactly to the $(k-1)$-simplices of C_X. Clearly X is uniquely determined by C_X, and every subcomplex of Δ^{n-1} occurs as C_X for some subcomplex X of T^n. Since every simplicial complex with n vertices is a subcomplex of Δ^{n-1}, we see that T^n has quite a large number of subcomplexes, if n is not too small. Of course, it may be that some of the resulting cohomology rings $H^*(X;\mathbb{Z})$ are isomorphic for different subcomplexes $X\subset T^n$. For example, one could just permute the factors of T^n to change X without affecting its cohomology ring. Whether there are less trivial examples is a harder algebraic problem.

Somewhat more elaborate examples could be produced by looking at subcomplexes of the product of n copies of $\mathbb{C}\mathrm{P}^\infty$, with cohomology rings isomorphic to polynomial rings modulo ideals generated by monomials. Or one could take subcomplexes of a product of

S^1's and $\mathbb{C}\mathrm{P}^\infty$'s. However, this is still a whole lot less complicated than the general case, where one takes free algebras modulo ideals generated by arbitrary polynomials which are homogeneous, having all their terms of the same dimension.

Example: Cohen-Macauley Rings. Consider the quotient space $X = \mathbb{C}\mathrm{P}^\infty/\mathbb{C}\mathrm{P}^{n-1}$. The quotient map $\mathbb{C}\mathrm{P}^\infty \to X$ induces an injection $H^*(X;\mathbb{Z}) \to H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ embedding $H^*(X;\mathbb{Z})$ in $\mathbb{Z}[\alpha]$ as the subring generated by $1, \alpha^n, \alpha^{n+1}, \cdots$. If we view this subring as a module over $\mathbb{Z}[\alpha^n]$, it is free with basis $\{1, \alpha^n, \alpha^{n+1}, \cdots, \alpha^{2n-1}\}$. Thus $H^*(X;\mathbb{Z})$ is an example of a *Cohen-Macauley* ring: A ring which contains a polynomial subring, over which it is a finitely-generated free module. While polynomial cup product rings are rather rare, Cohen-Macauley cup product rings occur much more frequently.

Exercises

1. Assuming as known the cup product structure on the torus $S^1 \times S^1$, compute the cup product structure in $H^*(M_g)$ for M_g the closed orientable surface of genus g, by using the quotient map from M_g to a wedge-sum of g tori, shown below.

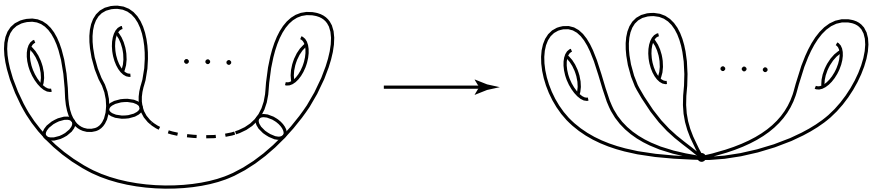

2. Using the cup product $H^k(X, A; R) \times H^\ell(X, B; R) \to H^{k+\ell}(X, A \cup B; R)$, show that if X is the union of contractible open subsets A and B, then all cup products of positive-dimensional classes in $H^*(X;R)$ are zero. This applies in particular if X is a suspension. Generalize to the situation that X is the union of n contractible open subsets, to show that all n-fold cup products of positive-dimensional classes are zero.

3. Using the cup product structure, show there is no map $\mathbb{R}\mathrm{P}^n \to \mathbb{R}\mathrm{P}^m$ inducing a nontrivial map $H^1(\mathbb{R}\mathrm{P}^m;\mathbb{Z}_2) \to H^1(\mathbb{R}\mathrm{P}^n;\mathbb{Z}_2)$ if $n > m$. Deduce the analogous statement for induced maps on π_1, thereby proving the general case of the Borsuk-Ulam theorem. What is the corresponding result for maps $\mathbb{C}\mathrm{P}^n \to \mathbb{C}\mathrm{P}^m$?

4. Apply the Lefschetz fixed point theorem to show that every map $f : \mathbb{C}\mathrm{P}^n \to \mathbb{C}\mathrm{P}^n$ has a fixed point if n is even, using the fact that $f^* : H^*(\mathbb{C}\mathrm{P}^n;\mathbb{Z}) \to H^*(\mathbb{C}\mathrm{P}^n;\mathbb{Z})$ is a ring homomorphism. When n is odd show there is a fixed point unless $f^*(\alpha) = -\alpha$, for α a generator of $H^2(\mathbb{C}\mathrm{P}^n;\mathbb{Z})$. [See §2.4c for an example of a map without fixed points in this exceptional case.]

5. Compute the ring structure in $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_4)$ and more generally $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_m)$.

6. Use cup products to compute the degree of the map $\mathbb{C}\mathrm{P}^n \to \mathbb{C}\mathrm{P}^n$ induced by the map $\mathbb{C}^{n+1} \to \mathbb{C}^{n+1}$ which raises each coordinate to the d^{th} power, $(z_0,\cdots,z_n) \mapsto (z_0^d,\cdots,z_n^d)$, for a fixed integer $d>0$. [First do the case $n=1$.]

7. Use cup products to show that $\mathbb{R}\mathrm{P}^3$ is not homotopy equivalent to $\mathbb{R}\mathrm{P}^2 \vee S^3$.

8. Let X be $\mathbb{C}\mathrm{P}^2$ with a cell e^3 attached by a map $S^2 \to \mathbb{C}\mathrm{P}^1 \subset \mathbb{C}\mathrm{P}^2$ of degree p, and let $Y = M(\mathbb{Z}_p,2) \vee S^4$. Thus X and Y have the same 3-skeleton but differ in the way their 4-cells are attached. Show that X and Y have isomorphic cohomology rings with $\mathbb{Z}$ coefficients but not with $\mathbb{Z}_p$ coefficients.

9. Show that if $H_n(X;\mathbb{Z})$ is finitely generated and free for each n, then $H^*(X;\mathbb{Z}_p) \approx H^*(X;\mathbb{Z}) \otimes \mathbb{Z}_p$ as rings, so in particular the ring structure with $\mathbb{Z}$ coefficients determines the ring structure with $\mathbb{Z}_p$ coefficients.

10. Using cup products, show that a map $S^{k+\ell} \to S^k \times S^\ell$ must induce the trivial homomorphism $H_{k+\ell}(S^{k+\ell}) \to H_{k+\ell}(S^k \times S^\ell)$.

11. Show that the spaces $(S^1 \times \mathbb{C}\mathrm{P}^\infty)/(S^1 \times \{x_0\})$ and $S^3 \times \mathbb{C}\mathrm{P}^\infty$ have isomorphic cohomology rings with $\mathbb{Z}$ or any other coefficients. [An exercise for §4.4k is to show these two spaces are not homotopy equivalent.]

12. Describe $H^*(\mathbb{C}\mathrm{P}^\infty/\mathbb{C}\mathrm{P}^1;\mathbb{Z})$ as a ring with finitely many multiplicative generators. How does this ring compare with $H^*(S^6 \times \mathbb{H}\mathrm{P}^\infty;\mathbb{Z})$?

13. Let $q:\mathbb{R}\mathrm{P}^\infty \to \mathbb{C}\mathrm{P}^\infty$ be the natural quotient map obtained by regarding both spaces as quotients of S^∞, modulo multiplication by real scalars in one case and complex scalars in the other. Show that the induced map $q^*:H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}) \to H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z})$ is surjective in even dimensions by showing first by a geometric argument that the restriction $q:\mathbb{R}\mathrm{P}^2 \to \mathbb{C}\mathrm{P}^1$ induces a surjection on H^2 and then appealing to cup product structures. Next, form a quotient space X of $\mathbb{R}\mathrm{P}^\infty \amalg \mathbb{C}\mathrm{P}^n$ by identifying each point $x \in \mathbb{R}\mathrm{P}^{2n}$ with $q(x) \in \mathbb{C}\mathrm{P}^n$. Show there are ring isomorphisms $H^*(X;\mathbb{Z}) \approx \mathbb{Z}[\alpha]/(2\alpha^{n+1})$ and $H^*(X;\mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha,\beta]/(\beta^2-\alpha^{n+1})$, where α has dimension 2 and β has dimension $2n+1$. Make a similar construction and analysis for the quotient map $q:\mathbb{C}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$.

14. For a fixed coefficient field F, define the **Poincaré series** of a space X to be the formal power series $p(t) = \sum_i a_i t^i$ where a_i is the dimension of $H^i(X;F)$ as a vector space over F, assuming this dimension is finite for all i. Show that $p(X \times Y) = p(X)p(Y)$. Compute the Poincaré series for S^n, $\mathbb{R}\mathrm{P}^n$, $\mathbb{R}\mathrm{P}^\infty$, $\mathbb{C}\mathrm{P}^n$, $\mathbb{C}\mathrm{P}^\infty$, and the spaces in the preceding three exercises.

15. Show that $H^*(J(S^n);\mathbb{Z})$ for n odd is the tensor product of an exterior algebra on an n-dimensional generator with a divided polynomial algebra on a $2n$-dimensional generator.

16. Verify that $H^*(J(S^{2k});\mathbb{Z}_p) \approx \mathbb{Z}_p[x_1, x_p, x_{p^2}, \cdots]/(x_1^p, x_p^p, x_{p^2}^p, \cdots)$ for p prime.

3. Poincaré Duality

Most often, algebraic topology is concerned with properties of spaces which depend only on homotopy type, and for this purpose CW complexes are the natural vehicle. Digressing somewhat from this viewpoint, in this section we study the algebraic topology of manifolds, which are determined instead by local homeomorphism properties, namely that their local topology is as simple as possible, being everywhere locally homeomorphic to $\mathbb{R}^n$. It is somewhat miraculous that just this local homegeneity property, together with global compactness, is enough to impose a strong symmetry on the homology and cohomology groups of such spaces, as well as strong nontriviality of cup products. This is the Poincaré duality theorem, one of the earliest theorems in the subject. In fact, Poincaré's original enunciations of the duality property came before homology and cohomology had even been properly defined, and it took many years for the concepts of homology and cohomology to be refined sufficiently to put Poincaré duality on a firm footing.

Let us proceed now to definitions. A **manifold of dimension** n, or **n-manifold**, is a Hausdorff space M in which each point has an open neighborhood homeomorphic to $\mathbb{R}^n$. The dimension of M is intrinsically characterized by the fact that for $x \in M$, the local homology group $H_i(M, M-\{x\};\mathbb{Z})$ is nonzero only for $i = n$, namely $H_i(M, M-\{x\};\mathbb{Z}) \approx H_i(\mathbb{R}^n, \mathbb{R}^n - \{0\};\mathbb{Z})$ by excision, and $H_i(\mathbb{R}^n, \mathbb{R}^n - \{0\};\mathbb{Z}) \approx \widetilde{H}_{i-1}(\mathbb{R}^n - \{0\};\mathbb{Z}) \approx \widetilde{H}_{i-1}(S^{n-1};\mathbb{Z})$ since $\mathbb{R}^n$ is contractible and $\mathbb{R}^n - \{0\}$ deformation retracts onto S^{n-1}.

A compact manifold is called **closed**, to distinguish it from the more general notion of a compact manifold with boundary which will be considered later in this section. For example S^n is obviously a closed manifold, as are $\mathbb{R}\mathrm{P}^n$ and lens spaces since they have S^n as a covering space. Another closed manifold we have encountered is $\mathbb{C}\mathrm{P}^n$. This is compact since it is a quotient space of S^{2n+1}, and the manifold property is satisfied since $\mathbb{C}\mathrm{P}^n$ has an open cover consisting of the sets $U_i = \{\, [z_0, \cdots, z_n] \in \mathbb{C}\mathrm{P}^n \mid z_i = 1 \,\}$ each of which is homeomorphic to $\mathbb{R}^{2n}$. Further examples of closed manifolds can be generated from these using the obvious fact that the product of closed manifolds of dimensions m and n is a closed manifold of dimension $m+n$.

Poincaré duality in its most primitive form asserts that for a closed orientable manifold M of dimension n, there are isomorphisms $H_k(M;\mathbb{Z}) \approx H^{n-k}(M;\mathbb{Z})$ for all k. Implicit here is the convention that homology and cohomology groups of negative dimension are zero, so all the nontrivial homology and cohomology of M lies in the dimension range from 0 to n. The definition of 'orientable' will be given below. Without the orientability hypothesis there is a weaker statement that $H_k(M;\mathbb{Z}_2) \approx H^{n-k}(M;\mathbb{Z}_2)$ for all k. As we show in Corollaries A.7 and A.8 in the Appendix, the homology groups of a closed manifold

are all finitely generated. So via the universal coefficient theorem, Poincaré duality for a closed orientable n-manifold M can be stated just in terms of homology: Modulo their torsion subgroups, $H_k(M;\mathbb{Z})$ and $H_{n-k}(M;\mathbb{Z})$ are isomorphic, and the torsion subgroups of $H_k(M;\mathbb{Z})$ and $H_{n-k-1}(M;\mathbb{Z})$ are isomorphic. However, the statement in terms of cohomology is really more natural.

Poincaré duality thus expresses a certain symmetry in the homology of closed orientable manifolds. For example, consider the n-dimensional torus T^n, the product of n circles. By induction on n it follows from the Künneth formula, or from the easy special case $H_i(X\times S^1;\mathbb{Z}) \approx H_i(X;\mathbb{Z})\oplus H_{i-1}(X;\mathbb{Z})$ which was an exercise in §2.2, that $H_k(T^n;\mathbb{Z})$ is isomorphic to the sum of $\binom{n}{k}$ copies of $\mathbb{Z}$. So Poincaré duality is reflected in the relation $\binom{n}{k} = \binom{n}{n-k}$. The reader might also check that Poincaré duality is consistent with our calculations of the homology of projective spaces and lens spaces, which are all orientable except for $\mathbb{R}\mathrm{P}^n$ with n even.

For many manifolds there is a very nice geometric proof of Poincaré duality using the notion of *dual cell structures*. The germ of this idea can be traced back to the five regular Platonic solids: the tetrahedron, cube, octahedron, dodecahedron, and icosahedron. Each of these polyhedra has a dual polyhedron whose vertices are the center points of the faces of the given polyhedron. Thus the dual of the cube is the octahedron, and vice versa. Similarly the dodecahedron and icosahedron are dual to each other, and the tetrahedron is its own dual. One can regard each of these polyhedra as defining a cell structure C on S^2 with a dual cell structure C^* determined by the dual polyhedron. Each vertex of C lies in a dual 2-cell of C^*, each edge of C crosses a dual edge of C^*, and each 2-cell of C contains a dual vertex of C^*. The first figure below shows the case of the cube and octahedron.

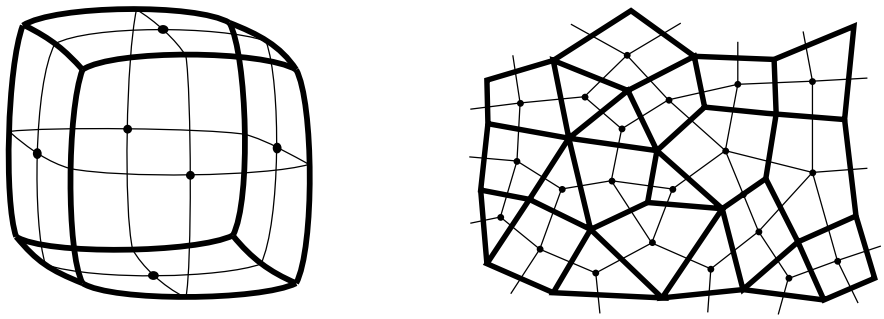

There is no need to restrict to regular polyhedra here, and we can generalize further by replacing S^2 by any surface. A portion of a more-or-less random pair of dual cell structures is shown in the second figure above. On the torus, if we lift a dual pair of cell structures

to the universal cover $\mathbb{R}^2$, we get a dual pair of periodic tilings of the plane, as in the next three figures.

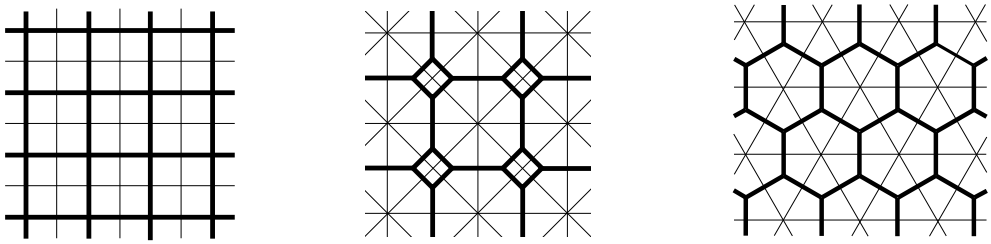

The standard CW structure on the surface of genus g, obtained from a $4g$-gon by identifying edges via the product of commutators $[a_1, b_1] \cdots [a_g, b_g]$, is homeomorphic to its own dual:

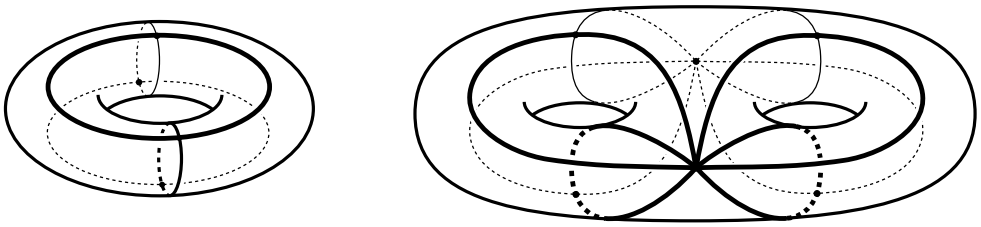

Given a pair of dual cell structures C and C^* on a closed surface M, the pairing of cells with dual cells gives identifications of cellular chain groups $C_0^* = C_2$, $C_1^* = C_1$, and $C_2^* = C_0$. If we use $\mathbb{Z}$ coefficients these identifications are not quite canonical since there is an ambiguity of sign for each cell, the choice of a generator for the corresponding $\mathbb{Z}$ summand of the cellular chain complex. We can avoid this ambiguity by considering the simpler situation of $\mathbb{Z}_2$ coefficients, where the identifications $C_i = C_{2-i}^*$ are completely canonical. The key observation now is that under these identifications, the cellular boundary map $\partial : C_i \to C_{i-1}$ becomes the cellular coboundary map $\delta : C_{2-i}^* \to C_{2-i+1}^*$ since ∂ assigns to a cell the sum of the cells which are its faces while δ assigns to a cell the sum of the cells of which it is a face. Thus $H_i(C; \mathbb{Z}_2) \approx H^{2-i}(C^*; \mathbb{Z}_2)$, and hence $H_i(M; \mathbb{Z}_2) \approx H^{2-i}(M; \mathbb{Z}_2)$ since C and C^* are cell structures on the same surface M.

To refine this argument to $\mathbb{Z}$ coefficients the problem of signs must be addressed. After analyzing the situation more closely, one sees that if M is orientable, it is possible to make consistent choices of orientations of all the cells of C and C^* so that the boundary maps in C agree with the coboundary maps in C^*, and therefore one gets $H_i(C; \mathbb{Z}) \approx H^{2-i}(C^*; \mathbb{Z})$, hence $H_i(M; \mathbb{Z}) \approx H^{2-i}(M; \mathbb{Z})$.

For higher dimensional manifolds the situation is entirely analogous. One would consider dual cell structures C and C^* on a closed n-manifold M, each i-cell of C being

dual to a unique $(n-i)$-cell of C^* which it intersects in one point 'transversely.' For example on the 3-dimensional torus $S^1 \times S^1 \times S^1$ one could take the standard cell structure which lifts to the decomposition of the universal cover $\mathbb{R}^3$ into cubes with vertices at the integer lattice points $\mathbb{Z}^3$, and then the dual cell structure is obtained by translating this by the vector $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. Each edge then has a dual 2-cell which it pierces orthogonally, and each vertex lies in a dual 3-cell.

All the manifolds one commonly meets, for example all differentiable manifolds, have dually paired cell structures with the properties needed to carry out the proof of Poincaré duality we have just sketched. However, to construct these cell structures requires a certain amount of manifold theory. To avoid this, and to get a theorem which applies to all manifolds, we shall take a completely different approach, using algebraic topology to replace the geometry of dual cell structures.

Orientations

Let us consider now how one might define orientability for manifolds. In linear algebra an orientation of $\mathbb{R}^n$ is defined to be an equivalence class of ordered bases, two ordered bases being equivalent if one is obtained from the other by a linear transformation of positive determinant. Thus an invertible linear transformation of $\mathbb{R}^n$ preserves or reverses orientation depending on the sign of its determinant. For a smooth surface $M \subset \mathbb{R}^3$ an orientation can be defined as a choice of an orientation for each tangent plane to M, varying continuously over M in the sense that at nearby points $x, y \in M$ whose tangent planes T_x and T_y are nearly parallel, the orientation of T_y is obtained from that for T_x by orthogonal projection of an ordered basis for T_x onto an ordered basis for T_y. Of course, it may not be possible to make such a continuous choice of orientations for all the tangent planes to M, for example for the Möbius band, so not all surfaces are orientable.

Using basic machinery of differential topology this definition can be generalized to higher dimensional smooth manifolds, but we shall take a different viewpoint, more in the realm of algebraic topology. Taking homology with $\mathbb{Z}$ coefficients, the idea is to regard the two possible choices of *generators* of $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\}) \approx \mathbb{Z}$ as the two choices of *orientations* of $\mathbb{R}^n$. This can be justified heuristically as follows. If we position the standard n-simplex Δ^n in $\mathbb{R}^n$ so that 0 lies in its interior, then the inclusion $\Delta^n \hookrightarrow \mathbb{R}^n$ represents a generator ι of $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\})$ since the inclusion $\Delta^n \hookrightarrow \mathbb{R}^n$ induces an isomorphism $H_n(\Delta^n, \partial\Delta^n) \approx H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\})$ and $H_n(\Delta^n, \partial\Delta^n)$ is generated by the identity map $\Delta^n \to \Delta^n$. If $f : \mathbb{R}^n \to \mathbb{R}^n$ is an invertible linear transformation, then $f \,|\, \Delta^n$ is a relative cycle representing the class $f_*(\iota) \in H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\})$. By an exercise in §2.2, $f_*(\iota)$ equals ι or $-\iota$ according to the sign of the determinant of f. Thus f_* preserves

generators of $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\})$ iff f preserves orientation of $\mathbb{R}^n$, so it becomes natural to identify the two generators of $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\})$ with the two orientations of $\mathbb{R}^n$.

Now fix an n-dimensional manifold M. In view of the preceding remarks it is reasonable to define a **local orientation** at a point $x \in M$ to be a choice of generator μ_x for $H_n(M, M - \{x\})$. A global **orientation** of M is a function $x \mapsto \mu_x$ assigning to each $x \in M$ a local orientation at x, subject to the following 'local consistency' condition: Each $x \in M$ has a neighborhood $\mathbb{R}^n \subset M$ containing an open ball B of finite radius about x, such that all the local orientations μ_y at points $y \in B$ are the images of one generator μ_B of $H_n(M, M - B) \approx H_n(\mathbb{R}^n, \mathbb{R}^n - B) \approx \mathbb{Z}$ under the natural maps $H_n(M, M - B) \to H_n(M, M - \{y\})$. If an orientation exists for M, then M is called **orientable**.

Every manifold M has a two-sheeted covering space $\widetilde{M}$ which is an orientable manifold. For example, $\mathbb{R}\mathrm{P}^2$ is covered by S^2, and the Klein bottle has the torus as a two-sheeted covering space. The general construction goes as follows. As a set, let

$$\widetilde{M} = \{\, (x, \mu_x) \mid x \in M \text{ and } \mu_x \text{ is a local orientation at } x \,\}.$$

Given an open ball $B \subset \mathbb{R}^n \subset M$ of finite radius and a generator $\mu_B \in H_n(M, M - B)$, let $U(\mu_B)$ be the set of pairs (x, μ_x) where $x \in B$ and μ_x is the image of μ_B under the natural map $H_n(M, M-B) \to H_n(M, M-\{x\})$. It is easy to check that these sets $U(\mu_B)$ form a basis for a topology on $\widetilde{M}$ such that the projection $\widetilde{M} \to M$, $(x, \mu_x) \mapsto x$, is a two-sheeted covering space. The manifold $\widetilde{M}$ is orientable since each point (x, μ_x) has the canonical local orientation μ_x, and by construction these local orientations satisfy the local consistency condition necessary to define a global orientation.

Proposition 3.14. *If M is connected, then M is orientable iff $\widetilde{M}$ has two components. In particular, M is orientable if it is simply-connected, or more generally if $\pi_1(M)$ has no subgroup of index two.*

Proof: If M is connected, $\widetilde{M}$ has either one or two components since it is a two-sheeted covering space of M. If it has two components, they are each mapped homeomorphically to M by the covering projection, so M is orientable, being homeomorphic to a component of the orientable manifold $\widetilde{M}$. Conversely, if M is orientable, it has exactly two orientations since it is connected, and each of these orientations defines a component of $\widetilde{M}$. The last statement of the proposition follows since two-sheeted covering spaces of M correspond to index-two subgroups of $\pi_1(M)$, by the classification of covering spaces. □

One can generalize the definition of orientation by replacing the coefficient group $\mathbb{Z}$ by any commutative ring R with identity. Then an ***R*-orientation** of M assigns to

each $x \in M$ a generator of $H_n(M, M-\{x\}; R) \approx R$, subject to the corresponding local consistency condition, where a 'generator' of R is an element u such that $Ru = R$. Since R has an identity element, this is equivalent to u being a unit, i.e., an invertible element of R. Since $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\}; R)$ is canonically isomorphic to $H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\}) \otimes R$, an orientable manifold is R-orientable for all R. Every manifold is $\mathbb{Z}_2$-orientable since $\mathbb{Z}_2$ has a unique generator.

In what follows we shall not restrict R to $\mathbb{Z}$ or $\mathbb{Z}_2$, but these are the two most important cases, and the reader should keep them foremost in mind.

Cap Product

When M is an R-orientable closed n-manifold, we would like to define homomorphisms $D: H^k(M; R) \to H_{n-k}(M; R)$ which Poincaré duality asserts are isomorphisms. The definition will be in terms of a more general construction called *cap product*.

For an arbitrary space X and coefficient ring R, define an R-bilinear cap product map $\frown : C_k(X; R) \times C^\ell(X; R) \to C_{k-\ell}(X; R)$ for $k \geq \ell$ by

$$\sigma \frown \varphi = \varphi(\sigma\big|[v_0, \cdots, v_\ell])\, \sigma\big|[v_\ell, \cdots, v_k] \qquad \text{for } \sigma: \Delta^k \to X \quad \text{and} \quad \varphi \in C^\ell(X; R).$$

One checks the formula

$$\partial(\sigma \frown \varphi) = (-1)^\ell(\partial\sigma \frown \varphi \;-\; \sigma \frown \delta\varphi)$$

by computing:

$$\begin{aligned}
\partial\sigma \frown \varphi &= \sum_{i=0}^{\ell} (-1)^i \varphi(\sigma\big|[v_0, \cdots, \widehat{v}_i, \cdots, v_{\ell+1}])\sigma\big|[v_{\ell+1}, \cdots, v_k] \\
&\qquad + \sum_{i=\ell+1}^{k} (-1)^i \varphi(\sigma\big|[v_0, \cdots, v_\ell])\sigma\big|[v_\ell, \cdots, \widehat{v}_i, \cdots, v_k] \\
\sigma \frown \delta\varphi &= \sum_{i=0}^{\ell+1} (-1)^i \varphi(\sigma\big|[v_0, \cdots, \widehat{v}_i, \cdots, v_{\ell+1}])\sigma\big|[v_{\ell+1}, \cdots, v_k] \\
\partial(\sigma \frown \varphi) &= \sum_{i=\ell}^{k} (-1)^{i-\ell} \varphi(\sigma\big|[v_0, \cdots, v_\ell])\sigma\big|[v_\ell, \cdots, \widehat{v}_i, \cdots, v_k]
\end{aligned}$$

From the relation $\partial(\sigma \frown \varphi) = \pm(\partial\sigma \frown \varphi \;-\; \sigma \frown \delta\varphi)$ we see that the cap product of a cycle σ and a cocycle φ is a cycle. Further, if $\partial\sigma = 0$ then $\partial(\sigma \frown \varphi) = \pm(\sigma \frown \delta\varphi)$, so the cap product of a cycle and a coboundary is a boundary. And if $\delta\varphi = 0$ then

$\partial(\sigma \frown \varphi) = \pm(\partial\sigma \frown \varphi)$, so the cap product of a boundary and a cocycle is a boundary. These facts imply that there is an induced cap product

$$H_k(X;R) \times H^\ell(X;R) \xrightarrow{\frown} H_{k-\ell}(X;R)$$

This is clearly R-linear in each variable.

Using the same formulas one checks that cap product has the relative forms

$$H_k(X,A;R) \times H^\ell(X;R) \xrightarrow{\frown} H_{k-\ell}(X,A;R)$$

$$H_k(X,A;R) \times H^\ell(X,A;R) \xrightarrow{\frown} H_{k-\ell}(X;R)$$

For example, in the second case the cap product of a relative cycle σ with a relative cocycle φ is an absolute cycle since $\varphi(\partial\sigma) = 0$. If A and B are open sets in X one can define a more general cap product

$$H_k(X, A\cup B;R) \times H^\ell(X,A;R) \xrightarrow{\frown} H_{k-\ell}(X,B;R),$$

using the fact that $H_k(X, A\cup B;R)$ can be computed using chains which are sums of chains in A and chains in B, as in the derivation of relative Mayer-Vietoris sequences in §2.2.

Cap product satisfies a naturality property which may look a bit strange because both covariant and contravariant functors are involved: $f_*\big(\alpha \frown f^*(\varphi)\big) = f_*(\alpha) \frown \varphi$ for $f : X \to Y$. The following diagram may help clarify the formula:

$$\begin{array}{ccccc} H_k(X) & \times & H^\ell(X) & \xrightarrow{\frown} & H_{k-\ell}(X) \\ \big\downarrow f_* & & \big\uparrow f^* & & \big\downarrow f_* \\ H_k(Y) & \times & H^\ell(Y) & \xrightarrow{\frown} & H_{k-\ell}(Y) \end{array}$$

The proof of the formula is left to the exercises. There are relative versions as well.

Now we can state Poincaré duality for closed manifolds:

Theorem 3.15. *If M is a closed R-orientable n-manifold, then there is a class $[M] \in H_n(M;R)$ whose image in $H_n(M, M-\{x\};R)$ is a generator for all $x \in M$, and for such a class $[M]$ the map $D : H^k(M;R) \longrightarrow H_{n-k}(M;R)$ defined by $D(\varphi) = [M] \frown \varphi$ is an isomorphism for all k.*

A class $[M]$ whose image in $H_n(M, M-\{x\};R)$ is a generator for all $x \in M$ is called a **fundamental class** or **orientation class** for M. If a fundamental class $[M] \in H_n(M;R)$ exists, then M is R-orientable since the natural map $H_n(M;R) \to H_n(M, M-\{x\};R)$ factors through $H_n(M, M-B;R)$ for B any open ball in M.

Example: Surfaces. Let M be the closed orientable surface of genus g, obtained from a wedge sum of $2g$ circles by attaching a 2-cell via the word $a_1b_1a_1^{-1}b_1^{-1}\cdots a_gb_ga_g^{-1}b_g^{-1}$. We give M the Δ-complex structure indicated in the first figure below for the case $g = 2$, by coning off a $4g$-gon to its center. We can compute cap products using simplicial homology and cohomology since cap products are defined for simplicial homology and cohomology by exactly the same formula as for singular homology and cohomology, so the isomorphism between the simplicial and singular theories respects cap products. A fundamental class $[M]$ generating $H_2(M)$ is represented by the 2-cycle which is the sum of all $4g$ 2-simplices with the signs indicated. The edges a_i and b_i form a basis for $H_1(M)$. Under the

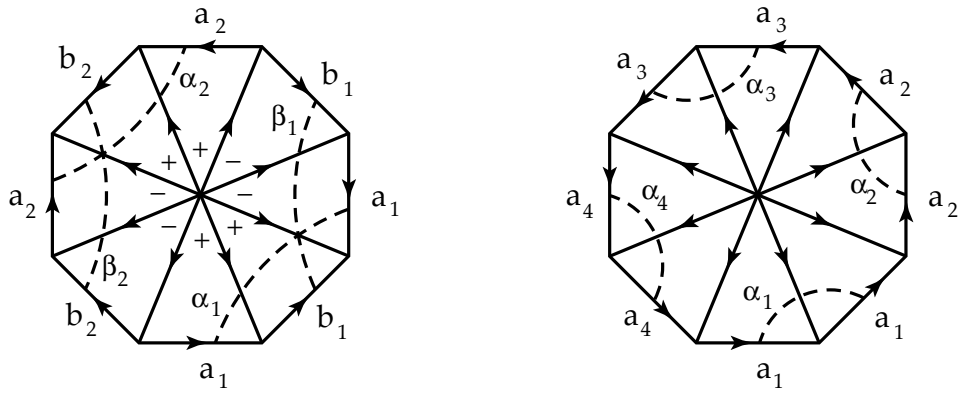

isomorphism $H^1(M) \approx Hom(H_1(M), \mathbb{Z})$, the cohomology class α_i corresponding to a_i assigns the value 1 to a_i and 0 to the other basis elements. This class α_i is represented by the cocycle φ_i assigning the value 1 to the 1-simplices meeting the arc labeled α_i in the figure and 0 to the other 1-simplices. Similarly we have a class β_i corresponding to b_i, represented by the cycle ψ_i assigning the value 1 to the 1-simplices meeting the arc β_i. Applying the definition of cap product, we have $[M] \frown \varphi_i = b_i$ and $[M] \frown \psi_i = -a_i$ since in both cases there is just one 2-simplex $[v_0, v_1, v_2]$ where φ_i or ψ_i is nonzero on the edge $[v_0, v_1]$. Thus b_i is the Poincaré dual of α_i and $-a_i$ is the Poincaré dual of β_i. If we interpret Poincaré duality entirely in terms of homology, identifying α_i with its Hom-dual a_i and β_i with b_i, then the classes a_i and b_i are duals of each other, up to sign at least. One can see from a picture of M that all the loops a_i and b_i in M can be perturbed slightly to be disjoint, except that a_i crosses b_i once transversely.

If we use $\mathbb{Z}_2$ coefficients, the closed nonorientable surface N of genus g can be treated in the same way. Here we form N from a $2g$-gon by identifying consecutive pairs of edges according to the word $a_1^2\cdots a_g^2$ as in the second figure above. We have classes $\alpha_i \in H^1(N; \mathbb{Z}_2)$ represented by cocycles φ_i assigning the value 1 to the edges meeting the arc α_i. One has $[N] \frown \varphi_i = a_i$, so a_i is the Poincaré dual of α_i, In terms of homology, a_i is its own dual. Geometrically, the loops a_i on N are homotopic to the loops α_i, which

are all disjoint from each other, while a_i crosses α_i transversely in a single point. So one sees again, just as in the orientable case, that the homology version of Poincaré duality has something to do with intersections of cycles.

In these examples the fundamental class $[M] \in H_n(M;R)$ consisted of the sum of all the n-dimensional simplices of M with appropriate signs, and this holds more generally for all closed manifolds M having simplicial or Δ-complex structures. Consider the case of $\mathbb{Z}$ coefficients, for example. Thinking of simplicial homology, the fundamental class $[M]$ must be represented by some linear combination $\sum_i k_i\sigma_i$ of the n-simplices σ_i of M. The condition that $[M]$ maps to a generator of $H_n(M, M-\{x\};\mathbb{Z})$ for points x in the interiors of the σ_i's means that each coefficient k_i must be ± 1. The k_i's must also be such that $\partial(\sum_i k_i\sigma_i) = 0$. This means that if σ_i and σ_j share a common $(n-1)$-dimensional face, then k_i determines k_j and vice versa. Analyzing the situation more closely, one can show that a choice of signs for the k_i's making $\sum_i k_i\sigma_i$ a cycle is possible iff M is orientable. However, we shall not pursue this simplicial viewpoint further since we are trying to avoid any extraneous conditions on M like the existence of simplicial structures.

Proof of Poincaré Duality

The first step is to construct fundamental classes.

Lemma 3.16. *Let M be a manifold of dimension n and let $A \subset M$ be a compact subset. Then:*

(a) *$H_i(M, M-A;R) = 0$ for $i > n$, and a class in $H_n(M, M-A;R)$ is zero iff its image in $H_n(M, M-\{x\};R)$ is zero for all $x \in A$.*

(b) *If $x \mapsto \mu_x$ is an R-orientation of M, then there is a unique class $\mu_A \in H_n(M, M-A;R)$ whose image in $H_n(M, M-\{x\};R)$ is μ_x for all $x \in A$.*

In particular if M is closed and orientable then taking $A = M$ in (b) we get a fundamental class for M.

Proof: The coefficient ring R will play no special role in the argument so we shall omit it from the notation. Note that the uniqueness part of (b) follows from (a). We break the proof up into several steps.

(1) First we observe that if the lemma is true for compact sets A, B, and $A \cap B$ then it is true for $A \cup B$. To see this, consider the Mayer-Vietoris sequence:

$$0 \longrightarrow H_n(M, M-A\cup B) \xrightarrow{\Phi} H_n(M, M-A) \oplus H_n(M, M-B) \xrightarrow{\Psi} H_n(M, M-A\cap B)$$

Here the zero on the left comes from the assumption $H_{n+1}(M, M-A\cap B) = 0$. The map Φ is $\Phi(\alpha) = (\alpha, -\alpha)$ and Ψ is $\Psi(\alpha,\beta) = \alpha + \beta$, where we omit notation for maps

on homology induced by inclusion. The terms $H_i(M, M - A \cup B)$ farther to the left in this sequence are sandwiched between groups which are zero, so $H_i(M, M - A \cup B) = 0$ for $i > n$. If a class $\alpha \in H_n(M, M - A \cup B)$ has image zero in $H_n(M, M - \{x\})$ for all $x \in A \cup B$, its images in $H_n(M, M - A)$ and $H_n(M, M - B)$ have the same property, hence are zero by hypothesis, so α itself must be zero since Φ is injective. This gives (a). For (b) the hypothesis gives classes $\mu_A \in H_n(M, M - A)$ and $\mu_B \in H_n(M, M - B)$ having image μ_x for all appropriate x. The images of μ_A and μ_B in $H_n(M, M - A \cap B)$ then satisfy the defining property of $\mu_{A\cap B}$, hence must equal $\mu_{A\cap B}$. By exactness of the sequence, $(\mu_A, -\mu_B)$ is therefore the image of an element $\mu_{A\cup B} \in H_n(M, M - A \cup B)$, and this $\mu_{A\cup B}$ has image μ_x for all $x \in A \cup B$ since μ_A and μ_B have this property.

(2) Next we reduce to the case $M = \mathbb{R}^n$. A compact set $A \subset M$ can be written as the union of finitely many compact sets $A_1, \cdots, A_m$ each contained in an open $\mathbb{R}^n \subset M$. We apply the result in (1) to $A_1 \cup \cdots \cup A_{m-1}$ and A_m, intersecting in $(A_1 \cap A_m) \cup \cdots \cup (A_{m-1} \cap A_m)$, a union of $m-1$ compact sets each contained in an $\mathbb{R}^n \subset M$. By induction on m this reduces us to the case $m = 1$, where by excision we may replace M by the neighborhood $\mathbb{R}^n \subset M$.

(3) When $M = \mathbb{R}^n$ and A is a finite simplicial complex K with simplices linearly embedded in $\mathbb{R}^n$, the result follows by induction on the number of open simplices of K, i.e., cells of K viewed as a CW complex, using (1) along with the easy case that K is a single simplex.

(4) For an arbitrary compact set $A \subset \mathbb{R}^n$ let $\alpha \in H_i(\mathbb{R}^n, \mathbb{R}^n - A)$ be represented by a relative cycle z, with $C \subset \mathbb{R}^n - A$ the union of the images of the singular simplices in ∂z. Since C is compact, it has a positive distance δ from A. Choose a large simplex $\Delta^n \subset \mathbb{R}^n$ containing A in its interior, and subdivide Δ^n, say by repeated barycentric subdivision, so that all its simplices have diameter less than $\delta/2$. Let K be the union of all the simplices of this subdivision which meet A, so K is a finite simplicial complex containing A and disjoint from C. The relative cycle z defines an element $\alpha_K \in H_i(\mathbb{R}^n, \mathbb{R}^n - K)$ mapping to the given $\alpha \in H_i(\mathbb{R}^n, \mathbb{R}^n - A)$. If $i > n$ then by (3) we have $H_i(\mathbb{R}^n, \mathbb{R}^n - K) = 0$, so $\alpha_K = 0$ implies $\alpha = 0$ and hence $H_i(\mathbb{R}^n, \mathbb{R}^n - A) = 0$. If $i = n$ and α_x is zero in $H_n(\mathbb{R}^n, \mathbb{R}^n - \{x\})$ for all $x \in A$ then in fact this holds for all $x \in K$ since K is a union of simplices σ meeting A and $H_n(\mathbb{R}^n, \mathbb{R}^n - \sigma) \to H_n(\mathbb{R}^n, \mathbb{R}^n - \{x\})$ is an isomorphism for all $x \in \sigma$. By (3), α_K is then zero, hence also α, so (a) is finished. Part (b) for $M = \mathbb{R}^n$ is obvious — just take a large simplex $\Delta^n \supset A$ and use the fact that $H_n(\mathbb{R}^n, \mathbb{R}^n - \Delta^n) \to H_n(\mathbb{R}^n, \mathbb{R}^n - \{x\})$ is an isomorphism for all $x \in \Delta^n$. □

Like the construction of fundamental classes, our proof of Poincaré duality will be

by an inductive argument using Mayer-Vietoris sequences. The induction step requires a version of Poincaré duality for open subsets of M, which are noncompact and can satisfy Poincaré duality only when a different kind of cohomology called *cohomology with compact supports* is used. We will define this as a limit of ordinary cohomology groups, so we first need to define a suitable notion of algebraic limit of groups.

Suppose one has abelian groups G_α indexed by some partially ordered index set I having the property that for each pair $\alpha, \beta \in I$ there exists $\gamma \in I$ with $\alpha \le \gamma$ and $\beta \le \gamma$. Such an I is called a **directed set**. Suppose also that for each pair $\alpha < \beta$ one has a homomorphism $f_{\alpha\beta} : G_\alpha \to G_\beta$, such that if $\alpha < \beta < \gamma$ then $f_{\alpha\gamma}$ is the composition of $f_{\alpha\beta}$ and $f_{\beta\gamma}$. Then the **direct limit** group $\varinjlim G_\alpha$ is defined to be the quotient of the direct sum $\bigoplus_\alpha G_\alpha$ by the subgroup generated by elements $a - f_{\alpha\beta}(a)$ for $a \in G_\alpha$. It follows from the directed set property that every element of $\varinjlim G_\alpha$ is represented by an element $a \in G_\alpha$ for some α, and two elements $a \in G_\alpha$ and $b \in G_\beta$ are equal in $\varinjlim G_\alpha$ iff $f_{\alpha\gamma}(a) = f_{\beta\gamma}(b)$ for some γ. We could in fact use this property to define $\varinjlim G_\alpha$ as a quotient of the disjoint union of the G_α's under the equivalence relation $a \sim b$ iff $f_{\alpha\gamma}(a) = f_{\beta\gamma}(b)$ for some γ, but then we would have to check that the axioms for a group are satisfied.

A useful consequence of the definition is that if we have a subset $J \subset I$ with the property that for each $\alpha \in I$ there exists a $\beta \in J$ with $\alpha \le \beta$, then $\varinjlim G_\alpha$ is the same whether we compute it with α varying over I or just over J. In particular, if I has a maximal element γ, we can take $J = \{\gamma\}$ and then $\varinjlim G_\alpha = G_\gamma$.

Suppose for example that a space X is the union of a collection of subspaces X_α which form a directed set with respect to the inclusion relation. Then the groups $H_i(X_\alpha; G)$ for fixed i and G form a directed system, and the natural maps $H_i(X_\alpha; G) \to H_i(X; G)$ induce a homomorphism $\varinjlim H_i(X_\alpha; G) \to H_i(X; G)$.

Proposition 3.17. *If a space X is the union of a directed set of subspaces X_α with the property that each compact set in X is contained in some X_α, then the natural map $\varinjlim H_i(X_\alpha; G) \to H_i(X; G)$ is an isomorphism for all i and G.*

Proof: For surjectivity, represent a cycle in X by a finite sum of singular simplices. The union of the images of these singular simplices is compact in X, hence lies in some X_α, so the map $\varinjlim H_i(X_\alpha; G) \to H_i(X; G)$ is surjective. Injectivity is similar: If a cycle in some X_α is a boundary in X, compactness implies it is a boundary in some $X_\beta \supset X_\alpha$, hence represents zero in $\varinjlim H_i(X_\alpha; G)$. □

The cohomology of direct limits is not so easily described (see §3.4f), but there is a

variant of cohomology which does behave reasonably well with respect to direct limits, and which we are now ready to define.

For a space X, the compact subsets $K \subset X$ form a directed set under inclusion since the union of two compact sets is compact. To each compact $K \subset X$ we associate the group $H^k(X, X-K; G)$, with a fixed k and coefficient group G, and to each inclusion $K \subset L$ of compact sets we associate the natural homomorphism $H^k(X, X-K; G) \to H^k(X, X-L; G)$. The resulting limit group $\varinjlim H^k(X, X-K; G)$ is denoted $H_c^k(X; G)$ and called **cohomology with compact supports**. This name refers to the fact that elements of $H_c^k(X; G)$ are represented by cochains φ which vanish on all singular simplices in the complement of a compact subset K_φ of X. Note that if X is compact, then $H_c^k(X; G) = H^k(X; G)$ since there is a unique maximal compact set $K \subset X$, namely X itself.

As an example, consider $H_c^k(\mathbb{R}^n; G)$. To compute $\varinjlim H^k(\mathbb{R}^n, \mathbb{R}^n - K; G)$ it suffices to let K range over balls B_i of integer radius i centered at the origin since every compact set is contained in such a ball. Since $H^k(\mathbb{R}^n, \mathbb{R}^n - B_i; G)$ is nonzero only for $k = n$, when it is G, and the maps $H^n(\mathbb{R}^n, \mathbb{R}^n - B_i; G) \to H^n(\mathbb{R}^n, \mathbb{R}^n - B_{i+1}; G)$ are isomorphisms, we deduce that $H_c^k(\mathbb{R}^n; G) = 0$ for $k \neq n$ and $H_c^n(\mathbb{R}^n; G) \approx G$.

Let M be an R-orientable n-manifold, possibly noncompact. For an open set $U \subset M$ and compact sets $K \subset L \subset U$ we then have a diagram

$$\begin{array}{ccccc} H_n(U, U-K; R) & \times & H^k(U, U-K; R) & \xrightarrow{\frown} & H_{n-k}(U; R) \\ \uparrow{\scriptstyle i_*} & & \downarrow{\scriptstyle i^*} & & \| \\ H_n(U, U-L; R) & \times & H^k(U, U-L; R) & \xrightarrow{\frown} & H_{n-k}(U; R) \end{array}$$

An R-orientation of M determines an R-orientation of U, so by Lemma 3.16 there are unique elements $\mu_K \in H_n(U, U-K; R)$ and $\mu_L \in H_n(U, U-L; R)$ restricting to the given orientation at each point of K and L, respectively, and $i_*(\mu_L) = \mu_K$. By naturality of cap product, $i_*(\mu_L) \frown x = \mu_K \frown i^*(x)$ for all $x \in H^k(U, U-K; R)$. Therefore, letting K vary over compact sets in U, the homomorphisms $H^k(U, U-K; R) \to H_{n-k}(U; R)$, $x \mapsto \mu_K \frown x$, induce in the limit a duality homomorphism $D_U : H_c^k(U; R) \to H_{n-k}(U; R)$.

Since $H_c^*(M; R) = H^*(M; R)$ if M is compact, the following result generalizes Poincaré duality for closed manifolds.

Theorem 3.18. *Let M an R-orientable n-manifold which is the union of a finite or countably infinite number of open neighborhoods homeomorphic to $\mathbb{R}^n$. Then the duality map $D_M : H_c^k(M; R) \to H_{n-k}(M; R)$ is an isomorphism for all k.*

To simplify notation we shall omit the coefficient ring R throughout the proof. The main remaining technical step is:

Lemma 3.19. *If M is the union of two open sets U and V, then there is a commutative diagram of Mayer-Vietoris sequences:*

$$\begin{array}{ccccccccc}
\longrightarrow & H_c^k(U\cap V) & \longrightarrow & H_c^k(U)\bigoplus H_c^k(V) & \longrightarrow & H_c^k(M) & \longrightarrow & H_c^{k+1}(U\cap V) & \longrightarrow \\
& \downarrow D_{U\cap V} & & \downarrow D_U \oplus -D_V & & \downarrow D_M & & \downarrow D_{U\cap V} & \\
\to & H_{n-k}(U\cap V) & \to & H_{n-k}(U)\bigoplus H_{n-k}(V) & \to & H_{n-k}(M) & \to & H_{n-k-1}(U\cap V) & \to
\end{array}$$

Proof: Compact sets $K \subset U$ and $L \subset V$ give rise to the Mayer-Vietoris sequence in the upper row of the following diagram, whose lower row is also a Mayer-Vietoris sequence.

$$\begin{array}{ccccccc}
\longrightarrow & H^k(M, M-K\cap L) & \longrightarrow & H^k(M,M-K)\bigoplus H^k(M,M-L) & \to & H^k(M,M-K\cup L) & \to \\
& \downarrow\approx & & \downarrow\approx & & & \\
& H^k(U\cap V, U\cap V - K\cap L) & & H^k(U,U-K)\bigoplus H^k(V,V-L) & & \downarrow \mu_{K\cup L}\frown & \\
& \downarrow \mu_{K\cap L}\frown & & \downarrow \mu_K\frown \oplus -\mu_L\frown & & & \\
\longrightarrow & H_{n-k}(U\cap V) & \longrightarrow & H_{n-k}(U)\bigoplus H_{n-k}(V) & \longrightarrow & H_{n-k}(M) & \longrightarrow
\end{array}$$

The two maps labelled isomorphisms come from excision. Assuming this diagram commutes, consider passing to the limit over compact sets $K \subset U$ and $L \subset V$. Since each compact set in $U \cap V$ is contained in an intersection $K \cap L$ of compact sets $K \subset U$ and $L \subset V$, and similarly for $U \cup V$, the diagram induces a limit diagram having the form stated in the lemma. The first row of this limit diagram is exact since a direct limit of exact sequences is exact; this is an exercise at the end of the section, and follows easily from the definition of direct limits.

It remains to consider the commutativity of the preceding diagram involving K and L. In the two squares shown, not involving boundary or coboundary maps, it is a triviality to check commutativity at the level of cycles and cocycles. Less trivial is the third square:

$$(*)\qquad \begin{array}{ccc}
H^k(M, M-K\cup L) & \xrightarrow{\delta} & H^{k+1}(M, M-K\cap L) \xrightarrow{\approx} H^{k+1}(U\cap V, U\cap V - K\cap L) \\
\downarrow \mu_{K\cup L}\frown & & \downarrow \mu_{K\cap L}\frown \\
H_{n-k}(M) & \xrightarrow{\qquad\qquad\partial\qquad\qquad} & H_{n-k-1}(U\cap V)
\end{array}$$

Letting $A = M - K$ and $B = M - L$, the map δ is the coboundary map in the Mayer-Vietoris sequence obtained from the short exact sequence of cochain complexes

$$0 \longrightarrow C^*(M, A+B) \longrightarrow C^*(M,A) \oplus C^*(M,B) \longrightarrow C^*(M, A\cap B) \longrightarrow 0$$

where $C^*(M, A+B)$ consists of cochains on M vanishing on sums of chains in A and chains in B. To evaluate δ on a cocycle $\varphi \in C^*(M, A\cap B)$ one first writes $\varphi = \varphi_A - \varphi_B$ for $\varphi_A \in C^*(M, A)$ and $\varphi_B \in C^*(M, B)$, and then $\delta[\varphi]$ is represented by the cocycle $\delta\varphi_A = \delta\varphi_B \in C^*(M, A+B)$. Similarly, the boundary map ∂ in the homology Mayer-Vietoris sequence is obtained by representing an element of $H_i(M)$ by a cycle z which is a sum of chains $z_U \in C_i(U)$ and $z_V \in C_i(V)$, and then $\partial[z] = [\partial z_U]$.

Via barycentric subdivision, the class $\mu_{K\cup L}$ can be represented by a chain α which is a sum $\alpha_{U-L}+\alpha_{U\cap V}+\alpha_{V-K}$ of chains in $U-L$, $U\cap V$, and $V-K$, respectively, since these three open sets cover M. Then $\alpha_{U\cap V}$ represents $\mu_{U\cap V}$ and $\alpha_{U-L}+\alpha_{U\cap V}$ represents μ_K.

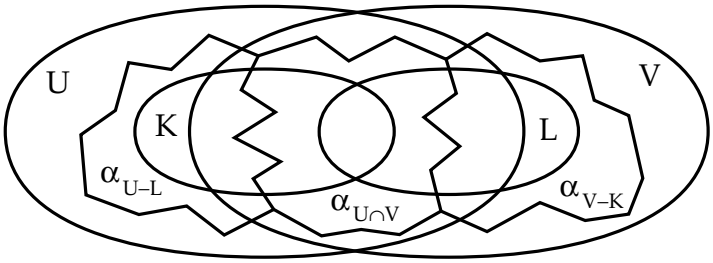

In the square $(*)$ let φ be a cocycle representing an element of $H^k(M, M-K\cup L)$. Under δ this maps to the cohomology class of $\delta\varphi_A$, and continuing on to $H_{n-k-1}(U\cap V)$ we obtain $\alpha_{U\cap V}\frown\delta\varphi_A$, which is in the same homology class as $\partial\alpha_{U\cap V}\frown\varphi_A$ since

$$\partial(\alpha_{U\cap V}\frown\varphi_A) = (-1)^k(\partial\alpha_{U\cap V}\frown\varphi_A \;-\; \alpha_{U\cap V}\frown\delta\varphi_A)$$

and $\alpha_{U\cap V}\frown\varphi_A$ is a chain in $U\cap V$.

Going around the square $(*)$ the other way, φ maps first to $\alpha\frown\varphi$. To apply the Mayer-Vietoris boundary map ∂ to this, we write $\alpha\frown\varphi$ as a sum of a chain in U and a chain in V,

$$\alpha\frown\varphi = (\alpha_{U-L}\frown\varphi) + (\alpha_{U\cap V}\frown\varphi + \alpha_{V-K}\frown\varphi)$$

then we take the boundary of the first of these two chains, obtaining $[\partial(\alpha_{U-L}\frown\varphi)] \in H_{n-k-1}(U\cap V)$. It remains to compare this with $[\partial\alpha_{U\cap V}\frown\varphi_A]$.

We have

$$\begin{aligned}
\partial(\alpha_{U-L}\frown\varphi) &= (-1)^k\partial\alpha_{U-L}\frown\varphi && \text{since } \delta\varphi = 0\\
&= (-1)^k\partial\alpha_{U-L}\frown\varphi_A && \text{since } \partial\alpha_{U-L}\frown\varphi_B = 0,\quad \varphi_B \text{ being}\\
& && \text{zero on chains in } B = M-L\\
&= (-1)^{k+1}(\partial\alpha_{U\cap V}\frown\varphi_A)
\end{aligned}$$

this last equality coming from the fact that $\partial(\alpha_{U-L} + \alpha_{U\cap V}) \frown \varphi_A = 0$ since $\partial(\alpha_{U-L} + \alpha_{U\cap V})$ is a chain in $U - K$ and φ_A vanishes on chains in $A = M - K$.

Thus the square commutes up to a sign depending only on k, so by incorporating this sign into δ or ∂ we get a commutative diagram. □

Proof of 3.18: There are two inductive steps, finite and infinite:

(A) If M is the union of open sets U and V and if D_U, D_V, and $D_{U\cap V}$ are isomorphisms, then so is D_M. Via the five-lemma, this is immediate from the preceding lemma.

(B) If U is the union of a sequence of open sets $U_1 \subset U_2 \subset \cdots$ and each $D_{U_i} : H_c^k(U_i) \to H_{n-k}(U_i)$ is an isomorphism, then so is D_U. To show this we notice first that by excision, $H_c^k(U_i)$ can be regarded as the limit of the groups $H^k(U, U-K)$ as K ranges over compact subsets of U_i. Then there are natural maps $H_c^k(U_i) \to H_c^k(U_{i+1})$ since the second of these groups is a limit over a larger collection of K's. Thus we can form $\varinjlim H_c^k(U_i)$ which is obviously isomorphic to $H_c^k(U)$ since the compact sets in U are just the compact sets in all the U_i's. By Proposition 3.17, $H_{n-k}(U) \approx \varinjlim H_{n-k}(U_i)$. The map D_U is thus the limit of the isomorphisms D_{U_i}, hence is an isomorphism.

Now after all these preliminaries we can prove the theorem in three easy steps:

(1) The case $U = \mathbb{R}^n$ can be proved by regarding $\mathbb{R}^n$ as the interior of Δ^n, and then the map D_U can be identified with the map $H^k(\Delta^n, \partial\Delta^n) \to H_{n-k}(\Delta^n)$ given by cap product with a unit times the generator $[\Delta^n] \in H_n(\Delta^n, \partial\Delta^n)$, where $[\Delta^n]$ is the homology class of the identity map of Δ^n. The only nontrivial value of k is $k = n$, when the cap product map is an isomorphism since a generator of $H^n(\Delta^n, \partial\Delta^n) \approx Hom(H_n(\Delta^n, \partial\Delta^n), R)$ is represented by the cocycle taking the value 1 on $[\Delta^n]$, which means that $[\Delta^n] \frown [\varphi]$ is a generator of $H_0(\Delta^n)$.

(2) More generally, D_U is an isomorphism for U an arbitrary open set in $\mathbb{R}^n$. To see this, first write U as a countable union of convex open sets U_i, for example balls, and let $V_i = \bigcup_{j<i} U_j$. We apply (A) to U_i and V_i. Both V_i and $U_i \cap V_i$ are unions of $i-1$ convex open sets, so by induction on the number of convex open sets in a cover, D_{V_i} and $D_{U_i\cap V_i}$ are isomorphisms. By (1), D_{U_i} is an isomorphism since U_i is homeomorphic to $\mathbb{R}^n$. So $D_{U_i\cup V_i}$ is an isomorphism. Since U is the increasing union of the V_i's and each D_{V_i} is an isomorphism, so is D_U by (B).

(3) For a manifold M expressible as a countable union of open sets U_i homeomorphic to $\mathbb{R}^n$, and in particular for closed M, the proposition now follows by the argument in (2), with the words 'convex open set' replaced by 'open set in $\mathbb{R}^n$.' □

A consequence of Poincaré duality is that $H_n(M;\mathbb{Z}) \approx H^0(M;\mathbb{Z}) \approx \mathbb{Z}$ if M is a connected closed orientable n-manifold. In fact, $H_n(M)$ can be used to detect whether M is closed and orientable:

Proposition 3.20. *If M is a connected n-manifold then $H_i(M;\mathbb{Z}) = 0$ for $i > n$, and $H_n(M;\mathbb{Z}) = 0$ unless M is closed and orientable.*

Proof: Represent an element of $H_i(M;\mathbb{Z})$ by a cycle z. This has compact image in M, so there is an open set $U \subset M$ containing the image of z and having compact closure $\overline{U} \subset M$. Let $V = M - \overline{U}$. Part of the long exact sequence of the triple $(M, U \cup V, V)$ fits into a commutative diagram

$$\begin{array}{ccccc} H_{i+1}(M, U\cup V;\mathbb{Z}) & \longrightarrow & H_i(U\cup V, V;\mathbb{Z}) & \longrightarrow & H_i(M,V;\mathbb{Z}) \\ & & \uparrow\approx & & \uparrow \\ & & H_i(U;\mathbb{Z}) & \longrightarrow & H_i(M;\mathbb{Z}) \end{array}$$

When $i > n$, the two groups on either side of $H_i(U \cup V, V;\mathbb{Z})$ are zero by Lemma 3.16 since $U \cup V$ and V are the complements of compact sets in M. Hence $H_i(U;\mathbb{Z}) = 0$, so z is a boundary in U and therefore in M, and we conclude that $H_i(M;\mathbb{Z}) = 0$.

When $i = n$ note first that since M is connected, the image $[z]_x$ of the class $[z] \in H_n(M;\mathbb{Z})$ in $H_n(M, M-\{x\};\mathbb{Z}) \approx \mathbb{Z}$ is independent of x, up to sign at least. If M is noncompact, $[z]_x$ must therefore be zero for all x since there are points x not in the image of z. By Lemma 3.16, z then represents zero in $H_n(M, V;\mathbb{Z})$, hence also in $H_n(U;\mathbb{Z})$ since the first term in the upper row of the diagram is zero when $i = n$, by Lemma 3.16 again. So $[z] = 0$ in $H_n(M;\mathbb{Z})$ and hence $H_n(M;\mathbb{Z}) = 0$.

Now suppose M is compact and $[z] \neq 0$ in $H_n(M;\mathbb{Z})$. As already noted, the image $[z]_x$ of $[z]$ in $H_n(M, M-\{x\};\mathbb{Z})$ is independent of x, up to sign; say it is k times a generator of $H_n(M, M-\{x\};\mathbb{Z})$ with $k \geq 0$. By Lemma 3.16, k is nonzero since we assume $[z] \neq 0$. Then $k^{-1}[z]_x$ defines an orientation of M. □

Cup Product and Poincaré Duality

There is a close connection between cup and cap product, arising from the formula

$(*)$ $\quad \psi(\alpha \frown \varphi) = (\varphi \smile \psi)(\alpha) \quad$ for $\alpha \in C_{k+\ell}(X;R)$, $\varphi \in C^k(X;R)$, and $\psi \in C^\ell(X;R)$.

This holds since for $\sigma : \Delta^{k+\ell} \to X$,

$$\begin{aligned} \psi(\sigma \frown \varphi) &= \psi(\varphi(\sigma\big|[v_0,\cdots,v_k])\sigma\big|[v_k,\cdots,v_{k+\ell}]) \\ &= \varphi(\sigma\big|[v_0,\cdots,v_k])\psi(\sigma\big|[v_k,\cdots,v_{k+\ell}]) \\ &= (\varphi \smile \psi)(\sigma). \end{aligned}$$

Passing to homology and cohomology, the formula $(*)$ gives a commutative diagram

$$\begin{array}{ccc} H^\ell(X;R) & \xrightarrow{\quad h \quad} & Hom_R(H_\ell(X;R),R) \\ \Big\downarrow{\scriptstyle \varphi\smile} & & \Big\downarrow{\scriptstyle (\frown\varphi)^*} \\ H^{k+\ell}(X;R) & \xrightarrow{\quad h \quad} & Hom_R(H_{k+\ell}(X;R),R) \end{array}$$

When the maps h are isomorphisms, for example when R is a field or when $R=\mathbb{Z}$ and the homology groups of X are free, then the map $\varphi\smile$ is the dual of $\frown\varphi$. Thus in these cases, cup and cap product determine each other, at least if one assume finite generation so that homology is the dual of cohomology, as well as vice versa. However, there are examples where cap and cup products are not equivalent when $R=\mathbb{Z}$ and there is torsion in homology.

By means of the formula $(*)$, Poincaré duality has nontrivial implications for the cup product structure of manifolds, which we now examine. For a closed R-orientable n-manifold M, consider the cup product pairing

$$H^k(M;R)\times H^{n-k}(M;R)\longrightarrow R, \qquad (\varphi,\psi)\mapsto(\varphi\smile\psi)[M]$$

Such a bilinear pairing $A\times B\to R$ is said to be **nonsingular** if both the maps $A\to Hom(B,R)$ and $B\to Hom(A,R)$ obtained by viewing the pairing as a function of each variable separately are isomorphisms.

Proposition 3.21. *The cup product pairing is nonsingular for closed R-orientable manifolds when R is a field, or when $R=\mathbb{Z}$ and torsion in $H^*(M;\mathbb{Z})$ is factored out.*

Proof: Consider the composition

$$H^{n-k}(M;R)\xrightarrow{h} Hom(H_{n-k}(M;R),R)\xrightarrow{D^*} Hom(H^k(M;R),R)$$

where h is the map appearing in the universal coefficient theorem, induced by evaluation of cochains on chains, and D^* is the Hom-dual of the Poincaré duality map $D\colon H^k\to H_{n-k}$. The composition D^*h sends $\psi\in H^{n-k}(M;R)$ to the homomorphism $\varphi\mapsto\psi([M]\frown\varphi)=(\varphi\smile\psi)[M]$. For field coefficients or for integer coefficients with torsion factored out, h is an isomorphism. Taking into account the commutativity formula for cup product, nonsingularity of the pairing is then equivalent to D being an isomorphism. □

Corollary 3.22. *If M is a closed orientable connected n-manifold, then for each element $\alpha\in H^k(M;\mathbb{Z})$ of infinite order which is not a proper multiple of another element, there*

exists an element $\beta \in H^{n-k}(M;\mathbb{Z})$ *such that* $\alpha \smile \beta$ *is a generator of* $H^n(M;\mathbb{Z}) \approx \mathbb{Z}$. *With coefficients in a field, the same conclusion holds for any nonzero* α.

Proof: The hypotheses on α mean that it generates a $\mathbb{Z}$ summand of $H^k(M;\mathbb{Z})$. There is then a homomorphism $\varphi: H^k(M;\mathbb{Z}) \to \mathbb{Z}$ with $\varphi(\alpha) = 1$. By the nonsingularity of the cup product pairing, φ is realized by taking cup product with an element $\beta \in H^{n-k}(M;\mathbb{Z})$ and evaluating on $[M]$, so $\alpha \smile \beta$ generates $H^n(M;\mathbb{Z})$. The case of field coefficients is similar. □

Example: Projective Spaces. The cup product structure $H^*(\mathbb{C}\mathrm{P}^n) \approx \mathbb{Z}[\alpha]/(\alpha^{n+1})$ with $\alpha \in H^2(\mathbb{C}\mathrm{P}^n)$ can easily be deduced from this, as follows. The inclusion $\mathbb{C}\mathrm{P}^{n-1} \subset \mathbb{C}\mathrm{P}^n$ induces an isomorphism on H^i for $i \leq 2n-2$, so by induction on n, $H^{2i}(\mathbb{C}\mathrm{P}^n)$ is generated by α^i for $i < n$. By the corollary, there is an integer m such that $\alpha \smile (m\alpha^{n-1}) = m\alpha^n$ generates $H^{2n}(\mathbb{C}\mathrm{P}^n)$, hence $m = \pm 1$, and therefore $H^*(\mathbb{C}\mathrm{P}^n) \approx \mathbb{Z}[\alpha]/(\alpha^{n+1})$. The same argument shows $H^*(\mathbb{H}\mathrm{P}^n) \approx \mathbb{Z}[\alpha]/(\alpha^{n+1})$ with $\alpha \in H^4$. For $\mathbb{R}\mathrm{P}^n$ one can use the same argument with $\mathbb{Z}_2$ coefficients to deduce that $H^*(\mathbb{R}\mathrm{P}^n;\mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha]/(\alpha^{n+1})$ with α now 1-dimensional. The cup product structure in infinite-dimensional projective spaces follows from the finite-dimensional case, as we saw in the proof of Theorem 3.11.

The cup product structure in $\mathbb{C}\mathrm{P}^n$ is as simple as one can imagine, consistent with Poincaré duality. Could there be a closed manifold whose cohomology is additively isomorphic to that of $\mathbb{C}\mathrm{P}^n$ but with a different cup product structure? For $n = 2$ the answer is no since duality implies that the square of a generator of H^2 must be a generator of H^4. For $n = 3$, duality says that the product of generators of H^2 and H^4 must be a generator of H^6, but nothing is said about the square of a generator of H^2. Indeed, for $S^2 \times S^4$, whose cohomology has the same additive structure as $\mathbb{C}\mathrm{P}^3$, the square of the generator of $H^2(S^2 \times S^4;\mathbb{Z})$ is zero since it is the pullback of a generator of $H^2(S^2;\mathbb{Z})$ under the projection $S^2 \times S^4 \to S^2$ and in $H^*(S^2;\mathbb{Z})$ the square of the generator of H^2 has to be zero. More generally, there are closed 6-manifolds having the same cohomology groups as $\mathbb{C}\mathrm{P}^3$ but where the square of the generator of H^2 is an arbitrary multiple of a generator of H^4. (These manifolds are fiber bundles over S^4 with fiber S^2; see [VBKT].)

Example: Lens Spaces. Cup products in lens spaces can be computed in the same way as in projective spaces. For a lens space L^{2n+1} of dimension $2n+1$ with fundamental group $\mathbb{Z}_m$, we computed $H_i(L^{2n+1})$ in §2.2 to be $\mathbb{Z}$ for $i = 0$ and $2n+1$, $\mathbb{Z}_m$ for odd $i < 2n+1$, and 0 otherwise. In particular, this implies that L^{2n+1} is orientable, which also follows since L^{2n+1} is the orbit space of an action of $\mathbb{Z}_m$ on S^{2n+1} by orientation preserving homeomorphisms; see the exercises at the end of the section. By the universal

coefficient theorem, $H^i(L^{2n+1};\mathbb{Z}_m) \approx \mathbb{Z}_m$ for $i \le 2n+1$. Let $\alpha \in H^1(L^{2n+1};\mathbb{Z}_m)$ and $\beta \in H^2(L^{2n+1};\mathbb{Z}_m)$ be generators. The assertion is:

$$\text{For } j \le 2n+1, \quad H^j(L^{2n+1};\mathbb{Z}_m) \approx \mathbb{Z}_m \quad \text{ is generated by } \begin{cases} \beta^i & \text{for } j = 2i \\ \alpha\beta^i & \text{for } j = 2i+1 \end{cases}$$

When m is prime this follows from the corollary, by induction on n, since the natural inclusion $L^{2n-1} \subset L^{2n+1}$ induces an isomorphism on H^j for $j \le 2n-1$, as one sees by comparing the cellular chain complexes for L^{2n-1} and L^{2n+1}. For nonprime m the corollary does not apply directly, but its proof does since the maps $h : H^i(L^{2n+1};\mathbb{Z}_m) \to Hom(H_i(L^{2n+1};\mathbb{Z}_m),\mathbb{Z}_m)$ are isomorphisms.

The rest of the cup product structure on $H^*(L^{2n+1};\mathbb{Z}_m)$ is determined once α^2 is expressed as a multiple of β. When m is odd, the commutativity formula for cup product implies $\alpha^2 = 0$. When m is even, commutativity implies only that α^2 is either zero or the unique class of order two. In fact, α^2 is nonzero since the 2-skeleton L^2 is the circle L^1 with a 2-cell attached by a map of degree m, and we computed the cup product in this 2-complex in Example 3 at the beginning of §3.2. It does not seem to be possible to deduce the nontriviality of α^2 from Poincaré duality alone, except when $m = 2$.

As with $\mathbb{R}\mathrm{P}^n$, the ring structure in $H^*(L^{2n+1};\mathbb{Z}_m)$ determines that in $H^*(L^{2n+1};\mathbb{Z})$, and likewise for infinite-dimensional lens spaces L^∞, where one has the slightly simpler structure $H^*(L^\infty;\mathbb{Z}) \approx \mathbb{Z}[\alpha]/(m\alpha)$ with α 2-dimensional. The case of L^{2n+1} is obtained from this by setting $\alpha^{n+1} = 0$ and adjoining the extra $\mathbb{Z} \approx H^{2n+1}(L^{2n+1};\mathbb{Z})$.

A different derivation of the cup product structure in lens spaces is given in Example 1 in §3.4e.

Using the ad hoc notation $H^k_{free}(M)$ for $H^k(M)$ modulo its torsion subgroup, the preceding proposition implies that for a closed orientable manifold M of dimension $2n$, the middle-dimensional cup product pairing $H^n_{free}(M) \times H^n_{free}(M) \to \mathbb{Z}$ is a nonsingular bilinear form on $H^n_{free}(M)$. This form is symmetric or skew-symmetric according to whether n is even or odd. The algebra in the skew-symmetric case is rather simple: With a suitable choice of basis, the matrix of a skew-symmetric nonsingular bilinear form over $\mathbb{Z}$ can be put into the standard form consisting of 2×2 blocks $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ along the diagonal and zeros elsewhere. The proof of this is an algebra exercise at the end of the section. In particular, it follows that the rank of $H^n(M^{2n})$ must be even when n is odd.

The symmetric case is much more interesting algebraically. There are only finitely many isomorphism classes of symmetric nonsingular bilinear forms over $\mathbb{Z}$ of a fixed rank, but this 'finitely many' grows rather rapidly, e.g., more than 80 million for rank 32; see [Serre] for an exposition of this beautiful chapter of number theory. It is known that for

each even $n \geq 2$, every symmetric nonsingular form actually occurs as the cup product pairing in some closed manifold M^{2n}. One can even take M^{2n} to be simply-connected and have the bare minimum of homology: $\mathbb{Z}$'s in dimensions 0 and $2n$ and a $\mathbb{Z}^k$ in dimension n. For $n = 2$ this is a difficult theorem of M. Freedman, who also showed that there are at most two nonhomeomorphic simply-connected closed 4-manifolds with the same bilinear form. Namely, there are two manifolds with the same form if the square $\alpha \smile \alpha$ of some $\alpha \in H^2(M^4)$ is an odd multiple of a generator of $H^4(M^4)$, e.g., if $M = \mathbb{C}\mathrm{P}^2$; and otherwise the M^4 is unique, e.g., if $M = S^4$ or $S^2 \times S^2$; see [Freedman-Quinn]. In §4.4c we take the first step in this direction by proving a classical result of J.H.C.Whitehead that the homotopy type of a simply-connected closed 4-manifold is uniquely determined by its cup product structure.

Other Forms of Duality

Generalizing the definition of a manifold, an ***n*-manifold with boundary** is a Hausdorff space M in which each point has an open neighborhood homeomorphic either to $\mathbb{R}^n$ or to the half-space $\mathbb{R}^n_+ = \{(x_1, \cdots, x_n) \in \mathbb{R}^n \mid x_n \geq 0\}$. If $x \in M$ corresponds under such a homeomorphism to a point $(x_1, \cdots, x_n) \in \mathbb{R}^n_+$ with $x_n = 0$, then by excision we have $H_n(M, M - \{x\}) \approx H_n(\mathbb{R}^n_+, \mathbb{R}^n_+ - \{0\}) = 0$, whereas if x corresponds to a point $(x_1, \cdots, x_n) \in \mathbb{R}^n_+$ with $x_n > 0$ or to a point of $\mathbb{R}^n$, then $H_n(M, M - \{x\}) \approx H_n(\mathbb{R}^n, \mathbb{R}^n - \{0\}) \approx \mathbb{Z}$. Thus the points x with $H_n(M, M - \{x\}) = 0$ form a well-defined subspace ∂M, the **boundary** of M, which is an $(n-1)$-dimensional manifold with empty boundary. For example, $\partial \mathbb{R}^n_+ = \mathbb{R}^{n-1}$ and $\partial D^n = S^{n-1}$.

If M is a manifold with boundary ∂M, then a **collar** neighborhood of ∂M is an open neighborhood homeomorphic to $\partial M \times [0, 1)$ by a homeomorphism taking ∂M to $\partial M \times \{0\}$.

Proposition 3.23. *If M is a compact manifold with boundary, then ∂M has a collar neighborhood.*

Proof: Let M' be M with an external collar attached, i.e., M' is the quotient of the disjoint union of M and $\partial M \times [0, 1]$ in which $x \in \partial M$ is identified with $(x, 0) \in \partial M \times [0, 1]$. It will suffice to construct a homeomorphism $h : M \to M'$ since $\partial M'$ clearly has a collar neighborhood.

Since M is compact, so is the closed subspace ∂M. This implies that we can choose a finite number of continuous functions $\varphi_i : \partial M \to [0, 1]$ such that the sets $V_i = \varphi_i^{-1}(0, 1]$ form an open cover of ∂M, and each V_i has closure contained in an open set $U_i \subset M$

homeomorphic to the half-space $\mathbb{R}^n_+$. After replacing each φ_i by $\varphi_i/\sum_j \varphi_j$ we may assume $\sum_i \varphi_i = 1$.

Let $\psi_k = \varphi_1 + \cdots + \varphi_k$ and let $M_k \subset M'$ be the union of M with the points $(x,t) \in \partial M \times [0,1]$ with $t \le \psi_k(x)$. By definition $\psi_0 = 0$ and $M_0 = M$. We construct a homeomorphism $h_k : M_{k-1} \to M_k$ as follows. The homeomorphism $U_i \approx \mathbb{R}^n_+$ gives a collar neighborhood $\partial U_i \times [-1,0]$ of ∂U_i in U_i, with $x \in \partial U_i$ corresponding to $(x,0) \in \partial U_i \times [-1,0]$. Via the external collar $\partial M \times [0,1]$ we then have an embedding $\partial U_i \times [-1,1] \subset M'$. Outside this $\partial U_i \times [-1,1]$, h_k is defined to be the identity, and for $x \in \partial U_k$, h_k is defined to stretch the segment $\{x\} \times [-1, \psi_{k-1}(x)]$ linearly onto $\{x\} \times [-1, \psi_k(x)]$. The composition of all the h_k's gives a homeomorphism $M \approx M'$. □

More generally, collars can be constructed for the boundaries of paracompact manifolds in the same way.

A compact manifold M with boundary is defined to be R-orientable if $M - \partial M$ is R-orientable as a manifold without boundary. If $\partial M \times [0,1) \subset M$ is a collar neighborhood of ∂M, then $H_i(M, \partial M; R)$ is naturally isomorphic to $H_i(M - \partial M, \partial M \times (0, {}^1\!/_2); R)$, so when M is R-orientable, Lemma 3.16 gives a relative fundamental class $[M] \in H_n(M, \partial M; R)$ restricting to a given orientation at each point of $M - \partial M$.

Generalizing Poincaré duality for closed manifolds, we have:

Theorem 3.24. *Suppose M is a compact R-orientable n-manifold whose boundary ∂M is decomposed as the union of two compact $(n-1)$-dimensional manifolds A and B such that $A \cap B = \partial A = \partial B$. Then cap product with a fundamental class $[M] \in H_n(M, \partial M; R)$ defines isomorphisms $D_M : H^k(M, A; R) \to H_{n-k}(M, B; R)$ for all k.*

The possibility that A, B, or $A \cap B$ is empty is not excluded. The cases $A = \varnothing$ and $B = \varnothing$ are sometimes called 'Lefschetz duality.'

Proof: The cap product map $D_M : H^k(M, A; R) \to H_{n-k}(M, B; R)$ is defined since the existence of collar neighborhoods of $A \cap B$ in A and B and ∂M in M implies that A and B are deformation retracts of open neighborhoods U and V in M such that $U \cup V$ deformation retracts onto $A \cup B = \partial M$ and $U \cap V$ deformation retracts onto $A \cap B$.

The case $B = \varnothing$ is proved by applying Theorem 3.18 to $M - \partial M$. Via a collar neighborhood of ∂M we see that $H^k(M, \partial M; R) \approx H^k_c(M - \partial M; R)$, and there are obvious isomorphisms $H_{n-k}(M; R) \approx H_{n-k}(M - \partial M; R)$.

The general case reduces to the case $B = \varnothing$ by applying the five-lemma to the

following diagram:

$$\begin{array}{ccccccccc}
\cdots \longrightarrow & H^k(M,\partial M) & \longrightarrow & H^k(M,A) & \longrightarrow & H^k(\partial M,A) & \longrightarrow & H^{k+1}(M,\partial M) & \longrightarrow \cdots \\
& \downarrow{\scriptstyle [M]\frown} & & \downarrow{\scriptstyle [M]\frown} & & \downarrow{\scriptstyle \approx} & & \downarrow{\scriptstyle [M]\frown} & \\
& & & & & H^k(B,\partial B) & & & \\
& & & & & \downarrow{\scriptstyle [B]\frown} & & & \\
\cdots \longrightarrow & H_{n-k}(M) & \longrightarrow & H_{n-k}(M,B) & \longrightarrow & H_{n-k-1}(B) & \longrightarrow & H_{n-k-1}(M) & \longrightarrow \cdots
\end{array}$$

For commutativity of the middle square one needs to check that the boundary map $H_n(M,\partial M) \to H_{n-1}(\partial M)$ sends a fundamental class for M to a fundamental class for ∂M. We leave this as an exercise at the end of the section. □

Next we turn to **Alexander duality**:

Theorem 3.25. *If K is a compact, locally contractible, nonempty, proper subspace of S^n, then $\widetilde{H}_i(S^n - K;\mathbb{Z}) \approx \widetilde{H}^{n-i-1}(K;\mathbb{Z})$ for all i.*

The special case that K is a sphere or disk was treated by more elementary means in Proposition 2.31. As remarked there, it is interesting that the homology of $S^n - K$ does not depend on the way that K is embedded in S^n. There can be local pathologies as in the case of the Alexander horned sphere, or global complications as with knotted circles in S^3, but these have no effect on the homology of the complement. The only requirement is that K is not too bad a space itself. An example where the theorem fails without the local contractibility assumption is the 'quasi-circle,' defined in an exercise for §1.3. This is a compact subspace $K \subset \mathbb{R}^2$ which we can regard as a subspace of S^2 by adding a point at infinity. Then we have $\widetilde{H}_0(S^2 - K;\mathbb{Z}) \approx \mathbb{Z}$ since $S^2 - K$ has two path-components, but $\widetilde{H}^1(K;\mathbb{Z}) = 0$.

Proof: For $i \neq 0$ the desired isomorphism will be the composition of five isomorphisms

$$\begin{aligned}
H_i(S^n - K) &\approx H_c^{n-i}(S^n - K) \\
&\quad \varinjlim H^{n-i}(S^n - K, U - K) \\
&\quad \varinjlim H^{n-i}(S^n, U) \\
&\quad \varinjlim \widetilde{H}^{n-i-1}(U) \qquad \text{if } i \neq 0 \\
&\quad \widetilde{H}^{n-i-1}(K)
\end{aligned}$$

where the direct limits are taken with respect to open neighborhoods U of K. The first isomorphism is Poincaré duality. The second is the definition of cohomology with compact supports. The third is excision. The fourth comes from the long exact sequences

of the pairs (S^n, U). For the final isomorphism, an easy special case is when K has a neighborhood which is a mapping cylinder of some map $X \to K$, as in the 'letter examples' at the beginning of Chapter 0. For in this special case we can compute the direct limit using neighborhoods U which are segments of the mapping cylinder which deformation retract to K.

To obtain the last isomorphism in the general case we need to quote Theorem A.6 in the Appendix, which says that K is a retract of some neighborhood U_0 in S^n since K is locally contractible. In computing the direct limits we can then restrict attention to open sets $U \subset U_0$, which all retract to K by restricting the retraction of U_0. The natural restriction map $\varinjlim H^*(U) \to H^*(K)$ is surjective since we can pull back each element of $H^*(K)$ to the direct limit via the retractions $U \to K$.

To see injectivity of the map $\varinjlim H^*(U) \to H^*(K)$, note first that each neighborhood $U \subset U_0$ of K contains a neighborhood V such that the inclusion $V \hookrightarrow U$ is homotopic to the retraction $V \to K \subset U$, namely, regarding U as a subspace of an $\mathbb{R}^n \subset S^n$, the linear homotopy $U \times I \to \mathbb{R}^n$ from the identity to the retraction $U \to K$ takes $K \times I$ to K, hence takes $V \times I$ to U for some neighborhood V of K, by compactness of I. Since the inclusion $V \hookrightarrow U$ is homotopic to the retraction $V \to K \subset U$, the restriction $H^*(U) \to H^*(V)$ factors through $H^*(K)$, so if an element of $H^*(U)$ restricts to zero in $H^*(K)$, it restricts to zero in $H^*(V)$. This implies that the map $\varinjlim H^*(U) \to H^*(K)$ is injective.

The only difficulty in the case $i = 0$ is that the fourth of the five isomorphisms above does not hold, and instead we have only a short exact sequence

$$0 \longrightarrow \widetilde{H}^{n-i-1}(U) \longrightarrow H^{n-i}(S^n, U) \longrightarrow \widetilde{H}^{n-i}(S^n) \longrightarrow 0$$

To get around this little problem, observe that all the groups involved in the first three isomorphisms map naturally to the corresponding groups for $K = \varnothing = U$. Then if we take the kernels of these maps we get an isomorphism $\widetilde{H}^0(S^n - K) \approx \varinjlim \widetilde{H}^n(U)$, and we have seen that the latter group is isomorphic to $\widetilde{H}^n(K)$. □

Corollary 3.26. *A closed non-orientable n-manifold cannot be embedded as a subspace of $\mathbb{R}^{n+1}$.*

Proof: Suppose we had an embedding $M \hookrightarrow \mathbb{R}^{n+1}$ with M a closed non-orientable n-manifold, which we may assume to be connected. Then M would embed in S^{n+1}, and Alexander duality would give an isomorphism of $H_{n-1}(M; \mathbb{Z})$ with $H^1(S^{n-1} - M; \mathbb{Z})$. The latter group is torsionfree by the universal coefficient theorem for cohomology. But

$H_{n-1}(M;\mathbb{Z})$ contains torsion by the universal coefficient theorem for homology, since nonorientability imples that $H_n(M;\mathbb{Z}) = 0$ and $H_n(M;\mathbb{Z}_2) \neq 0$, forcing $Tor(H_{n-1}(M;\mathbb{Z}),\mathbb{Z}_2)$ to be nonzero. □

Here is another kind of duality which generalizes the calculation of the local homology groups $H_i(\mathbb{R}^n, \mathbb{R}^n - \{x\})$:

Proposition 3.27. *If K is a compact, locally contractible subspace of an orientable n-manifold M, then there are isomorphisms $H_i(M, M-K;\mathbb{Z}) \approx H^{n-i}(K;\mathbb{Z})$ for all i.*

Proof: Let U be an open neighborhood of K in M and let V be the complement of a compact set in M. We assume $U \cap V = \varnothing$. Then cap product with fundamental classes gives a commutative diagram with exact rows

$$\begin{array}{ccccccccc}
\longrightarrow & H_i(M-K) & \longrightarrow & H_i(M) & \longrightarrow & H_i(M, M-K) & \approx & H_i(U, U-K) & \longrightarrow \\
 & \uparrow & & \uparrow & & & & \uparrow & \\
\longrightarrow & H^{n-i}(M, U\cup V) & \longrightarrow & H^{n-i}(M,V) & \longrightarrow & H^{n-i}(U\cup V, V) & \approx & H^{n-i}(U) & \longrightarrow
\end{array}$$

Passing to the direct limit over decreasing $U \supset K$ and V, the first two vertical arrows become the Poincaré duality isomorphisms $H_i(M-K) \approx H_c^{n-i}(M-K)$ and $H_i(M) \approx H_c^{n-i}(M)$. The five-lemma then gives an isomorphism $H_i(M, M-K) \approx \varinjlim H^{n-i}(U)$. The latter group will be isomorphic to $H^{n-i}(K)$ by the argument in the proof of Theorem 3.25, provided that K is a retract of some neighborhood in M. To obtain such a retraction we can first construct a map $M \hookrightarrow \mathbb{R}^k$ which is an embedding near the compact set K, for some large k, by the method of Corollary A.8 in the Appendix. Then a neighborhood of K in $\mathbb{R}^k$ retracts onto K by Theorem A.6, so the restriction of this retraction to a neighborhood of K in M finishes the job. □

There is a way of extending Alexander duality and the duality in the preceding proposition to compact sets K which are not locally contractible, by replacing the singular cohomology of K by another kind of cohomology called Čech cohomology. This is defined in the following way. To each open cover $\mathcal{U} = \{U_\alpha\}$ of a given space X we can associate a simplicial complex $N(\mathcal{U})$ called the **nerve** of $\mathcal{U}$, having a vertex v_α for each U_α and having a set of $k+1$ vertices spanning a k-simplex whenever the $k+1$ corresponding U_α's have nonempty intersection. When another cover $\mathcal{V} = \{V_\beta\}$ is a refinement of $\mathcal{U}$, so each V_β is contained in some U_α, then these inclusions induce a simplicial map $N(\mathcal{V}) \to N(\mathcal{U})$ which is well-defined up to homotopy. We can then form the direct limit $\varinjlim H^i(N(\mathcal{U});G)$ with respect to finer and finer open covers $\mathcal{U}$. This limit group is by

definition the **Čech cohomology group** $\check{H}^i(X;G)$. For a full exposition of this cohomology theory see [Eilenberg-Steenrod]. With an analogous definition of relative groups, Čech cohomology turns out to satisfy the same axioms as singular cohomology, and indeed a stronger form of excision: a map $(X,A) \to (Y,B)$ which restricts to a homeomorphism $X - A \to Y - B$ induces isomorphisms on Čech cohomology groups. For spaces homotopy equivalent to CW complexes, Čech cohomology coincides with singular cohomology, but for spaces with local complexities it often behaves more reasonably. For example, if X is the subspace of $\mathbb{R}^3$ consisting of the spheres of radius $1/n$ and center $(1/n, 0, 0)$ for $n = 1, 2, \cdots$ then contrary to what one might expect, $H^3(X;\mathbb{Z})$ is nonzero, as shown in [Barratt-Milnor]. But $\check{H}^3(X;\mathbb{Z}) = 0$ and $\check{H}^2(X;\mathbb{Z}) = \mathbb{Z}^\infty$, the direct sum of countably many copies of $\mathbb{Z}$.

Oddly enough, the corresponding Čech homology groups defined using inverse limits are not so well-behaved, because the exactness axiom fails due to the algebraic fact that an inverse limit of exact sequences need not be exact, as a direct limit would be; see §3.4f. However, there is a way around this problem using a more refined definition. This is 'Steenrod homology theory,' which the reader can find out about in [Milnor 95].

Exercises

1. Show that every covering space of an orientable manifold is orientable.

2. Given a covering space action of a group G on an orientable manifold M by orientation-preserving homeomorphisms, show that the quotient manifold M/G is also orientable.

3. Show that $M \times N$ is orientable iff M and N are both orientable.

4. For a map $f : M \to N$ between connected closed orientable n-manifolds with fundamental classes $[M]$ and $[N]$, the **degree** of f is defined to be the integer d such that $f_*([M]) = d[N]$, so the sign of the degree depends on the choice of fundamental classes. Show that for any connected closed orientable n-manifold M there is a degree 1 map $M \to S^n$.

5. For a map $f : M \to N$ between connected closed orientable n-manifolds, suppose there is a ball $B \subset N$ such that $f^{-1}(B)$ is the disjoint union of of balls B_i each mapped homeomorphically by f onto B. Show the degree of f is $\sum_i \varepsilon_i$ where ε_i is $+1$ or -1 according to whether $f : B_i \to B$ preserves or reverses local orientations induced from given fundamental classes $[M]$ and $[N]$.

6. Show that a p-sheeted covering space projection $M \to N$ has degree $\pm p$, when M and N are connected closed orientable manifolds.

7. Show that for a degree 1 map $f:M\to N$ of connected closed orientable manifolds, the induced map $f_*:\pi_1M\to\pi_1N$ is surjective, hence also $f_*:H_1(M)\to H_1(N)$. [Hint: If f is not surjective, factor it through the covering space $\widetilde N\to N$ corresponding to the subgroup $Im(f_*)\subset\pi_1N$, then consider the two cases that this covering is finite-sheeted or infinite-sheeted.]

8. If M_g denotes the closed orientable surface of genus g, show that degree 1 maps $M_g\to M_h$ exist iff $g\ge h$.

9. As an algebraic application of the preceding problem, show that in a free group F with basis $x_1,\cdots,x_{2k}$, the product of commutators $[x_1,x_2]\cdots[x_{2k-1},x_{2k}]$ is not equal to a product of fewer than k commutators $[v_i,w_i]$ of elements $v_i,w_i\in F$. [Recall that the 2-cell of M_k is attached by the product $[x_1,x_2]\cdots[x_{2k-1},x_{2k}]$. From a relation $[x_1,x_2]\cdots[x_{2k-1},x_{2k}]=[v_1,w_1]\cdots[v_j,w_j]$ in F, construct a degree 1 map $M_j\to M_k$.]

10. Let $M_h'\subset M_g$ be a compact subsurface of genus h with one boundary circle, i.e., M_h' is homeomorphism to M_h with an open disk removed. Show there is no retraction $M_g\to M_h'$ if $h>g/2$. [Apply the previous problem, using the fact that M_g-M_h' has genus $g-h$.]

11. Let X be the subspace of $\mathbb{R}^2$ which is the union of the circles of radius $1/n$ and center $(1/n,0)$ for $n=1,2,\cdots$. If $f_n:I\to X$ is the loop based at the origin winding once around the n^{th} circle, show that the infinite product of commutators $[x_1,x_2][x_3,x_4]\cdots$ defines a loop in X which is nontrivial in $H_1(X)$. [Use problem 9.]

12. Verify that for a map $f:X\to Y$, cap products in X and Y are related by the formula $f_*\big(a\frown f^*(b)\big)=f_*(a)\frown b$.

13. Show that $(\alpha\frown\varphi)\frown\psi=\alpha\frown(\varphi\smile\psi)$ for $\alpha\in C_k(X;R)$, $\varphi\in C^\ell(X;R)$, and $\psi\in C^m(X;R)$. Deduce that cap product makes $H_*(X;R)$ into a right $H^*(X;R)$-module.

14. Show that a direct limit of exact sequences is exact. More generally, show that homology commutes with direct limits: If $\{C_\alpha,f_{\alpha\beta}\}$ is a directed system of chain complexes, with the maps $f_{\alpha\beta}:C_\alpha\to C_\beta$ chain maps, then $H_n(\varinjlim C_\alpha)=\varinjlim H_n(C_\alpha)$.

15. Show that $H_c^k(X;G)$ can be defined as the k^{th} cohomology group of the subcomplex of the cochain complex $C^k(X;G)$ consisting of cochains $\varphi:C_k(X)\to G$ for which there exists a compact set $K=K_\varphi\subset X$ such that φ is zero on chains in $X-K$.

16. Using Poincaré duality, show that a closed manifold of odd dimension has Euler characteristic zero.

17. Show that if a closed orientable manifold M of dimension $2k$ has $\widetilde H_i(M)=0$ for $i<k$, then $H_k(M)$ has no torsion.

18. Show that the homology groups of a closed connected orientable 3-manifold M are determined by $\pi_1(M)$. In particular, note that $H_2(M)$ has no torsion.

19. Show that if M is a closed, connected, non-orientable n-manifold then the torsion subgroup of $H_{n-1}(M;\mathbb{Z})$ is $\mathbb{Z}_2$.

20. Show that with respect to a suitable basis, a skew-symmetric nonsingular bilinear form over $\mathbb{Z}$ can be represented by a matrix consisting of 2×2 blocks $\left(\begin{smallmatrix}0&-1\\1&0\end{smallmatrix}\right)$ along the diagonal and zeros elsewhere. [For the matrix of a bilinear form, the following operation can be realized by a change of basis: add an integer multiple of the i^{th} row to the j^{th} row and add the same integer multiple of the i^{th} column to the j^{th} column. Use this to fix up each column in turn. Note that a skew-symmetric matrix must have zeros on the diagonal.]

21. Show that a nonsingular symmetric or skew-symmetric bilinear pairing $F^n\times F^n\to F$, with F a field, cannot be identically zero when restricted to all pairs v, w in a k-dimensional subspace $V\subset F^n$ if $k>n/2$.

22. Show that if the closed orientable surface M_g of genus g retracts onto a graph $X\subset M_g$, then $H_1(X)$ has rank at most g. Deduce an alternative proof of Problem 9 from this, and construct a retraction of M_g onto a wedge of k circles for each $k\le g$.

23. Show that the closed surface M_g of genus g is not homotopy equivalent to a wedge sum $X\vee Y$ of CW complexes with nontrivial reduced homology.

24. Show that if M is a compact R-orientable n-manifold, then the boundary map $H_n(M,\partial M;R)\to H_{n-1}(\partial M;R)$ sends a fundamental class for $(M,\partial M)$ to a fundamental class for ∂M.

25. Show that a compact manifold does not retract onto its boundary.

26. For a compact manifold M verify that the following diagram relating Poincaré duality for M and ∂M is commutative, up to sign at least:

$$\begin{array}{ccccccc}
H^{k-1}(\partial M;R) & \longrightarrow & H^k(M,\partial M;R) & \longrightarrow & H^k(M;R) & \longrightarrow & H^k(\partial M;R)\\
\downarrow{\scriptstyle [\partial M]\frown} & & \downarrow{\scriptstyle [M]\frown} & & \downarrow{\scriptstyle [M]\frown} & & \downarrow{\scriptstyle [\partial M]\frown}\\
H_{n-k}(\partial M;R) & \longrightarrow & H_{n-k}(M;R) & \longrightarrow & H_{n-k}(M,\partial M;R) & \longrightarrow & H_{n-k-1}(\partial M;R)
\end{array}$$

27. For some reasonably broad class of noncompact R-orientable n-manifolds M with boundary, show there are Poincaré duality isomorphisms $H^k_c(M;R)\approx H_{n-k}(M,\partial M;R)$

for all k, using the five-lemma and a diagram

$$\begin{array}{ccccccc}
H_c^{k-1}(\partial M;R) & \longrightarrow & H_c^k(M,\partial M;R) & \longrightarrow & H_c^k(M;R) & \longrightarrow & H_c^k(\partial M;R) \\
\downarrow{\scriptstyle D_{\partial M}} & & \downarrow{\scriptstyle D_M} & & \downarrow{\scriptstyle D_M} & & \downarrow{\scriptstyle D_{\partial M}} \\
H_{n-k}(\partial M;R) & \longrightarrow & H_{n-k}(M;R) & \longrightarrow & H_{n-k}(M,\partial M;R) & \longrightarrow & H_{n-k-1}(\partial M;R)
\end{array}$$

4. Additional Topics

a. Basepoints and Homotopy

The action of π_1 on π_n can be used to describe the difference between $\pi_n(X, x_0)$ and the set of homotopy classes of maps $S^n \to X$ without conditions on basepoints. It will be no extra effort to consider a more general situation, in fact, comparing the set $\langle Z, X\rangle$ of basepoint-preserving homotopy classes of maps $(Z, z_0) \to (X, x_0)$ with the set $[Z, X]$ of unrestricted homotopy classes of maps $Z \to X$, for Z any CW complex with basepoint z_0 a 0-cell.

The first step is to construct an action of $\pi_1(X, x_0)$ on $\langle Z, X\rangle$.

Proposition 4.49. *An action of $\pi_1(X, x_0)$ on $\langle Z, X\rangle$ is defined by setting $[\gamma][f_1] = [f_0]$ whenever there exists a homotopy $f_s : Z \to X$ from f_0 to f_1 such that $f_s(z_0)$ is the loop γ.*

Proof: Given a loop γ in X based at x_0 and a map $f_1 : (Z, z_0) \to (X, x_0)$, then since we assume z_0 is a 0-cell, the homotopy extension property implies that there is a homotopy $f_s : Z \to X$ ending with f_1 and such that $f_s(z_0)$ is the loop γ. Usually one thinks of homotopies as starting rather than ending with a given map, though of course there is no intrinsic difference between the two cases since the ends of the homotopy parameter interval can always be switched. We will see in a moment the rather trivial reason why we need to say 'ending' instead of 'starting' in the present situation.

The existence of the homotopy f_s allows us to define $[\gamma][f_1] = [f_0]$, although the question immediately arises whether this is well-defined. To show that it is, suppose we have homotopies f_s and g_s such that $f_1 \simeq g_1$ fixing z_0 and such that the loops $f_s(z_0)$ and $g_s(z_0)$ are homotopic. These various homotopies define a map $H : Z \times I \times \partial I \cup Z \times \{1\} \times I \cup \{z_0\} \times I \times I \longrightarrow X$ which is f_s on $Z \times I \times \{0\}$, g_s on $Z \times I \times \{1\}$, the homotopy $f_1 \simeq g_1$ on $Z \times \{1\} \times I$, and the homotopy between $f_s(z_0)$ and $g_s(z_0)$ on $\{z_0\} \times I \times I$. We would like to extend H over $Z \times I \times I$. The pair $(I \times I, I \times \partial I \cup \{1\} \times I)$ is homeomorphic to $(I \times I, I \times \{0\})$, and via this homeomorphism we can view H as a map $Z \times I \times \{0\} \cup \{z_0\} \times I \times I \longrightarrow X$, i.e., a map $Z \times I \to X$ with a homotopy on the subcomplex $\{z_0\} \times I$. This means the homotopy extension property can be applied to produce an extension of the original H to $Z \times I \times I$. Restricting this extended H to $Z \times \{0\} \times I$ gives a homotopy $f_0 \simeq g_0$, which shows that $[\gamma][f_1]$ is well-defined.

Clearly $\big([\eta][\gamma]\big)[f_1] = [\eta]\big([\gamma][f_1]\big)$ since we can construct a homotopy of $[f_1]$ in two stages, first using γ and then using η. Note that if we had defined $[\gamma][f_1]$ via a homotopy starting rather than ending with f_1, this formula would be $[\eta]\big([\gamma][f_1]\big) = \big([\gamma][\eta]\big)[f_1]$ since

multiplication of loops reads from left to right. However, we have arranged our definition so that $\big([\eta][\gamma]\big)[f_1] = [\eta]\big([\gamma][f_1]\big)$, hence we have an action of $\pi_1(X, x_0)$ on $\langle Z, X\rangle$. □

When $Z = S^n$ this is the same as the action of $\pi_1(X, x_0)$ on $\pi_n(X, x_0)$ since, in the original definition of γf in terms of I^n, a homotopy from γf to f is obtained by restricting γf to smaller and smaller concentric cubes, and on the 'basepoint' ∂I^n this homotopy traces out the path γ.

Proposition 4.50. *If (Z, z_0) is a CW pair and X is a path-connected space, then the natural map $\langle Z, X\rangle \to [Z, X]$ induces a bijection of the orbit set $\langle Z, X\rangle/\pi_1(X, x_0)$ onto $[Z, X]$.*

In particular, this implies that $[Z, X] = \langle Z, X\rangle$ if X is simply-connected.

Proof: Since X is path-connected, any map $f : Z \to X$ can be homotoped to take z_0 to the basepoint x_0, via homotopy extension, so the map $\langle Z, X\rangle \to [Z, X]$ is onto. If f_0 and f_1 are basepoint-preserving maps which are homotopic via the homotopy $f_s : Z \to X$, then by definition $[f_0] = [\gamma][f_1]$ for the loop $\gamma(s) = f_s(z_0)$, so $[f_0]$ and $[f_1]$ are in the same orbit under the action of $\pi_1(X, x_0)$. □

Example. If X is an H-space with identity element x_0, then the action of $\pi_1(X, x_0)$ on $\langle Z, X\rangle$ is always trivial. This is because, given a map $f : (Z, z_0) \to (X, x_0)$ and a loop γ in X based at x_0, the multiplication in X allows us to define a homotopy $F(z, s) = f(z)\gamma(s)$ which starts and ends with a map which is basepoint-preserving homotopic to f, and such that the loop $F(z_0, s)$ is homotopic to $\gamma(s)$.

One final remark: The set of orbits of the π_1 action on π_n does not generally inherit a group structure from π_n. For example, when $n = 1$ the orbits are just the conjugacy classes in π_1, and these form a group only when π_1 is abelian, i.e., the conjugation action is trivial. Basepoints are thus a necessary technical device for producing the group structure in homotopy groups, though as we have shown, they can be ignored in simply-connected spaces.

Exercises

1. Show directly that if X is a topological group with identity element x_0, then any two maps $f, g : (Z, z_0) \to (X, x_0)$ which are homotopic are homotopic through basepoint-preserving maps.

2. Show that under the map $\langle X, Y\rangle \to Hom(\pi_n(X, x_0), \pi_n(Y, y_0))$, the action of $\pi_1(Y, y_0)$ on $\langle X, Y\rangle$ corresponds to composing with the action on $\pi_n(Y, y_0)$, i.e., $(\gamma f)_* = \beta f_*$.

Deduce a bijection of $[X, K(\pi,1)]$ with the set of orbits of $Hom(\pi_1(X), \pi)$ under composition with inner automorphisms of π. In particular, if π is abelian then $[X, K(\pi,1)] = \langle X, K(\pi,1)\rangle = Hom(\pi_1(X), \pi)$.

3. For a space X let $Aut(X)$ denote the group of homotopy classes of homotopy equivalences $X \to X$. Show that for a CW complex $K(\pi,1)$, $Aut(K(\pi,1))$ is isomorphic to the group of outer automorphisms of π, that is, automorphisms modulo inner automorphisms.

4. With the notation of the preceding problem, show that $Aut(\bigvee_n S^k) \approx \mathrm{GL}_n(\mathbb{Z})$ for $k > 1$, where $\bigvee_n S^k$ denotes the wedge sum of n copies of S^k and $\mathrm{GL}_n(\mathbb{Z})$ is the group of $n \times n$ matrices with entries in $\mathbb{Z}$ having an inverse matrix of the same form. [$\mathrm{GL}_n(\mathbb{Z})$ is the automorphism group of $\mathbb{Z}^n \approx \pi_k(\bigvee_n S^k) \approx H_k(\bigvee_n S^k)$.]

b. The Hopf Invariant

In §2.2 we used homology to distinguish different homotopy classes of maps $S^n \to S^n$ via the notion of degree. Using cup product it is possible to do something similar for maps $S^{2n-1} \to S^n$.

In general, given a map $f: S^m \to S^n$ with $m \geq n$, we can form a CW complex C_f by attaching a cell e^{m+1} to S^n via f. The homotopy type of C_f depends only on the homotopy class of f, by Proposition 0.3. Thus for maps $f, g: S^m \to S^n$, any invariant of homotopy type which distinguishes C_f from C_g will show that f is not homotopic to g. For example, if $m = n$ and f has degree d, then from the cellular chain complex of C_f we see that $H_n(C_f) \approx \mathbb{Z}_{|d|}$, so the homology of C_f detects the degree of f, up to sign. When $m > n$, however, the homology of C_f consists of $\mathbb{Z}$'s in dimensions 0, n, and $m+1$, independent of f. The same is true of cohomology groups, but cup products have a chance of being nontrivial in $H^*(C_f)$ when $m = 2n-1$. In this case, letting $\alpha \in H^n(C_f)$ and $\beta \in H^{2n}(C_f)$ be generators, the multiplicative structure of $H^*(C_f)$ is determined by a relation $\alpha^2 = H(f)\beta$ for an integer $H(f)$ called the **Hopf invariant** of f. The sign of $H(f)$ depends on the choice of generator β, but this can be specified by requiring β to correspond to a fixed generator of $H^{2n}(D^{2n}, \partial D^{2n})$ under the map $H^{2n}(C_f) \approx H^{2n}(C_f, S^n) \to H^{2n}(D^{2n}, \partial D^{2n})$ induced by the characteristic map of the cell e^{2n}, which is determined by f. We can then change the sign of $H(f)$ by composing f with a reflection of S^{2n-1}, of degree -1. If $f \simeq g$, then under the homotopy equivalence $C_f \simeq C_g$ constructed above, the chosen generators β for $H^{2n}(C_f)$ and $H^{2n}(C_g)$ correspond, so $H(f)$ depends only on the homotopy class of f.

If f is a constant map then $C_f = S^n \vee S^{2n}$ and $H(f) = 0$ since C_f retracts onto S^n. Also, $H(f)$ is always zero for odd n since in this case $\alpha \smile \alpha = (-1)^{n^2} \alpha \smile \alpha = -\alpha \smile \alpha$ by the commutativity formula for cup product.

We don't have to look far to find maps of nonzero Hopf invariant. The space $\mathbb{C}\mathrm{P}^2$ is obtained from S^2 by attaching a 4-cell by a map $f: S^3 \to S^2$ with $H(f) = 1$ since $H^*(\mathbb{C}\mathrm{P}^2; \mathbb{Z}) \approx \mathbb{Z}[\alpha]/(\alpha^3)$ by Theorem 3.11. Similarly, $\mathbb{H}\mathrm{P}^2$ gives rise to a map $S^7 \to S^4$ of Hopf invariant 1. There is also an octonionic projective plane $\mathbb{O}\mathrm{P}^2$ built from a map $S^{15} \to S^8$ of Hopf invariant 1, described in §4.3.

It is a fundamental theorem of Adams that a map $f: S^{2n-1} \to S^n$ of Hopf invariant 1 exists only when $n = 2, 4, 8$. This has a number of very interesting consequences, for example:

1. $\mathbb{R}^n$ is a division algebra only for $n = 1,\ 2,\ 4,\ 8$.
2. S^n is an H-space only for $n = 0,\ 1,\ 3,\ 7$.

3. S^n is parallelizable only for $n = 0,\ 1,\ 3,\ 7$.
4. The only fiber bundles $S^p \to S^q \to S^r$ occur when $(p,q,r) = (0,1,1),\ (1,3,2),\ (3,7,4)$, and $(7,15,8)$.

In each case the problem easily reduces to Adams' theorem. A full discussion of all this, and a proof of Adams' theorem, will be given in [VBKT].

Though maps of Hopf invariant 1 are rare, there are maps $S^{2n-1} \to S^n$ of Hopf invariant 2 for all even n. Namely, consider the space $J_2(S^n)$ constructed in §3.2. This has a CW structure with three cells, of dimensions 0, n, and $2n$, so $J_2(S^n)$ has the form C_f for some $f : S^{2n-1} \to S^n$. We showed that if n is even, the square of a generator of $H^n(J_2(S^n);\mathbb{Z})$ is twice a generator of $H^{2n}(J_2(S^n);\mathbb{Z})$, so $H(f) = \pm 2$.

From this example we can get maps of any even Hopf invariant when n is even via the following fact.

Proposition 4.51. *The map $\pi_{2n-1}(S^n) \to \mathbb{Z}$ defined by the Hopf invariant is a homomorphism.*

This implies that $\pi_{2n-1}(S^n)$ contains a $\mathbb{Z}$ direct summand when n is even, since in these cases we have a surjective homomorphism from the abelian group $\pi_{2n-1}(S^n)$ onto an infinite cyclic group.

Proof: Given two maps $f, g : S^{2n-1} \to S^n$, let us compare C_{f+g} with the space $C_{f\vee g}$ obtained from S^n by attaching two $2n$-cells via f and g. There is a natural quotient map $q : C_{f+g} \to C_{f\vee g}$ collapsing the equatorial disk of the $2n$-cell of C_{f+g} to a point. The induced cellular chain map q_* sends e^{2n}_{f+g} to $e^{2n}_f + e^{2n}_g$. In cohomology this implies that $q^*(\beta_f) = q^*(\beta_g) = \beta_{f+g}$ where β_f, β_g, and β_{f+g} are the cohomology classes dual to the $2n$-cells. Letting α_{f+g} and $\alpha_{f\vee g}$ be the cohomology classes corresponding to the n-cells, we have $q^*(\alpha_{f\vee g}) = \alpha_{f+g}$ since q is a homeomorphism on the n-cells. By restricting to the subspaces C_f and C_g of $C_{f\vee g}$ we see that $\alpha^2_{f\vee g} = H(f)\beta_f + H(g)\beta_g$. Therefore $\alpha^2_{f+g} = q^*(\alpha^2_{f\vee g}) = H(f)q^*(\beta_f) + H(g)q^*(\beta_g) = \big(H(f) + H(g)\big)\beta_{f+g}$. □

Exercises

1. Show that the Hopf invariant of a composition $S^{2n-1} \xrightarrow{f} S^{2n-1} \xrightarrow{g} S^n$ is given by $H(gf) = deg(f)H(g)$, and for a composition $S^{2n-1} \xrightarrow{f} S^n \xrightarrow{g} S^n$ the Hopf invariant satisfies $H(gf) = (deg(g))^2 H(f)$.

2. Show that if $S^k \longrightarrow S^m \xrightarrow{p} S^n$ is a fiber bundle, then $m = n + k = 2n - 1$, hence $k = n - 1$, and, if $n > 1$, the Hopf invariant of p is ± 1. [Show that C_p is a manifold and apply Poincaré duality.]

c. Minimal Cell Structures

We can apply Corollary 4.12, the homology version of Whitehead's theorem, to show that a simply-connected CW complex with finitely-generated homology groups is always homotopy equivalent to a CW complex having the minimum number of cells consistent with its homology, namely, one n-cell for each $\mathbb{Z}$ summand of H_n, and for each $\mathbb{Z}_k$ summand of H_n an n-cell and an $(n+1)$-cell.

Proposition 4.52. *Given a simply-connected CW complex Y and a decomposition of $H_n(Y)$ as a direct sum of cyclic groups G_{ni} for each n, then there is a CW complex X and a cellular homotopy equivalence $f: X \to Y$ such that each cell of X is either:*

(a) a 'generator' n-cell e_i^n, which is a cycle in cellular homology mapped by f to a cellular cycle representing a chosen generator of G_{ni}; or

(b) a 'relator' $(n+1)$-cell E_i^{n+1}, with cellular boundary equal to $k_{ni}e_i^n$ where k_{ni} is the order of G_{ni}, if G_{ni} is finite.

In the nonsimply-connected case this result can easily be false, counterexamples being provided by acyclic spaces and the space $Y = S^1 \vee S^n \cup e^{n+1}$ constructed in the Example following Corollary 4.12, which has the same homology as S^1 but which must have cells of dimension greater than 1 in order to have π_n nontrivial.

Proof: We build X inductively over skeleta, starting with X^1 a point since Y is simply-connected. For the inductive step, suppose we have constructed $f: X^n \to Y$ inducing an isomorphism on H_i for $i < n$ and a surjection on H_n. For the mapping cylinder M_f we then have $H_i(M_f, X^n) = 0$ for $i \le n$ and $H_{n+1}(M_f, X^n) \approx \pi_{n+1}(M_f, X^n)$ by the Hurewicz theorem. Part of the exact sequence in homology for (M_f, X^n) is contained in the following diagram:

$$\begin{array}{ccccccc}
H_{n+1}(Y) & & \pi_{n+1}(M_f, X^n) & & & & H_n(Y) \\
\wr\uparrow & & \wr\uparrow & & & & \wr\uparrow \\
H_{n+1}(M_f) & \longrightarrow & H_{n+1}(M_f, X^n) & \longrightarrow & H_n(X^n) & \longrightarrow & H_n(M_f) \longrightarrow 0 \\
\uparrow & & \uparrow & & \| & & \uparrow \\
H_{n+1}(X^{n+1}) & \longrightarrow & H_{n+1}(X^{n+1}, X^n) & \longrightarrow & H_n(X^n) & \longrightarrow & H_n(X^{n+1}) \longrightarrow 0
\end{array}$$

By induction we know the map $H_n(X^n) \to H_n(M_f) \approx H_n(Y)$ exactly, namely, X^n has generator n-cells, which are cellular cycles mapping to the given generators of $H_n(Y)$, and relator n-cells which do not contribute to $H_n(X^n)$. Thus $H_n(X^n)$ is free with basis the generator n-cells, and the kernel of $H_n(X^n) \to H_n(Y)$ is free with basis given by certain multiples of some of the generator n-cells. Choose 'relator' elements $\rho_i \in H_{n+1}(M_f, X^n)$

mapping to this basis for the kernel, and let the 'generator' elements $\gamma_i \in H_{n+1}(M_f, X^n)$ be the images of the chosen generators of $H_{n+1}(M_f) \approx H_{n+1}(Y)$.

Via the Hurewicz isomorphism $H_{n+1}(M_f, X^n) \approx \pi_{n+1}(M_f, X^n)$, the homology classes ρ_i and γ_i are represented by maps $r_i, g_i : (D^{n+1}, S^n) \to (M_f, X^n)$. We form X^{n+1} from X^n by attaching $(n+1)$-cells via the restrictions of the maps r_i and g_i to S^n. The maps r_i and g_i themselves then give an extension of the inclusion $X^n \hookrightarrow M_f$ to a map $X^{n+1} \to M_f$, whose composition with the retraction $M_f \to Y$ is the extended map $f : X^{n+1} \to Y$. This gives us the lower row of the preceding diagram, with commutative squares. By construction, the subgroup of $H_{n+1}(X^{n+1}, X^n)$ generated by the relator $(n+1)$-cells maps injectively to $H_n(X^n)$, with image the kernel of $H_n(X^n) \to H_n(Y)$, so $f_* : H_n(X^{n+1}) \to H_n(Y)$ is an isomorphism. The elements of $H_{n+1}(X^{n+1}, X^n)$ represented by the generator $(n+1)$-cells map to the γ_i's, hence map to zero in $H_n(X^n)$, so by exactness of the second row these generator $(n+1)$-cells are cellular cycles representing elements of $H_{n+1}(X^{n+1})$ mapped by f_* to the given generators of $H_{n+1}(Y)$. In particular, $f_* : H_{n+1}(X^{n+1}) \to H_{n+1}(Y)$ is surjective, and the induction step is done.

Doing this for all n, we produce a CW complex X and a map $f : X \to Y$ with the desired properties. □

Example. Suppose Y is a simply-connected CW complex such that for some $n \geq 2$,

$$\widetilde{H}_i(Y) \text{ is } \begin{cases} \text{finitely generated for } i = n \\ \text{finitely generated and free for } i = n+1 \\ 0 \text{ for } i \neq n, n+1. \end{cases}$$

Then the proposition says that Y is homotopy equivalent to a CW complex X obtained from a wedge of n-spheres by attaching $(n+1)$-cells. The attaching maps of these cells are determined by the cellular boundary map $H_{n+1}(X^{n+1}, X^n) \to H_n(X^n)$ since $\pi_n(X^n) \approx H_n(X^n)$. This implies that the homotopy type of Y is determined by its homology groups. Hence Y is homotopy equivalent to the wedge of a number of Moore spaces $M(\mathbb{Z}_{k_i}, n)$ of the form $S^n \cup e^{n+1}$ together with a number of spheres S^{n+1}.

This result was extended by Whitehead to the case that X has cells in three adjacent dimensions n, $n+1$, and $n+2$; see [Baues] for a full treatment of this. An interesting special case is:

Proposition 4.53. *Let X be a simply-connected space homotopy equivalent to a CW complex, such that the only nontrivial reduced homology groups of X are $H_2(X) \approx \mathbb{Z}^k$ and $H_4(X) \approx \mathbb{Z}$. Then the homotopy type of X is uniquely determined by the cup product ring $H^*(X; \mathbb{Z})$. In particular, this applies to any simply-connected closed 4-manifold.*

Proof: By the previous proposition we may assume X is of the form $X_\varphi = (\bigvee_i S_i^2) \cup_\varphi e^4$ for some $\varphi : S^3 \to \bigvee_i S_i^2$, where the subscript i ranges from 1 to k. As shown in Example 4 following Proposition 4.32, $\pi_3(\bigvee_i S_i^2)$ is free with basis the Hopf maps $\eta_i : S^3 \to S_i^2$ and the Whitehead products $[\iota_i, \iota_j]$, $i < j$, where ι_i is the inclusion $S_i^2 \hookrightarrow \bigvee_i S_i^2$. Since a homotopy of φ does not change the homotopy type of X, we may assume φ is a linear combination $\sum_i a_i \eta_i + \sum_{i<j} a_{ij}[\iota_i, \iota_j]$. We need to see how the coefficients a_i and a_{ij} determine the cup product $H^2(X;\mathbb{Z}) \times H^2(X;\mathbb{Z}) \to H^4(X;\mathbb{Z})$.

This cup product can be represented by a $k \times k$ symmetric matrix (b_{ij}) where the cup product of the cohomology classes dual to the i^{th} and j^{th} 2-cells is b_{ij} times the class dual to the 4-cell. We claim that $b_{ij} = a_{ij}$ for $i < j$ and $b_{ii} = a_i$. If φ is one of the generators η_i or $[\iota_i, \iota_j]$ this is clear, since if $\varphi = \eta_i$ then X is the wedge sum of $\mathbb{C}\mathrm{P}^2$ with $k-1$ 2-spheres, while if $\varphi = [\iota_i, \iota_j]$ then X_φ is the wedge sum of $S_i^2 \times S_j^2$ with $k-2$ 2-spheres. The claim is also true when φ is $-\eta_i$ or $-[\iota_i, \iota_j]$ since changing the sign of φ amounts to composing φ with a reflection of S^3, and this changes the generator of $H^4(X_\varphi;\mathbb{Z})$ to its negative. The general case now follows by induction from the assertion that the matrix (b_{ij}) for $X_{\varphi+\psi}$ is the sum of the corresponding matrices for X_φ and X_ψ. This assertion can be proved as follows. One has a quotient map

$$q : X_{\varphi+\psi} \longrightarrow X_{\varphi,\psi} = (\bigvee_k S_i^2) \cup_\varphi e^4 \cup_\psi e^4 = X_\varphi \cup X_\psi$$

which is a homeomorphism on the 2-skeleton and collapses an equatorial 3-disk in the 4-cell of $X_{\varphi+\psi}$ to a point. The map $q^* : H^4(X_{\varphi,\psi}) \to H^4(X_{\varphi+\psi})$ sends each of the two generators corresponding to the 4-cells of $X_{\varphi,\psi}$ to a generator, and the assertion follows.

Now suppose X_φ and X_ψ have isomorphic cup product rings. This means bases for $H^*(X_\varphi;\mathbb{Z})$ and $H^*(X_\psi;\mathbb{Z})$ can be chosen so that the matrices specifying the cup product $H^2 \times H^2 \to H^4$ with respect to these bases are the same. The preceding proposition says that any choice of basis can be realized as the dual basis to a cell structure on a CW complex homotopy equivalent to the given one. Therefore we may assume the matrices (b_{ij}) for X_φ and X_ψ are the same. By what we have shown above, this means φ and ψ are homotopic, hence X_φ and X_ψ are homotopy equivalent.

For the statement about simply-connected closed 4-manifolds, Corollaries A.7 and A.8 and Proposition A.10 in the Appendix say that such a manifold M has the homotopy type of a CW complex and finitely generated homology groups. Then Poincaré duality and the universal coefficient theorem imply that the only nontrivial homology groups $H_i(M)$ are $\mathbb{Z}$ for $i = 0, 4$ and $\mathbb{Z}^k$ for $i = 2$, for some $k \geq 0$. □

d. Cohomology of Fiber Bundles

While the homotopy groups of a fiber bundle fit into a nice long exact sequence, the situation for homology and cohomology is much more complicated. The Künneth formula shows that there are some subtleties even for a product bundle, and the general case requires the machinery of spectral sequences, described in Volume II. However there is one special situation which is easily analyzed and which suffices for a number of important calculations of cohomology rings for spaces closely related to Lie groups. This is the Leray-Hirsch theorem:

Theorem 4.54. *Let $p: E \to B$ be a fiber bundle with fiber F, and suppose that $H^*(F;R)$ is a free R-module which is finitely generated in each dimension. Suppose further that there exist classes $c_j \in H^{k_j}(E;R)$ whose restrictions $i^*(c_j)$ form a basis for $H^*(F;R)$ in each fiber F, where $i: F \to E$ is the inclusion. Then the map $\Phi: H^*(B;R) \otimes_R H^*(F;R) \to H^*(E;R)$ sending $\sum_{ij} \alpha_i \otimes i^*(c_j)$ to $\sum_{ij} p^*(\alpha_i) \smile c_j$ is an isomorphism.*

In other words, $H^*(E;R)$ is a free $H^*(B;R)$-module with basis $\{c_j\}$, where we view $H^*(E;R)$ as a module over the ring $H^*(B;R)$ by defining scalar multiplication by $\rho\alpha = p^*(\rho) \smile \alpha$ for $\rho \in H^*(B;R)$ and $\alpha \in H^*(E;R)$.

In the case of a product $E = B \times F$ with $H^*(F;R)$ free over R, we can pull back a basis for $H^*(F;R)$ via the projection $E \to F$ to obtain the classes c_j. Thus the Leray-Hirsch theorem generalizes the version of the Künneth formula involving cup products, Theorem 3.9, at least as far as the additive structure and the module structure over $H^*(B;R)$ are concerned. However, the Leray-Hirsch theorem does not assert that the isomorphism $H^*(E;R) \approx H^*(B;R) \otimes_R H^*(F;R)$ is a ring isomorphism, and in fact this need not be true, as we shall see by an example below.

An example of a bundle where the classes c_j do not exist is the Hopf bundle $S^1 \to S^3 \to S^2$, since $H^*(S^3) \not\approx H^*(S^2) \otimes H^*(S^1)$.

Proof: The proof will be similar to the proof of the cup product Künneth formula, Theorem 3.9. We first prove the result for finite-dimensional CW complexes B by induction on their dimension. The case that B is 0-dimensional is trivial. For the induction step, suppose B has dimension n, and let $B' \subset B$ be the subspace obtained by deleting a point x_α from the interior of each n-cell e^n_α of B. Let $E' = p^{-1}(B')$. Then we have a commutative diagram, with coefficients in R understood:

$$
(*) \qquad \begin{array}{ccccccc}
\longrightarrow & H^*(B,B') \otimes_R H^*(F) & \longrightarrow & H^*(B) \otimes_R H^*(F) & \longrightarrow & H^*(B') \otimes_R H^*(F) & \longrightarrow \\
& \downarrow \Phi & & \downarrow \Phi & & \downarrow \Phi & \\
\longrightarrow & H^*(E,E') & \longrightarrow & H^*(E) & \longrightarrow & H^*(E') & \longrightarrow
\end{array}
$$

The map Φ involving relative cohomology groups is defined by exactly the same formula as in the absolute case, using the relative cup product $H^*(E,E')\otimes_R H^*(E)\to H^*(E,E')$. The first row of the diagram is exact since tensoring with a free module preserves exactness. The second row is of course exact also. The commutativity of the diagram follows from the evident naturality of Φ in the case of the two squares shown. For the other square involving coboundary maps, if we start with an element $\beta\otimes i^*(c_j)\in H^*(B')\otimes_R H^*(F)$ and map this horizontally we get $\delta\beta\otimes i^*(c_j)$ which maps vertically to $p^*(\delta\beta)\smile c_j$, whereas if we first map vertically we get $p^*(\beta)\smile c_j$ which maps horizontally to $\delta(p^*(\beta)\smile c_j) = \delta p^*(\beta)\smile c_j = p^*(\delta\beta)\smile c_j$ since $\delta c_j = 0$.

The space B' deformation retracts onto the skeleton B^{n-1}, and the following lemma implies that the inclusion $p^{-1}(B^{n-1})\hookrightarrow E'$ induces an isomorphism on all cohomology groups:

Lemma 4.55. *Given a fiber bundle $p:E\to B$ and a subspace $A\subset B$ such that (B,A) is k-connected, then $(E,p^{-1}(A))$ is also k-connected.*

Proof: For a map $g:(D^i,\partial D^i)\to(E,p^{-1}(A))$ with $i\le k$, there is by hypothesis a homotopy $f_t:(D^i,\partial D^i)\to(B,A)$ of $f_0 = pg$ to a map f_1 with image in A. The homotopy lifting property then gives a homotopy $g_t:(D^i,\partial D^i)\to(E,p^{-1}(A))$ of g to a map with image in $p^{-1}(A)$. □

Hence the theorem for finite-dimensional B will follow by induction on n and the five-lemma once we show that the left-hand Φ in the diagram is an isomorphism.

By the fiber bundle property there are open disk neighborhoods $U_\alpha\subset e^n_\alpha$ of the points x_α, over which the bundle is a product. Let $U=\bigcup_\alpha U_\alpha$ and $U' = U\cap B'$. By excision we have $H^*(B,B')\approx H^*(U,U')$, and the lemma implies that $H^*(E,E')\approx H^*(p^{-1}(U),p^{-1}(U'))$. Thus we are reduced to showing that $\Phi:H^*(U,U')\otimes_R H^*(F)\to H^*(U\times F,U'\times F)$ is an isomorphism. For this we can either appeal to Theorem 3.9 or we can argue again by induction, applying the five-lemma to the diagram $(*)$ with (B,B') replaced by (U,U'), induction implying that the theorem holds for U and U' since they deformation retract onto complexes of dimensions 0 and $n-1$, respectively.

Next there is the case that B is an infinite-dimensional CW complex. Since (B,B^n) is n-connected, the lemma implies that the same is true of $(E,p^{-1}(B^n))$. Hence in the commutative diagram

$$\begin{array}{ccc} H^*(B)\otimes_R H^*(F) & \longrightarrow & H^*(B^n)\otimes_R H^*(F) \\ \downarrow\Phi & & \downarrow\Phi \\ H^*(E) & \longrightarrow & H^*\big(p^{-1}(B^n)\big) \end{array}$$

the horizontal maps are isomorphisms below dimension n, so the fact that the right-hand Φ is an isomorphism, as we showed above, implies that the left-hand Φ is an isomorphism below dimension n. Since n is arbitrary, this gives the theorem for all CW complexes B.

To extend to the case of arbitrary base spaces B we need the notion of a **pullback bundle** which is used quite frequently in bundle theory. Given a fiber bundle $p: E \to B$ and a map $f: A \to B$, let $f^*(E) = \{(a, e) \in A \times E \mid f(a) = p(e)\}$, so there is a commutative diagram

$$\begin{array}{ccc} f^*(E) & \longrightarrow & E \\ \downarrow & & \downarrow p \\ A & \xrightarrow{f} & B \end{array}$$

where the two maps from $f^*(E)$ are $(a, e) \mapsto a$ and $(a, e) \mapsto e$. It is a simple exercise to verify that the projection $f^*(E) \to A$ is a fiber bundle with the same fiber as $E \to B$.

Now for a non-CW base space B, let $f: A \to B$ be a CW approximation. The map $f^*(E) \to E$ then induces an isomorphism on homotopy groups by the five-lemma applied to the long exact sequences of homotopy groups for the two bundles $E \to B$ and $f^*(E) \to A$ with fiber F. Hence $f^*(E) \to E$ is also an isomorphism on cohomology. The classes c_j pull back to classes in $H^*(f^*(E); R)$ which still restrict to a basis in each fiber, and so the naturality of Φ reduces the theorem for $E \to B$ to the case of $f^*(E) \to A$. □

Corollary 4.56.

(a) $H^*(U(n); \mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_1, x_3, \cdots, x_{2n-1}]$, *the exterior algebra on generators* x_i *of odd dimension* i.

(b) $H^*(SU(n); \mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_3, x_5, \cdots, x_{2n-1}]$.

(c) $H^*(Sp(n); \mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_3, x_7, \cdots, x_{4n-1}]$.

This rather simple structure is in marked contrast with the cohomology of $O(n)$ and $SO(n)$ which is considerably more complicated, as shown in §3.4d.

Proof: For (a), assume inductively that the result holds for $U(n-1)$. From the bundle $U(n-1) \to U(n) \to S^{2n-1}$ we see that the pair $(U(n), U(n-1))$ is $(2n-2)$-connected, so $H^i(U(n); \mathbb{Z}) \to H^i(U(n-1); \mathbb{Z})$ is onto for $i \le 2n-3$ and the classes $x_1, \cdots, x_{2n-3} \in H^*(U(n-1); \mathbb{Z})$ are the restrictions of classes $c_i \in H^*(U(n); \mathbb{Z})$. The products of distinct x_i's form a basis for $H^*(U(n-1); \mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_1, \cdots, x_{2n-3}]$, and these products are restrictions of the corresponding products of c_i's, so the Leray-Hirsch theorem applies, yielding $H^*(U(n); \mathbb{Z}) \approx H^*(U(n-1); \mathbb{Z}) \otimes H^*(S^{2n-1}; \mathbb{Z})$. In view of the commutativity property of cup product, this tensor product is the exterior algebra on odd-dimensional generators $x_1, \cdots, x_{2n-1}$.

The same argument works in the case of $Sp(n)$ using the bundle $Sp(n-1) \to Sp(n) \to S^{4n-1}$. In the case of $SU(n)$ one uses the bundle $SU(n-1) \to SU(n) \to S^{2n-1}$. Since $SU(1)$ is the trivial group, the bundle $SU(1) \to SU(2) \to S^3$ shows that $SU(2) = S^3$, so the first generator is x_3. □

It is illuminating to look more closely at how the homology and cohomology of $O(n)$, $U(n)$, and $Sp(n)$ are related to their bundle structures. For $U(n)$ one has the sequence of bundles

$$\begin{array}{ccccccccc} S^1 = U(1) & \hookrightarrow & U(2) & \hookrightarrow & U(3) & \hookrightarrow \cdots \hookrightarrow & U(n-1) & \hookrightarrow & U(n) \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ & & S^3 & & S^5 & \cdots & S^{2n-3} & & S^{2n-1} \end{array}$$

If all these were product bundles, $U(n)$ would be homeomorphic to $S^1 \times S^3 \times \cdots \times S^{2n-1}$. In actuality the bundles are nontrivial, but the homology and cohomology of $U(n)$ are the same as for this product of spheres, including the cup product structure. For $Sp(n)$ the situation is quite similar, with the corresponding product of spheres $S^3 \times S^7 \times \cdots \times S^{4n-1}$. For $O(n)$ the corresponding sequence of bundles is:

$$\begin{array}{ccccccccc} S^0 = O(1) & \hookrightarrow & O(2) & \hookrightarrow & O(3) & \hookrightarrow \cdots \hookrightarrow & O(n-1) & \hookrightarrow & O(n) \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ & & S^1 & & S^2 & \cdots & S^{n-2} & & S^{n-1} \end{array}$$

The calculations in §3.4d show that $H_*(O(n);\mathbb{Z}_2) \approx H_*(S^0 \times S^1 \times \cdots \times S^{n-1};\mathbb{Z}_2)$, but with $\mathbb{Z}$ coefficients this no longer holds. Instead, consider the coarser sequence of bundles

$$\begin{array}{ccccccccc} S^0 = O(1) & \hookrightarrow & O(3) & \hookrightarrow & O(5) & \hookrightarrow \cdots \hookrightarrow & O(2k-1) & \hookrightarrow & O(2k) \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ & & V_2(\mathbb{R}^3) & & V_2(\mathbb{R}^5) & \cdots & V_2(\mathbb{R}^{2k-1}) & & S^{2k-1} \end{array}$$

where the last bundle $O(2k) \to S^{2k-1}$ is omitted if $n = 2k-1$. The calculations show that $O(n)$ has the same integral homology and cohomology groups as if these bundles were products. However, the cup product structure for $O(n)$ is not the same as in this product, even with $\mathbb{Z}_2$ coefficients.

Cohomology of Grassmannians

Here is an important application of the Leray-Hirsch theorem, generalizing the calculation of the cohomology rings of projective spaces:

Theorem 4.57. *If $G_n(\mathbb{C}^\infty)$ is the Grassmann space of n-dimensional vector subspaces of $\mathbb{C}^\infty$, then $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ is a polynomial ring $\mathbb{Z}[c_1,\cdots,c_n]$ on generators c_i of dimension $2i$. Similarly, $H^*(G_n(\mathbb{R}^\infty);\mathbb{Z}_2)$ is a polynomial ring $\mathbb{Z}_2[w_1,\cdots,w_n]$ on generators w_i of dimension i, and $H^*(G_n(\mathbb{H}^\infty);\mathbb{Z})$ is a polynomial ring $\mathbb{Z}[q_1,\cdots,q_n]$ on generators q_i of dimension $4i$.*

The plan of the proof is to apply the Leray-Hirsch theorem to a fiber bundle $F \longrightarrow E \xrightarrow{p} G_n(\mathbb{C}^\infty)$ where E has the same cohomology ring as the product of n copies of $\mathbb{C}\mathrm{P}^\infty$, a polynomial ring $\mathbb{Z}[x_1,\cdots,x_n]$ with each x_i 2-dimensional. The induced map $p^*: H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}) \to H^*(E;\mathbb{Z})$ will be injective, and we will show its image consists of the symmetric polynomials in $\mathbb{Z}[x_1,\cdots,x_n]$, the polynomials invariant under permutations of the variables x_i. It is a classical theorem in algebra that the symmetric polynomials themselves form a polynomial ring $\mathbb{Z}[\sigma_1,\cdots,\sigma_n]$ where σ_i is a certain symmetric polynomial of degree i, namely the sum of all products of i distinct x_j's. This gives the result for $G_n(\mathbb{C}^\infty)$, and the same argument will also apply in the real and quaternionic cases.

Proof: Define an ***n*-flag** in $\mathbb{C}^k$ to be an ordered n-tuple of orthogonal one-dimensional vector subspaces of $\mathbb{C}^k$. This is equivalent to a chain of vector subspaces $V_1 \subset \cdots \subset V_n$ of $\mathbb{C}^k$ where V_i has dimension i. Why either of these objects should be called a flag is not exactly clear, but that is the traditional name. The set of all n-flags in $\mathbb{C}^k$ forms a subspace $F_n(\mathbb{C}^k)$ of the product of n copies of $\mathbb{C}P^{k-1}$. There is a natural fiber bundle

$$F_n(\mathbb{C}^n) \longrightarrow F_n(\mathbb{C}^k) \xrightarrow{p} G_n(\mathbb{C}^k)$$

where p sends an n-tuple of orthogonal lines to the n-plane they span. The local triviality property can be verified just as for the analogous Stiefel bundle $V_n(\mathbb{C}^n) \to V_n(\mathbb{C}^k) \to G_n(\mathbb{C}^k)$ in Example 5 in the fiber bundle subsection of §4.3. The case $k=\infty$ is covered by the same argument, and this case will be the bundle $F \to E \to G_n(\mathbb{C}^\infty)$ alluded to in the paragraph preceding the proof.

The first step in the proof is to show that $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}) \approx \mathbb{Z}[x_1,\cdots,x_n]$ where x_i is the pullback of a generator of $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ under the map $F_n(\mathbb{C}^\infty) \to \mathbb{C}\mathrm{P}^\infty$ projecting an n-flag onto its i^{th} line. This can be seen by considering the fiber bundle

$$\mathbb{C}\mathrm{P}^\infty \longrightarrow F_n(\mathbb{C}^\infty) \xrightarrow{p} F_{n-1}(\mathbb{C}^\infty)$$

where p projects an n-flag onto the $(n-1)$-flag obtained by ignoring its last line. The local triviality property can be verified by the argument in Example 6 in the fiber bundle subsection of §4.3. The Leray-Hirsch theorem applies since the powers of x_n restrict to a

basis for $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ in the fibers $\mathbb{C}\mathrm{P}^\infty$, these fibers being the spaces of lines in vector subspaces $\mathbb{C}^\infty$ of the standard $\mathbb{C}^\infty$. The elements x_i for $i < n$ are the pullbacks via p of similarly-defined elements of $H^*(F_{n-1}(\mathbb{C}^\infty);\mathbb{Z})$. By induction the latter ring is a polynomial ring on these similarly-defined elements. From the Leray-Hirsch theorem we conclude that the products of powers of the x_i's for $1 \le i \le n$ form an additive basis for $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$, hence this ring is the polynomial ring on the x_i's.

There is a corresponding result for $F_n(\mathbb{C}^k)$, that $H^*(F_n(\mathbb{C}^k);\mathbb{Z})$ is free with basis the monomials $x_1^{i_1}\cdots x_n^{i_n}$ with $i_j \le k-j$ for each j. This is proved in exactly the same way, using induction on n and the fiber bundle $\mathbb{C}\mathrm{P}^{k-n} \to F_n(\mathbb{C}^k) \to F_{n-1}(\mathbb{C}^k)$.

After these preliminaries we can start the main argument, using the fiber bundle $F_n(\mathbb{C}^n) \longrightarrow F_n(\mathbb{C}^\infty) \xrightarrow{p} G_n(\mathbb{C}^\infty)$. The preceding calculations show that the Leray-Hirsch theorem applies, so $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$ is a free module over $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ with basis the monomials $x_1^{i_1}\cdots x_n^{i_n}$ with $i_j \le n-j$ for each j. In particular, the homomorphism p^* is injective since 1 is among the basis elements. Thus we are reduced to computing the image of p^*.

To show that the image of p^* is contained in the symmetric polynomials, consider a permutation of the numbers $1,\cdots,n$. This corresponds to a map $\pi : F_n(\mathbb{C}^\infty) \to F_n(\mathbb{C}^\infty)$ permuting the lines in each n-flag, and the induced map π^* is the corresponding permutation of the variables x_i. Since permuting the lines in an n-flag has no effect on the n-plane they span, we have $p\pi = p$, hence $\pi^* p^* = p^*$, which says that polynomials in the image of p^* are invariant under permutations of the variables.

As remarked earlier, the symmetric polynomials in $\mathbb{Z}[x_1,\cdots,x_n]$ form a polynomial ring $\mathbb{Z}[\sigma_1,\cdots,\sigma_n]$ where σ_i has degree i. Showing that p^* maps onto the symmetric polynomials is equivalent to proving the corresponding result for cohomology with $\mathbb{Z}_p$ coefficients for all primes p, since for spaces whose $\mathbb{Z}$ cohomology is finitely generated and free in each dimension, passing to $\mathbb{Z}_p$ coefficients amounts to tensoring with $\mathbb{Z}_p$, and the symmetric polynomials in $\mathbb{Z}_p[x_1,\cdots,x_n]$ are $\mathbb{Z}_p[\sigma_1,\cdots,\sigma_n]$ just as over $\mathbb{Z}$.

To show that the image of $p^* : H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}_p) \to H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}_p)$ is $\mathbb{Z}_p[\sigma_1,\cdots,\sigma_n]$ we will use a sort of counting argument. For a graded $\mathbb{Z}_p$-module $A = \bigoplus A_i$, define its **Poincaré series** to be the formal power series $p_A(t) = \sum_i a_i t^i$ where a_i is the dimension of A_i as a vector space over $\mathbb{Z}_p$ and we assume a_i is finite for all i. The basic formula we need is that $p_{A\otimes B}(t) = p_A(t)\, p_B(t)$, which is immediate from the definition of the graded tensor product.

In the case at hand all nonzero cohomology is in even dimensions, so let us simplify notation by taking A_i to be the $2i$-dimensional cohomology of the space in question. Since the Poincaré series of $\mathbb{Z}_p[x]$ is $\sum_i t^i = (1-t)^{-1}$, the Poincaré series of $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}_p)$

is $(1-t)^{-n}$. For $H^*(F_n(\mathbb{C}^n);\mathbb{Z}_p)$ the Poincaré series is

$$(1+t)(1+t+t^2)\cdots(1+t+\cdots+t^{n-1}) = \prod_{i=1}^n\Big(\frac{1-t^i}{1-t}\Big) = \Big[\prod_{i=1}^n(1-t^i)\Big](1-t)^{-n}$$

From the additive isomorphism $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}_p) \approx H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}_p)\otimes H^*(F_n(\mathbb{C}^n);\mathbb{Z}_p)$ we get that the Poincaré series $p(t)$ of $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}_p)$ satisfies

$$p(t)\Big[\prod_{i=1}^n(1-t^i)\Big](1-t)^{-n} = (1-t)^{-n} \qquad \text{and hence} \qquad p(t) = \prod_{i=1}^n(1-t^i)^{-1}$$

This is exactly the Poincaré series of $\mathbb{Z}[\sigma_1,\cdots,\sigma_n]$ since σ_i has degree i, and it follows that the image of p^* is all the symmetric polynomials.

This finishes the proof for $G_n(\mathbb{C}^\infty)$. The same arguments apply in the other two cases, using $\mathbb{Z}_2$ coefficients throughout in the real case. □

These calculations show that the isomorphism $H^*(E;R) \approx H^*(B;R)\otimes_R H^*(F;R)$ of the Leray-Hirsch theorem is not generally a ring isomorphism, for if it were, then the polynomial ring $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$ would contain a copy of $H^*(F_n(\mathbb{C}^n);\mathbb{Z})$ as a subring, but in the latter ring every positive-dimensional element is nilpotent — some power of the element is zero — while in a polynomial ring there are no nilpotent elements.

The Gysin Sequence

Besides the Leray-Hirsch theorem which deals with fiber bundles which are cohomologically like products, there is another special class of fiber bundles which permit an elementary analysis of their cohomology structure. These are fiber bundles $S^{n-1} \longrightarrow E \xrightarrow{p} B$ satisfying an orientability hypothesis which will be defined in due course, and which always holds if B is simply-connected or if we take cohomology with $\mathbb{Z}_2$ coefficients. For such bundles we will show there is an exact sequence, called the **Gysin sequence**,

$$\cdots \longrightarrow H^{i-n}(B;R) \xrightarrow{\smile e} H^i(B;R) \xrightarrow{p^*} H^i(E;R) \longrightarrow H^{i-n+1}(B;R) \longrightarrow \cdots$$

where e is a certain 'Euler class' in $H^n(B;R)$. The convention is that $H^i = 0$ for $i<0$, so the initial portion of the Gysin sequence gives isomorphisms $p^*: H^i(B;R) \xrightarrow{\approx} H^i(E;R)$ for $i < n-1$, and the more interesting part of the sequence begins

$$0 \longrightarrow H^{n-1}(B;R) \xrightarrow{p^*} H^{n-1}(E;R) \longrightarrow H^0(B;R) \xrightarrow{\smile e} H^n(B;R) \xrightarrow{p^*} H^n(E;R) \longrightarrow$$
$$H^1(B;R) \xrightarrow{\smile e} H^{n+1}(B;R) \longrightarrow \cdots$$

In the case of a product bundle $E = S^{n-1} \times B$ there is a section, a map $s : B \to E$ with $ps = \mathbb{1}$, so the Gysin sequence breaks up into split short exact sequences

$$0 \longrightarrow H^i(B;R) \xrightarrow{p^*} H^i(S^{n-1} \times B;R) \longrightarrow H^{i-n+1}(B;R) \longrightarrow 0$$

which agrees with the Künneth formula $H^*(S^{n-1} \times B;R) \approx H^*(S^{n-1};R) \otimes H^*(B;R)$. The splitting holds whenever the bundle has a section, and this can happen even when the bundle is not a product. For example, consider the bundle $S^{n-1} \to V_2(\mathbb{R}^{n+1}) \xrightarrow{p} S^n$. Points of $V_2(\mathbb{R}^{n+1})$ are pairs (v_1, v_2) of orthogonal unit vectors in $\mathbb{R}^{n+1}$, and $p(v_1, v_2) = v_1$. If we think of v_1 as a point of S^n and v_2 as a unit vector tangent to S^n at v_1, then $V_2(\mathbb{R}^{n+1})$ is exactly the bundle of unit tangent vectors to S^n. A section of this bundle is a field of unit tangent vectors to S^n, and such a vector field exists iff n is odd by Theorem 2.21. The fact that the Gysin sequence splits when there is a section then says that $V_2(\mathbb{R}^{n+1})$ has the same cohomology as the product $S^{n-1} \times S^n$ if n is odd, at least when $n > 1$ so that the base space S^n is simply-connected and the orientability hypothesis is satisfied. When n is even, the calculations at the end of §3.4d show that $H^*(V_2(\mathbb{R}^{n+1});\mathbb{Z})$ consists of $\mathbb{Z}$'s in dimensions 0 and $2n-1$ and a $\mathbb{Z}_2$ in dimension n. The latter group appears in the Gysin sequence as

$$\begin{array}{ccccccc} H^0(S^n) & \xrightarrow{\smile e} & H^n(S^n) & \longrightarrow & H^n(V_2(\mathbb{R}^{n+1})) & \longrightarrow & H^1(S^n) \\ \| & & \| & & \| & & \| \\ \mathbb{Z} & & \mathbb{Z} & & \mathbb{Z}_2 & & 0 \end{array}$$

hence the Euler class e must be twice a generator of $H^n(S^n)$ in the case that n is even. When n is odd it must be zero in order for the Gysin sequence to split.

In fact, whenever a bundle $S^{n-1} \to E \xrightarrow{p} B$ has a section, the Euler class e must be zero from exactness of $H^0(B) \xrightarrow{\smile e} H^n(B) \xrightarrow{p^*} H^n(E)$ since p^* is injective if there is a section. Thus the Euler class can be viewed as an obstruction to the existence of a section — if it is nonzero, there can be no section. This qualitative statement can be made more precise by bringing in the machinery of obstruction theory; see [Milnor-Stasheff] or [VBKT].

Before deriving the Gysin sequence let us look at some examples of how it can be used to compute cup products.

Example 1. Consider a bundle $S^{n-1} \longrightarrow E \xrightarrow{p} B$ with E contractible, for example the bundle $S^1 \to S^\infty \to \mathbb{C}\mathrm{P}^\infty$ or its real or quaternionic analogues. The long exact sequence of homotopy groups for the bundle shows that B is $(n-1)$-connected. Thus if $n > 1$, B is simply-connected and we have a Gysin sequence for cohomology with $\mathbb{Z}$ coefficients,

while for $n = 1$ we take $\mathbb{Z}_2$ coefficients. If $n > 1$ then since E is contractible the Gysin sequence implies that $H^i(B;\mathbb{Z}) = 0$ for $0 < i < n$ and that $\smile e : H^i(B;\mathbb{Z}) \to H^{i+n}(B;\mathbb{Z})$ is an isomorphism for all $i \geq 0$. It follows that $H^*(B;\mathbb{Z})$ is the polynomial ring $\mathbb{Z}[e]$ when $n > 1$. When $n = 1$ the map $p^* : H^{n-1}(B;\mathbb{Z}_2) \to H^{n-1}(E;\mathbb{Z}_2)$ in the Gysin sequence is surjective, so we see that $\smile e : H^i(B;\mathbb{Z}_2) \to H^{i+n}(B;\mathbb{Z}_2)$ is again an isomorphism for all $i \geq 0$, and hence $H^*(B;\mathbb{Z}_2) \approx \mathbb{Z}_2[e]$ when $n = 1$. Thus the Gysin sequence gives a new derivation of the cup product structure in projective spaces. Also, since polynomial rings $\mathbb{Z}[e]$ are realizable as $H^*(X;\mathbb{Z})$ only when e has dimension 2 or 4, as we show in Chapter 4B, we can conclude that there exist bundles $S^{n-1} \to E \to B$ with B contractible only when n is 1, 2, or 4.

Example 2. For the real Grassmann space $G_n = G_n(\mathbb{R}^\infty)$ we have $\pi_1(G_n) \approx \pi_0(O(n)) \approx \mathbb{Z}_2$, so the universal cover of G_n gives a bundle $S^0 \to \widetilde{G}_n \to G_n$. One can view $\widetilde{G}_n$ as the space of oriented n-planes in $\mathbb{R}^\infty$, which is obviously a 2-sheeted covering space of G_n, hence the universal cover since it is path-connected, being the quotient $V_n(\mathbb{R}^\infty)/SO(n)$ of the contracible space $V_n(\mathbb{R}^\infty)$. A portion of the Gysin sequence for the bundle $S^0 \to \widetilde{G}_n \to G_n$ is $H^0(G_n;\mathbb{Z}_2) \xrightarrow{\smile e} H^1(G_n;\mathbb{Z}_2) \longrightarrow H^1(\widetilde{G}_n;\mathbb{Z}_2)$. This last group is zero since $\widetilde{G}_n$ is simply-connected, and $H^1(G_n;\mathbb{Z}_2) \approx \mathbb{Z}_2$ since $H^*(G_n;\mathbb{Z}_2) \approx \mathbb{Z}_2[w_1,\cdots,w_n]$ as we showed earlier in this section, so $e = w_1$ and the map $\smile e : H^*(G_n;\mathbb{Z}_2) \to H^*(G_n;\mathbb{Z}_2)$ is injective. The Gysin sequence then breaks up into short exact sequences $0 \to H^i(G_n;\mathbb{Z}_2) \xrightarrow{\smile e} H^{i+1}(G_n;\mathbb{Z}_2) \to H^{i+1}(\widetilde{G}_n;\mathbb{Z}_2) \to 0$, from which it follows that $H^*(\widetilde{G}_n;\mathbb{Z}_2)$ is the quotient ring $\mathbb{Z}_2[w_1,\cdots,w_n]/(w_1) \approx \mathbb{Z}_2[w_2,\cdots,w_n]$.

Example 3. The complex analogue of the bundle in the preceding example is a bundle $S^1 \to \widetilde{G}_n(\mathbb{C}^\infty) \to G_n(\mathbb{C}^\infty)$ with $\widetilde{G}_n(\mathbb{C}^\infty)$ 2-connected. This can be constructed in the following way. There is a determinant homomorphism $U(n) \to S^1$ with kernel $SU(n)$, the unitary matrices of determinant 1, so S^1 is the coset space $U(n)/SU(n)$ and by restricting the action of $U(n)$ on $V_n(\mathbb{C}^\infty)$ to $SU(n)$ we obtain the second row of the following commutative diagram:

$$\begin{array}{ccccc} U(n) & \longrightarrow & V_n(\mathbb{C}^\infty) & \longrightarrow & G_n(\mathbb{C}^\infty) \\ \downarrow & & \downarrow & & \| \\ S^1 & \longrightarrow & V_n(\mathbb{C}^\infty)/SU(n) & \longrightarrow & G_n(\mathbb{C}^\infty) \end{array}$$

The second row is a fiber bundle by the usual argument of choosing continuously-varying orthonormal bases in n-planes near a given n-plane. One sees that the space $\widetilde{G}_n(\mathbb{C}^\infty) = V_n(\mathbb{C}^\infty)/SU(n)$ is 2-connected by looking at the relevant portion of the diagram of homo-

topy groups associated to these two bundles:

$$\begin{array}{ccccccccc} & & 0 & \longrightarrow & \pi_2(G_n) & \xrightarrow[\approx]{\partial} & \pi_1(U(n))) & \longrightarrow & 0 \\ & & & & \| & & \downarrow\approx & & \\ 0 & \longrightarrow & \pi_2(\widetilde{G}_n) & \longrightarrow & \pi_2(G_n) & \xrightarrow{\partial} & \pi_1(S^1) & \longrightarrow & \pi_1(\widetilde{G}_n) \longrightarrow 0 \end{array}$$

The second vertical map is an isomorphism since S^1 embeds in $U(n)$ as the subgroup $U(1)$. Since the boundary map in the upper row is an isomorphism, so also is the boundary map in the lower row, and then exactness implies that $\widetilde{G}_n$ is 2-connected.

The analysis of the Gysin sequence for the bundle $S^1 \to \widetilde{G}_n(\mathbb{C}^\infty) \to G_n(\mathbb{C}^\infty)$ proceeds just as in Example 2 above. Part of the sequence is $H^0(G_n;\mathbb{Z}) \xrightarrow{\smile e} H^2(G_n;\mathbb{Z}) \longrightarrow H^2(\widetilde{G}_n;\mathbb{Z})$, and this last group is zero since $\widetilde{G}_n$ is 2-connected, so e must be a generator of $H^2(G_n;\mathbb{Z}) \approx \mathbb{Z}$. Since $H^*(G_n;\mathbb{Z})$ is a polynomial algebra $\mathbb{Z}[c_1,\cdots,c_n]$, we must have $e = \pm c_1$, so the map $\smile e : H^*(G_n;\mathbb{Z}) \to H^*(G_n;\mathbb{Z})$ is injective, the Gysin sequence breaks up into short exact sequences, and $H^*(\widetilde{G}_n;\mathbb{Z})$ is the quotient ring $\mathbb{Z}[c_1,\cdots,c_n]/(c_1) \approx \mathbb{Z}[c_2,\cdots,c_n]$.

The spaces $\widetilde{G}_n$ in the last two examples are often denoted $BSO(n)$ and $BSU(n)$, expressing the fact that they fit into bundles $SO(n) \to V_n(\mathbb{R}^\infty) \to BSO(n)$ and $SU(n) \to V_n(\mathbb{C}^\infty) \to BSU(n)$ with contractible total spaces V_n. There is no quaternion analogue of these spaces since for $n = 2$ this would give a space with cohomology ring $\mathbb{Z}[x]$ on an 8-dimensional generator, which is impossible.

The derivation of the Gysin sequence follows a rather roundabout route:

(1) Deduce a relative version of the Leray-Hirsch theorem from the absolute case.
(2) Specialize this to the case of bundles with fiber a disk, yielding a basic result called the Thom isomorphism.
(3) Show this applies to all orientable disk bundles.
(4) Deduce the Gysin sequence by plugging the Thom isomorphism into the long exact sequence of cohomology groups for the pair consisting of a disk bundle and its boundary sphere bundle.

A **fiber bundle pair** consists of a fiber bundle $p : E \to B$ with fiber F, together with a subspace $E' \subset E$ such that $p\big|E'$ is a bundle with fiber a subspace $F' \subset F$, and moreover, near each fiber of E there is a local trivialization of E which restricts to a local trivialization of E'. For example, if $E \to B$ is a bundle with fiber D^n and $E' \subset E$ is the union of the boundary spheres of the fibers, then (E, E') is a fiber bundle pair; local trivializations of E restrict to local trivializations of E' since homeomorphisms of D^n

must preserve boundary, boundary and interior points of D^n being distinguished by the local homology groups $H_n(D^n, D^n - \{x\})$.

Theorem 4.58. *Suppose that* $(F, F') \to (E, E') \xrightarrow{p} B$ *is a fiber bundle pair such that* $H^*(F, F'; R)$ *is a free* R*-module, finitely generated in each dimension. If there exist classes* $c_j \in H^*(E, E'; R)$ *whose restrictions form a basis for* $H^*(F, F'; R)$ *in each fiber, then* $H^*(E, E'; R)$*, as a module over* $H^*(B; R)$*, is free with basis* $\{c_j\}$.

The module structure is defined just as in the absolute case by $\rho\alpha = p^*(\rho) \smile \alpha$, but now we use the relative cup product $H^*(E; R) \times H^*(E, E'; R) \to H^*(E, E'; R)$.

Proof: Construct a bundle $\widehat{E} \to B$ from E by attaching a copy of the mapping cylinder M of $p : E' \to B$ to E by identifying the subspaces $E' \subset E$ and $E' \subset M$. Thus the fibers $\widehat{F}$ of $\widehat{E}$ are obtained from the fibers F by attaching cones CF' on the subspaces $F' \subset F$. We have $H^*(E, E'; R) \approx H^*(\widehat{E}, M; R)$ via excision and the obvious deformation retraction of $\widehat{E} - B$ onto E, where B is included in $\widehat{E}$ as one end of the mapping cylinder M. Also, the five-lemma gives $H^*(\widehat{E}, M; R) \approx H^*(\widehat{E}, B; R)$ since M deformation retracts to B. All these isomorphisms are $H^*(B; R)$-module isomorphisms. Since B is a retract of $\widehat{E}$ via the bundle projection $\widehat{E} \to B$, we have a splitting $H^*(\widehat{E}; R) \approx H^*(\widehat{E}, B; R) \bigoplus H^*(B; R)$ as $H^*(B; R)$-modules. Let $\widehat{c}_j \in H^*(\widehat{E}; R)$ correspond to $c_j \in H^*(E, E'; R) \approx H^*(\widehat{E}, B; R)$ in this splitting. The classes $\widehat{c}_j$ together with 1 restrict to a basis for $H^*(\widehat{F}; R)$ in each fiber $\widehat{F} = F \cup CF'$, so the absolute form of the Leray-Hirsch theorem implies that $H^*(\widehat{E}; R)$ is a free $H^*(B; R)$-module with basis $\{1, \widehat{c}_j\}$. It follows that $\{c_j\}$ is a basis for the free $H^*(B; R)$-module $H^*(E, E'; R)$. □

Now we specialize to the case of a fiber bundle pair $(D^n, S^{n-1}) \longrightarrow (E, E') \xrightarrow{p} B$. An element $c \in H^n(E, E'; R)$ which restricts in each fiber (D^n, S^{n-1}) to a generator of $H^n(D^n, S^{n-1}; R) \approx R$ is called a **Thom class** for the bundle. We are mainly interested in the cases $R = \mathbb{Z}$ and $\mathbb{Z}_2$, but R could be any commutative ring with identity, in which case a 'generator' is an element with a multiplicative inverse, so all elements of R are multiples of the generator. A Thom class with $\mathbb{Z}$ coefficients gives rise to a Thom class with any other coefficent ring R under the homomorphism $H^n(E, E'; \mathbb{Z}) \to H^n(E, E'; R)$ induced by the homomorphism $\mathbb{Z} \to R$ sending 1 to the identity element of R.

Corollary 4.59. *If the bundle* $(D^n, S^{n-1}) \longrightarrow (E, E') \xrightarrow{p} B$ *has a Thom class* $c \in H^n(E, E'; R)$*, then the map* $\Phi : H^i(B; R) \to H^{i+n}(E, E'; R)$*,* $\Phi(\alpha) = p^*(\alpha) \smile c$*, is an isomorphism for all* $i \geq 0$*, and* $H^i(E, E'; R) = 0$ *for* $i < n$. □

The isomorphism Φ is called the **Thom isomorphism**. The corollary can be made

into a statement about absolute cohomology by defining the **Thom space** $T(E)$ to be the quotient E/E'. Each disk fiber D^n of E becomes a sphere S^n in $T(E)$, and all these spheres coming from different fibers are disjoint except for the common basepoint $x_0 = E'/E'$. A Thom class can be regarded as an element of $H^n(T(E), x_0; R) \approx H^n(T(E); R)$ which restricts to a generator of $H^n(S^n; R)$ in each 'fiber' S^n in $T(E)$, and the Thom isomorphism becomes $H^i(B; R) \approx \widetilde{H}^{n+i}(T(E); R)$.

The major remaining step in the derivation of the Gysin sequence is to relate the existence of a Thom class for a disk bundle $D^n \to E \to B$ to orientability of the bundle. First we define orientability for a sphere bundle $S^{n-1} \to E' \to B$. In the proof of Proposition 4.33 we described a procedure for lifting paths γ in B to homotopy equivalences L_γ between the fibers above the endpoints of γ. We did this for fibrations rather than fiber bundles, but the method applies equally well to fiber bundles whose fiber is a CW complex since the homotopy lifting property was used only for the fiber and for the fiber cross I. Then the sphere bundle E' is called **orientable** if L_γ induces the identity map on $H^{n-1}(S^{n-1}; \mathbb{Z})$ for each loop γ in B. For example, the Klein bottle, regarded as a bundle over S^1 with fiber S^1, is non-orientable since as we follow a path looping once around the base circle, the corresponding fiber circles sweep out the full Klein bottle, ending up where they started but with orientation reversed. The same reasoning shows that the torus, viewed as a circle bundle over S^1 in the same way, is orientable. Indeed, any sphere bundle which is a product is orientable.

One could define orientability for a disk bundle $D^n \to E \to B$ by relativizing the previous definition, constructing lifts L_γ which are homotopy equivalences of the fiber pairs (D^n, S^{n-1}), but since $H^n(D^n, S^{n-1}; \mathbb{Z})$ is canonically isomorphic to $H^{n-1}(S^{n-1}; \mathbb{Z})$ via the coboundary map in the long exact sequence of the pair, it is simpler and amounts to the same thing just to define E to be orientable if its boundary sphere subbundle E' is orientable.

Theorem 4.60. *Every disk bundle has a Thom class with $\mathbb{Z}_2$ coefficients, and orientable disk bundles have Thom classes with $\mathbb{Z}$ coefficients.*

Proof: We will show that if the disk bundle $D^n \to E \to B$ is orientable and B is path-connected, then for cohomology with $\mathbb{Z}$ coefficients the following is true:

(∗) The restriction map $H^i(E, E') \to H^i(D^n_b, S^{n-1}_b)$ is an isomorphism for all fibers D^n_b, $b \in B$, and for all $i \le n$.

Similarly, for $\mathbb{Z}_2$ coefficients our proof will show that (∗) holds without any orientability hypothesis. Hence in both cases a generator of $H^n(E, E') \approx H^n(D^n_b, S^{n-1}_b)$ is a Thom class.

For an orientable disk bundle over a path-connected base, choosing an isomorphism $H^n(D^n_b, S^{n-1}_b) \approx \mathbb{Z}$ for one fiber D^n_b determines such isomorphisms for all fibers since the isomorphisms L^*_γ depend only on the endpoints of γ. Having made such a choice, then if $(*)$ is true we have an isomorphism $H^n(E, E') \approx \mathbb{Z}$ which restricts to the chosen isomorphism $H^n(D^n_b, S^{n-1}_b) \approx \mathbb{Z}$ for all fibers. We will use this preferred isomorphism $H^n(E, E') \approx \mathbb{Z}$ in the inductive proof of $(*)$ given below. In the case of $\mathbb{Z}_2$ coefficients, there can be only one isomorphism of a group with $\mathbb{Z}_2$ so no choices are necessary and orientability is irrelevant. We will prove $(*)$ in the $\mathbb{Z}$ coefficient case, leaving it to the reader to replace all $\mathbb{Z}$'s in the proof by $\mathbb{Z}_2$'s to obtain a proof in the $\mathbb{Z}_2$ case.

Suppose first that B is a connected CW complex of finite dimension k. Let $U \subset B$ be the subspace obtained by deleting one point from the interior of each k-cell of B, and let $V \subset B$ be the union of the open k-cells. For a subspace $A \subset B$ let $E_A \to A$ and $E'_A \to A$ be the disk and sphere bundles obtained by taking the subspaces of E and E' projecting to A. Consider the following portion of a Mayer-Vietoris sequence:

$$H^n(E, E') \longrightarrow H^n(E_U, E'_U) \bigoplus H^n(E_V, E'_V) \xrightarrow{\Psi} H^n(E_{U\cap V}, E'_{U\cap V})$$

The first map is injective since the preceding term in the sequence is zero by induction on k, since $U \cap V$ deformation retracts onto a disjoint union of $(k-1)$-spheres and we can apply Lemma 4.55 to replace $E_{U\cap V}$ by the restriction of E over this union of $(k-1)$-spheres. Thus we have an isomorphism $H^n(E, E') \approx Ker\Psi$. Similarly, by Lemma 4.55 and induction each of the terms $H^n(E_U, E'_U)$, $H^n(E_V, E'_V)$, and $H^n(E_{U\cap V}, E'_{U\cap V})$ is a product of $\mathbb{Z}$'s, with one $\mathbb{Z}$ factor for each component of the spaces involved, projection onto the $\mathbb{Z}$ factor being given by restriction to any fiber in the component. Elements of $Ker\Psi$ are pairs $(\alpha, \beta) \in H^n(E_U, E'_U) \bigoplus H^n(E_V, E'_V)$ having the same restriction to $H^n(E_{U\cap V}, E'_{U\cap V})$. Since B is connected, this means that all the $\mathbb{Z}$-coordinates of α and β in the previous direct product decompositions must be equal, since between any two components of U or V one can interpolate a finite sequence of components of U and V alternately, each component in the sequence having nontrivial intersection with its successor. Thus $Ker\Psi$ is a copy of $\mathbb{Z}$, with restriction to a fiber being the isomorphism $H^n(E, E') \approx \mathbb{Z}$.

To finish proving $(*)$ for finite-dimensional B it remains to see that $H^i(E, E') = 0$ for $i < n$, but this follows immediately by looking at an earlier stage of the Mayer-Vietoris sequence, where the two terms adjacent to $H^i(E, E')$ vanish by induction.

Proving $(*)$ for an infinite-dimensional CW complex B reduces to the finite dimensional case as in the Leray-Hirsch theorem since we are only interested in cohomology in a finite range of dimensions, through dimension n. To handle the case of a non-CW base

space B we can reduce to the CW case by pulling back over a CW approximation to B, again as in the Leray-Hirsch theorem, applying the five-lemma to say that the pullback bundle has isomorphic homotopy groups, hence isomorphic absolute and relative cohomology groups. From the definition of the pullback bundle it is immediate that the pullback of an orientable sphere bundle is orientable, so the case of a CW base implies the general case. □

Now we can derive the Gysin sequence for a sphere bundle $S^{n-1} \to E \xrightarrow{p} B$. The mapping cylinder M_p is then a disk bundle $D^n \to M_p \xrightarrow{p} B$ with E as its boundary sphere bundle. Assuming that a Thom class $c \in H^n(M_p, E; R)$ exists, as is the case if E is orientable or if $R = \mathbb{Z}_2$, then the long exact sequence of cohomology groups for the pair (M_p, E) gives the first row of the following commutative diagram, with R coefficients implicit:

$$\begin{array}{ccccccccc}
\cdots \longrightarrow & H^i(M_p, E) & \xrightarrow{j} & H^i(M_p) & \longrightarrow & H^i(E) & \longrightarrow & H^{i+1}(M_p, E) & \longrightarrow \cdots \\
& \approx\uparrow\Phi & & \approx\uparrow p^* & & \| & & \approx\uparrow\Phi & \\
\cdots \longrightarrow & H^{i-n}(B) & \xrightarrow{\smile e} & H^i(B) & \xrightarrow{p^*} & H^i(E) & \longrightarrow & H^{i-n+1}(B) & \longrightarrow \cdots
\end{array}$$

The maps Φ are the Thom isomorphism, and the vertical map p^* is an isomorphism since M_p deformation retracts onto B. The **Euler class** $e \in H^n(B; R)$ is defined to be $(p^*)^{-1}j^*(c)$, c being a Thom class. The square containing the map $\smile e$ commutes since for $\alpha \in H^{i-n}(B)$ we have $j^*\Phi(\alpha) = j^*(p^*(\alpha) \smile c) = j^*p^*(\alpha) \smile j^*(c) = p^*(\alpha) \smile j^*(c)$, which equals $p^*(\alpha \smile e) = p^*(\alpha) \smile p^*(e)$ since $p^*(e) = j^*(c)$. Another way of defining e is as the class corresponding to $c \smile c$ under the Thom isomorphism, since $\Phi(e) = p^*(e) \smile c = j^*(c) \smile c = c \smile c$.

The lower row of the diagram is the Gysin sequence.

Corollary 4.61. *If the sphere bundle $S^{2k-1} \longrightarrow E \xrightarrow{p} B$ is orientable and $H^*(E; R)$ is a polynomial ring $R[x_1, \cdots, x_\ell]$ on even-dimensional generators x_i, then $H^*(B; R) = R[y_1, \cdots, y_\ell, e]$ where e is the Euler class of the bundle and $p^*(y_i) = x_i$ for each i.*

Proof: Consider the three terms $H^i(B; R) \xrightarrow{\smile e} H^{i+2k}(B; R) \longrightarrow H^{i+2k}(E; R)$ of the Gysin sequence. If i is odd, the third term is zero since E has no odd-dimensional cohomology. Hence the map $\smile e$ is surjective, and by induction on dimension this implies that $H^*(B; R)$ is zero in odd dimensions. This means the Gysin sequence reduces to short exact sequences

$$0 \longrightarrow H^{2i}(B; R) \xrightarrow{\smile e} H^{2i+2k}(B; R) \xrightarrow{p^*} H^{2i+2k}(E; R) \longrightarrow 0$$

Since p^* is surjective, we can choose elements $y_j \in H^*(B;R)$ with $p^*(y_j) = x_j$. It remains to check that $H^*(B;R) = R[y_1, \cdots, y_\ell, e]$, which is elementary algebra: Given $\alpha \in H^*(B;R)$, $p^*(\alpha)$ must be a polynomial $\varphi(x_1, \cdots, x_\ell)$, so $\alpha - \varphi(y_1, \cdots, y_n)$ is in the kernel of p^* and exactness gives an equation $\alpha - \varphi(y_1, \cdots, y_n) = \beta \smile e$ for some $\beta \in H^*(B;R)$. Since β has lower dimension than α, we may assume by induction that β is a polynomial in $y_1, \cdots, y_\ell, e$, hence so is α. Thus the natural map $R[y_1, \cdots, y_\ell, e] \to H^*(B;R)$ is surjective. To see that it is injective, suppose there is a polynomial relation $\varphi(y_1, \cdots, y_\ell, e) = 0$. Applying p^*, we see that this polynomial must have the form $\psi(y_1, \cdots, y_\ell, e) \smile e$ for some polynomial ψ. Since $\smile e$ is injective, we then have a polynomial relation $\psi(y_1, \cdots, y_\ell, e) = 0$. But ψ has lower degree than φ, so by induction we deduce that ψ must be the zero polynomial, hence also φ. □

Example 4. Let us apply this to give another proof that $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ is a polynomial ring $\mathbb{Z}[c_1, \cdots, c_n]$ with $|c_i| = 2i$. We use two fiber bundles:

$$S^{2n-1} \longrightarrow E \longrightarrow G_n(\mathbb{C}^\infty) \qquad \text{and} \qquad S^\infty \longrightarrow E \longrightarrow G_{n-1}(\mathbb{C}^\infty)$$

The total space E in both cases is the space of pairs (P, v) where P is an n-plane in $\mathbb{C}^\infty$ and v is a unit vector in P. In the first bundle the map $E \to G_n(\mathbb{C}^\infty)$ is $(P, v) \mapsto P$, with fiber S^{2n-1}, and for the second bundle the map $E \to G_{n-1}(\mathbb{C}^\infty)$ sends (P, v) to the $(n-1)$-plane in P orthogonal to v, with fiber S^∞, all the unit vectors in $\mathbb{C}^\infty$ orthogonal to a given $(n-1)$-plane. Local triviality for the two bundles is verified in the usual way. Since S^∞ is contractible, the map $E \to G_{n-1}(\mathbb{C}^\infty)$ induces isomorphisms on all homotopy groups, hence also on all cohomology groups. By induction on n we then have $H^*(E;\mathbb{Z}) \approx \mathbb{Z}[c_1, \cdots, c_{n-1}]$. The first bundle is orientable since $G_n(\mathbb{C}^\infty)$ is simply-connected, so the corollary gives $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}) \approx \mathbb{Z}[c_1, \cdots, c_n]$ with $c_n = e$.

The same argument works in the quaternionic case. For a version of this argument in the real case see [VBKT].

Before giving our next example, let us observe that the Gysin sequence with a fixed coefficient ring R is valid for any orientable fiber bundle $F \longrightarrow E \xrightarrow{p} B$ whose fiber is a CW complex F with $H^*(F;R) \approx H^*(S^{n-1};R)$. Orientability is defined just as before in terms of induced maps $L_\gamma^* : H^{n-1}(F;R) \to H^{n-1}(F;R)$. No changes are needed in the derivation of the Gysin sequence to get this more general case. (The associated 'disk' bundle is again the mapping cylinder $CF \to M_p \to B$.)

Example 5. We have computed the cohomology of $\widetilde{G}(\mathbb{R}^\infty)$ with $\mathbb{Z}_2$ coefficients, finding it to be a polynomial ring on generators in dimensions 2 through n, and now we compute

the cohomology with $\mathbb{Z}_p$ coefficients for p an odd prime. The answer will again be a polynomial algebra, but this time on even-dimensional generators, depending on the parity of n. Consider first the case that n is odd, say $n = 2k+1$. There are two fiber bundles

$$V_2(\mathbb{R}^{2k+1}) \longrightarrow E \longrightarrow \widetilde{G}_{2k+1}(\mathbb{R}^\infty) \qquad \text{and} \qquad V_2(\mathbb{R}^\infty) \longrightarrow E \longrightarrow \widetilde{G}_{2k-1}(\mathbb{R}^\infty)$$

where E is the space of triples (P, v_1, v_2) with P an oriented $(2k+1)$-plane in $\mathbb{R}^\infty$ and v_1 and v_2 two orthogonal unit vectors in P. The projection map in the first bundle is $(P, v_1, v_2) \mapsto P$, and for the second bundle the projection sends (P, v_1, v_2) to the oriented $(n-1)$-plane in P orthogonal to v_1 and v_2, with the orientation specified by saying for example that v_1, v_2 followed by a positively-oriented basis for the orthogonal $(n-1)$-plane is a positively-oriented basis for P. Both bundles are orientable since their base spaces $\widetilde{G}_n(\mathbb{R}^\infty)$ are simply-connected, from the bundle $SO(n) \to V_n(\mathbb{R}^\infty) \to \widetilde{G}_n(\mathbb{R}^\infty)$.

The fiber $V_2(\mathbb{R}^\infty)$ of the second bundle is contractible, so E has the same cohomology as $\widetilde{G}_{2k-1}(\mathbb{R}^\infty)$. The fiber of the first bundle has the same $\mathbb{Z}_p$-cohomology as S^{4k-1} if p is odd, by the calculation at the end of §3.4d. So if we assume inductively that $H^*(\widetilde{G}_{2k-1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_{k-1}]$ with $|p_i| = 4i$, then Corollary 4.61 above implies that $H^*(\widetilde{G}_{2k+1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_k]$ where $p_k = e$ has dimension $4k$. The induction can start with $\widetilde{G}_1(\mathbb{R}^\infty)$ which is just S^∞ since an oriented line in $\mathbb{R}^\infty$ contains a unique unit vector in the positive direction.

To handle the case of $\widetilde{G}_n(\mathbb{R}^\infty)$ with $n = 2k$ even, we proceed just as in Example 4, considering the bundles

$$S^{2k-1} \longrightarrow E \longrightarrow \widetilde{G}_{2k}(\mathbb{R}^\infty) \qquad \text{and} \qquad S^\infty \longrightarrow E \longrightarrow \widetilde{G}_{2k-1}(\mathbb{R}^\infty)$$

By the preceding case we have $H^*(\widetilde{G}_{2k-1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_{k-1}]$ with $|p_i| = 4i$, so the corollary implies that $H^*(\widetilde{G}_{2k}(\mathbb{R}^\infty); \mathbb{Z}_p)$ is a polynomial ring on these generators and also a generator in dimension $2k$. Summarizing, for p an odd prime we have shown:

$$\begin{aligned} &H^*(\widetilde{G}_{2k+1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_k] \\ &H^*(\widetilde{G}_{2k}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_{k-1}, e] \end{aligned} \qquad |p_i| = 4i,\ |e| = 2k$$

The same result holds also with $\mathbb{Q}$ coefficients. In fact, our proof applies for any coefficient ring in which 2 has a multiplicative inverse, since all that is needed is that $H^*(V_2(\mathbb{R}^{2k+1}); R) \approx H^*(S^{4k-1}; R)$. For a calculation of the cohomology of $\widetilde{G}_n(\mathbb{R}^\infty)$ with $\mathbb{Z}$ coefficients, see [VBKT]. It turns out that all torsion elements have order 2, and modulo this torsion the integral cohomology is again a polynomial ring on the generators p_i and e. Similar results hold also for the cohomology of the unoriented Grassmann space $G_n(\mathbb{R}^\infty)$, but with the generator e replaced by p_k when n is even.

Exercises

1. For the flag space $F_n(\mathbb{C}^n)$ show that $H^*(F_n(\mathbb{C}^n);\mathbb{Z}) \approx \mathbb{Z}[x_1,...,x_n]/(\sigma_1,...,\sigma_n)$ where σ_i is the i^{th} elementary symmetric polynomial.

2. Apply the Leray-Hirsch theorem to the bundle $S^1 \to S^\infty/\mathbb{Z}_p \to \mathbb{C}\mathrm{P}^\infty$ to compute $H^*(K(\mathbb{Z}_p,1);\mathbb{Z}_p)$ from $H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}_p)$.

3. Use the Leray-Hirsch theorem as in Corollary 4.56 to compute $H^*(V_n(\mathbb{C}^k);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_{2k-2n+1}, x_{2k-2n+3}, \cdots, x_{2k-1}]$ and similarly in the quaternionic case.

4. Use the Gysin sequence to compute the cohomology of bundles $S^{n-1} \to E \to S^k$ with $k \neq n$.

5. Use the Gysin sequence to show that for a fiber bundle $S^k \to S^m \xrightarrow{p} S^n$ we must have $k = n-1$ and $m = 2n-1$. Then use the Thom isomorphism to show that the Hopf invariant of p must be ± 1. [Hence $n = 1, 2, 4, 8$ by Adams' theorem.]

6. Show that if M is a manifold of dimension $2n$ for which there exists a fiber bundle $S^1 \to S^{2n+1} \to M$, then M is simply-connected and $H^*(M;\mathbb{Z}) \approx H^*(\mathbb{C}\mathrm{P}^n;\mathbb{Z})$ as rings. Conversely, if M is simply-connected and $H^*(M;\mathbb{Z}) \approx H^*(\mathbb{C}\mathrm{P}^n;\mathbb{Z})$ as rings, show there is a bundle $S^1 \to E \to M$ where $E \simeq S^{2n+1}$. [When $n > 1$ there are examples where M is not homeomorphic to $\mathbb{C}\mathrm{P}^n$.]

7. Show that if a disk bundle $D^n \to E \to B$ has a Thom class with $\mathbb{Z}$ coefficients, then it is orientable.

8. Show that the inclusion $T^n \hookrightarrow U(n)$ of the n-torus of diagonal matrices is homotopic to the map $T^n \to U(1) \hookrightarrow U(n)$ sending an n-tuple of unit complex numbers $(z_1, \cdots, z_n)$ to the 1×1 matrix $(z_1 \cdots z_n)$. Do the same for the diagonal subgroup of $Sp(n)$. [Hint: Diagonal matrices in $U(n)$ are compositions of scalar multiplication in n lines in $\mathbb{C}^n$, and $\mathbb{C}\mathrm{P}^{n-1}$ is connected.]

9. Fill in the details of the following argument which shows that every $n \times n$ matrix A with entries in $\mathbb{H}$ has an eigenvalue in $\mathbb{H}$. (The usual argument over $\mathbb{C}$ involving roots of the characteristic polynomial does not work due to the lack of a good quaternionic determinant function.) For $t \in [0,1]$ and $\lambda \in S^3 \subset \mathbb{H}$, consider the matrix $t\lambda I + (1-t)A$. If A has no eigenvalues, this is invertible for all t. Thus the map $S^3 \to GL_n(\mathbb{H})$, $\lambda \mapsto \lambda I$, is nullhomotopic. But by the preceding problem and exercise 10(b) in §3.4c, this map represents n times a generator of $\pi_3(GL_n(\mathbb{H}))$.

e. Cohomology Theories and Ω-Spectra

In view of the bijection $H^n(X;G) \approx [X, K(G,n)]$ of Theorem 4.17, it is natural to ask what special property of Eilenberg-MacLane spaces makes the sets $[X, K(G,n)]$ for fixed G and varying n into a cohomology theory. As we shall see, the answer turns out to be rather simple: Each space in the sequence $K(G,1), K(G,2), \cdots$ is the loopspace of the next one, up to weak homotopy equivalence at least.

With this example in mind, we define an **Ω-spectrum** to be a sequence of CW complexes $Z_1, Z_2, \cdots$ together with weak homotopy equivalences $Z_n \to \Omega Z_{n+1}$. Our first goal in this section is to show that for CW complexes X, the sets $h^n(X) = \langle X, Z_n \rangle$ are abelian groups which form a reduced cohomology theory. Using basepoint-preserving homotopy classes $\langle -, Z_n \rangle$ rather than free homotopy classes $[-, Z_n]$ turns out to have many technical advantages, though for $Z_n = K(G,n)$ there is no real distinction. After showing that Ω-spectra give rise to cohomology theories in this way, we shall prove the converse statement, that every cohomology theory on the category of CW pairs comes from an Ω-spectrum. Moreover, the Ω-spectrum is uniquely determined, up to homotopy equivalence, by the cohomology theory. Thus cohomology theories are equivalent to Ω-spectra.

The main task in associating a cohomology theory to an Ω-spectrum is to construct the long exact sequence for a CW pair (X, A). As a first step we build the following diagram:

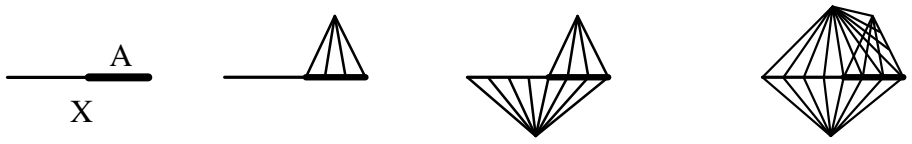

$$
\begin{array}{lccccccc}
 & A & \hookrightarrow & X & \hookrightarrow & X \cup CA & \hookrightarrow & (X \cup CA) \cup CX \hookrightarrow \big((X \cup CA) \cup CX\big) \cup C(X \cup CA) \\
(1) & \| & & \| & & \simeq\downarrow & & \simeq\downarrow\uparrow \qquad\qquad \simeq\downarrow\uparrow \\
 & A & \hookrightarrow & X & \longrightarrow & X/A & \longrightarrow & SA \xrightarrow{\quad\subset\quad} SX
\end{array}
$$

The first row is obtained from the inclusion $A \hookrightarrow X$ by iterating the rule, 'attach a cone on the preceding subspace,' as shown in the pictures above this row. The three downward arrows are quotient maps collapsing the most recently attached cone to a point. Since cones are contractible, these downward maps are homotopy equivalences. The second and third of them have homotopy inverses the evident inclusion maps, indicated by the upward arrows. In the lower row of the diagram the maps are the obvious ones, except for the map $X/A \to SX$ which is the composition of a homotopy inverse of the quotient map $X \cup CA \to X/A$ followed by the maps $X \cup CA \to (X \cup CA) \cup CX \to SA$. Thus the square containing this map commutes up to homotopy. It is easy to check that the same is true

of the right-hand square as well.

The whole construction can now be repeated with $SA \hookrightarrow SX$ in place of $A \hookrightarrow X$, then with double suspensions, etc. The resulting infinite sequence can be written in either of the following two forms:

$$A \to X \to X \cup CA \to SA \to SX \to S(X \cup CA) \to S^2A \to S^2X \to \cdots$$

$$A \to X \to X/A \to SA \to SX \to SX/SA \to S^2A \to S^2X \to \cdots$$

In the first version we use the obvious equality $SX \cup CSA = S(X \cup CA)$. The first version has the advantage that the map $X \cup CA \to SA$ is easily described and canonical, whereas in the second version the corresponding map $X/A \to SA$ is only defined up to homotopy since it depends on choosing a homotopy inverse to the quotient map $X \cup CA \to X/A$. The second version does have the advantage of conciseness, however.

When basepoints are important it is generally more convenient to use reduced cones and reduced suspensions, obtained from ordinary cones and suspensions by collapsing the segment $\{x_0\} \times I$ where x_0 is the basepoint. The image point of this segment in the reduced cone or suspension then serves as a natural basepoint in the quotient. Assuming x_0 is a 0-cell, these collapses of $\{x_0\} \times I$ are homotopy equivalences. Using reduced cones and suspensions in the preceding construction yields a sequence

$$A \hookrightarrow X \to X/A \to \Sigma A \hookrightarrow \Sigma X \to \Sigma(X/A) \to \Sigma^2 A \hookrightarrow \Sigma^2 X \to \cdots \tag{2}$$

where we identity $\Sigma X/\Sigma A$ with $\Sigma(X/A)$ and all the later maps in the sequence are suspensions of the first three maps. This sequence, or its unreduced version, is called the **cofibration sequence** or **Puppe sequence** of the pair (X, A). It has an evident naturality property, namely, a map $(X, A) \to (Y, B)$ induces a map between the cofibration sequences of these two pairs, with homotopy-commutative squares:

$$\begin{array}{ccccccccccc}
A & \longrightarrow & X & \longrightarrow & X/A & \longrightarrow & \Sigma A & \longrightarrow & \Sigma X & \longrightarrow & \Sigma(X/A) & \longrightarrow & \Sigma^2 A & \longrightarrow & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
B & \longrightarrow & Y & \longrightarrow & Y/B & \longrightarrow & \Sigma B & \longrightarrow & \Sigma Y & \longrightarrow & \Sigma(Y/B) & \longrightarrow & \Sigma^2 B & \longrightarrow & \cdots
\end{array}$$

Now let Z be a fixed space with basepoint. Taking basepoint-preserving homotopy classes of maps from the spaces in (2) to Z gives a sequence

$$\langle A, Z\rangle \leftarrow \langle X, Z\rangle \leftarrow \langle X/A, Z\rangle \leftarrow \langle \Sigma A, Z\rangle \leftarrow \langle \Sigma X, Z\rangle \leftarrow \cdots \tag{3}$$

whose maps are defined by composition with those in (2), e.g., $\langle X, Z\rangle \to \langle A, Z\rangle$ sends a map $X \to Z$ to $A \to X \to Z$. For any space Y the set $\langle \Sigma Y, Z\rangle$ has a natural group

structure, generalizing the group structure on homotopy groups, defined by $[f]+[g]=[(f\vee g)\circ c]$ where $c:\Sigma Y\to\Sigma Y\vee\Sigma Y$ collapses the 'equator' $Y\times\{\frac{1}{2}\}$ to the basepoint. The proof that π_2 is abelian immediately generalizes to show that $\langle\Sigma^2Y,Z\rangle$ is abelian, hence also the groups $\langle\Sigma^nY,Z\rangle$ for $n\geq 2$. Thus the sets in (3) are groups starting with $\langle\Sigma A,Z\rangle$, and abelian groups from $\langle\Sigma^2A,Z\rangle$ onward. It is easy to see that the maps between these groups are homomorphisms since the maps in (2) are suspensions from $\Sigma A\to\Sigma X$ onward. In general the first three terms of (3) are only sets with distinguished 'zero' elements, the constant maps.

Proposition 4.62. *The sequence* (3) *is exact.*

Proof: The diagram (1) shows that, up to homotopy equivalence, each term is obtained from its two predecessors by the same procedure of forming a mapping cone, so it suffices to show that $\langle A,Z\rangle\leftarrow\langle X,Z\rangle\leftarrow\langle X\cup CA,Z\rangle$ is exact. A map $f:X\to Z$ goes to zero in $\langle A,Z\rangle$ iff its restriction to A is nullhomotopic, fixing basepoint. This is equivalent to f extending to a map $X\cup CA\to Z$. □

The first clue that loopspaces are relevant is the natural bijection $\langle\Sigma Y,Z\rangle\approx\langle Y,\Omega Z\rangle$, generalizing the isomorphism $\pi_n(Z)\approx\pi_{n-1}(\Omega Z)$ of Proposition 4.36 which is the case $Y=S^{n-1}$. The proof in the general case is the same as for spheres, using the fact that maps $Y\times I\to Z$ are the same as maps $Y\to Z^I$.

Next, suppose we have a weak homotopy equivalence $Z\to\Omega Z'$ for some space Z'. By Proposition 4.21 the induced map $\langle Y,Z\rangle\to\langle Y,\Omega Z'\rangle$ is a bijection for all CW complexes Y, not just spheres. The sequence (3) can then be continued three steps to the left:

$$\begin{array}{ccccccccc} & & & & \langle A,Z\rangle & \longleftarrow & \langle X,Z\rangle & \longleftarrow & \langle X/A,Z\rangle\longleftarrow\cdots \\ & & & & & & \downarrow\approx & & \\ \langle A,Z'\rangle & \longleftarrow & \langle X,Z'\rangle & \longleftarrow & \langle X/A,Z'\rangle & \longleftarrow & \langle A,\Omega Z'\rangle=\langle\Sigma A,Z'\rangle & & \end{array}$$

Thus if we have a sequence of spaces Z_n and weak homotopy equivalences $Z_n\to\Omega Z_{n+1}$, we can extend the sequence (3) to the left indefinitely, producing a long exact sequence

$$(4)\qquad \cdots\leftarrow\langle A,Z_n\rangle\leftarrow\langle X,Z_n\rangle\leftarrow\langle X/A,Z_n\rangle\leftarrow\langle A,Z_{n-1}\rangle\leftarrow\langle X,Z_{n-1}\rangle\leftarrow\cdots$$

All the terms here are abelian groups and the maps homomorphisms since $\langle Y,Z_n\rangle=\langle Y,\Omega^2Z_{n+2}\rangle=\langle\Sigma^2Y,Z_{n+2}\rangle$.

Since the sequence (4) looks formally like a long exact sequence of cohomology groups, we are led to ask whether setting $h^n(X)=\langle X,Z_n\rangle$ gives a cohomology theory for CW complexes X, or more precisely a reduced cohomology theory since $\langle X,Z_n\rangle$ is trivial when

X is a point. Note that $h^n(X)$ is defined even for negative n if we take $Z_n = \Omega Z_{n+1}$ for $n < 0$.

Before showing that we do indeed have a cohomology theory, let us observe that the preceding ideas allow us to recast the definition of a reduced cohomology theory in a slightly more concise form: A reduced cohomology theory on the category $\mathcal{C}$ whose objects are CW complexes with a chosen basepoint 0-cell and whose morphisms are basepoint-preserving maps, is a sequence of functors h^n, $n \in \mathbb{Z}$, from $\mathcal{C}$ to abelian groups, together with natural isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$, defined for all $X \in \mathcal{C}$, such that the following axioms hold for each h^n:

(*i*) If $f \simeq g : X \to Y$ through basepoint-preserving maps, then $f^* = g^* : h^n(Y) \to h^n(X)$.

(*ii*) For each inclusion $A \hookrightarrow X$ in $\mathcal{C}$ the sequence $h^n(X/A) \to h^n(X) \to h^n(A)$ is exact.

(*iii*) For a wedge sum $X = \bigvee_\alpha X_\alpha$ with inclusions $i_\alpha : X_\alpha \hookrightarrow X$, the product map $\prod_\alpha i_\alpha^* : h^n(X) \to \prod_\alpha h^n(X_\alpha)$ is an isomorphism.

To see that these axioms suffice to define a cohomology theory, the main thing to note is that the cofibration sequence (2) allows us to construct the long exact sequence of a pair, just as we did in the case of the functors $h^n(X) = \langle X, Z_n \rangle$. In the converse direction, if we have natural long exact sequences of pairs, then by applying these to pairs of the form (CX, X) we get natural isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$. Note that these natural isomorphisms coming from coboundary maps of pairs (CX, X) uniquely determine the coboundary maps for all pairs (X, A), via the diagram

$$\begin{array}{ccccc} h^n(A) & & & & \\ \downarrow & \searrow & & \searrow & \\ h^{n+1}(CA/A) & \xleftarrow{\approx} & h^{n+1}(CX/A) & \longrightarrow & h^{n+1}(X/A) \end{array}$$

where the vertical and diagonal maps are coboundary maps of pairs, and the diagram commutes by naturality of these coboundary maps.

Proposition 4.63. *For an Ω-spectrum $\{Z_n\}$, the functors $X \mapsto h^n(X) = \langle X, Z_n \rangle$ define a reduced cohomology theory on the category of basepointed CW complexes and basepoint-preserving maps.*

Proof: The isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$ are the equivalences $\langle Y, Z_n \rangle \approx \langle Y, \Omega Z_{n+1} \rangle \approx \langle \Sigma Y, Z_{n+1} \rangle$, which are certainly natural, so it remains to check the axioms (*i*)-(*iii*). The first of these is trivial: A basepoint-preserving map $f : X \to Y$ induces a homomorphism $f^* : h^n(Y) \to h^n(X)$ by composition, $(Y \to Z_n) \mapsto (X \to Y \to Z_n)$, and f^* depends only on the basepoint-preserving homotopy class of f. The second axiom is Proposition

4.62. And the wedge axiom holds since in the realm of basepoint-preserving maps, a map $\bigvee_\alpha X_\alpha \to Z_n$ is the same as a collection of maps $X_\alpha \to Z_n$. □

There is no essential difference between cohomology theories on basepointed CW complexes and cohomology theories on nonbasepointed CW complexes. Given a reduced basepointed cohomology theory $\widetilde{h}^*$, one gets an unreduced theory by setting $h^n(X,A) = \widetilde{h}^n(X/A)$, where $X/\varnothing = X_+$, the union of X with a disjoint basepoint, and this is a nonbasepointed theory since an arbitrary map $X \to Y$ induces a basepoint-preserving map $X_+ \to Y_+$. Furthermore, a nonbasepointed unreduced theory h^* gives a nonbasepointed reduced theory by setting $\widetilde{h}^n(X) = Coker\big(h^n(pt) \to h^n(X)\big)$, where the map is induced by the constant map $X \to pt$. One could also give an argument using suspension, which is always an isomorphism for reduced theories, and which takes one from the nonbasepointed to the basepointed category.

One can use these general constructions to give another proof that $H^n(X;G) \approx \langle X, K(G,n)\rangle$ for CW complexes X, in the same spirit as our proof of the equivalence of singular and simplicial homology of Δ-complexes in §2.1. One checks first that for a fundamental class $\iota \in H^n(K(G,n);G)$, the function $\chi : \langle X, K(G,n)\rangle \to H^n(X;G)$, $\chi(f) = f^*(\iota)$, is a natural transformation between the cohomology theories $h^n(X,A) = \langle X/A, K(G,n)\rangle$ and $H^n(X,A;G)$, i.e., the maps χ fit into a commutative diagram

$$\begin{array}{ccccccccc}
\cdots \longrightarrow & h^n(X,A) & \longrightarrow & h^n(X) & \longrightarrow & h^n(A) & \longrightarrow & h^{n+1}(X,A) & \longrightarrow \cdots \\
& \downarrow\chi & & \downarrow\chi & & \downarrow\chi & & \downarrow\chi & \\
\cdots \longrightarrow & H^n(X,A;G) & \longrightarrow & H^n(X;G) & \longrightarrow & H^n(A;G) & \longrightarrow & H^{n+1}(X,A;G) & \longrightarrow \cdots
\end{array}$$

for each CW pair (X,A). Then one takes $(X,A) = (X^k, X^{k-1})$ and argues by induction that χ is an isomorphism for finite-dimensional X, using the five-lemma and the easily-checked fact that χ is an isomorphism for spheres, hence for wedges of spheres. For infinite-dimensional X, restriction to X^k induces isomorphisms on H^n and h^n if $k > n$, so this case follows from the finite-dimensional case. Filling in the details of this argument is left as an exercise.

Now we turn to showing that the preceding construction of cohomology theories from Ω-spectra produces all cohomology theories on CW complexes. The bulk of the work is in realizing a single functor h^n of a cohomology theory as $\langle -, Z_n\rangle$ for some space Z_n. For this, it will not be necessary to assume Z_n is a loopspace. So let us consider what properties the functor $h(X) = \langle X, Z\rangle$ has, where Z is a fixed space with basepoint. First of all, it is a contravariant functor from the category of basepointed CW complexes to the

category of pointed sets, i.e., sets with a distinguished element, the homotopy class of the constant map in the present case. Morphisms in the category of pointed sets are maps preserving the distinguished element. We have already seen that $h(X)$ satisfies the three axioms (i)-(iii). A further property is the following **Mayer-Vietoris axiom**:

> Suppose the CW complex X is the union of subcomplexes A and B containing the basepoint. Then if $a \in h(A)$ and $b \in h(B)$ restrict to the same element of $h(A \cap B)$, there exists an element $x \in h(X)$ whose restrictions to A and B are the given elements a and b.

Here and in what follows we use the term 'restriction' to mean the map induced by inclusion. In the case that $h(X) = \langle X, Z \rangle$, this axiom is an immediate consequence of the homotopy extension property. The functors h^n in any cohomology theory also satisfy this axiom since there are Mayer-Vietoris exact sequences in any cohomology theory, as we observed in §2.3 in the analogous setting of homology theories.

Now we can state the **Brown Representability Theorem**:

Theorem 4.64. *If h is a contravariant functor from the category of connected basepointed CW complexes to the category of pointed sets, satisfying the homotopy axiom (i), the Mayer-Vietoris axiom, and the wedge axiom (iii), then there exists a connected CW complex K and an element $u \in h(K)$ such that the transformation $T_u : \langle X, K \rangle \to h(X)$, $T_u(f) = f^*(u)$, is a bijection for all X.*

Such a pair (K, u) is called **universal** for the functor h. It is automatic from the definition that the space K in a universal pair (K, u) is unique up to homotopy equivalence. For if (K', u') is also universal for h, then there are maps $f : (K, u) \to (K', u')$ and $g : (K', u') \to (K, u)$ which are unique up to homotopy, and likewise the compositions $gf : (K, u) \to (K, u)$ and $fg : (K', u') \to (K', u')$ are unique up to homotopy, hence are homotopic to the identity maps.

If we apply the theorem with $h(X) = H^n(X; G)$ we obtain a connected CW complex K with bijections $\langle X, K \rangle \approx H^n(X; G)$ for all connected CW complexes X. In particular, taking $X = S^i$, we get bijections $\pi_i(K) \approx H^n(S^i; G)$ for all $i > 0$, so the only nontrivial homotopy group of K is π_n. And in fact K is a $K(G, n)$ since the bijection $\pi_n(K) \approx H^n(S^n; G)$ is a group isomorphism by Lemma 4.68 later in this section.

Before starting the proof of the theorem we make two preliminary comments on the axioms.

(1) The wedge axiom implies that $h(pt)$ is trivial — just use the fact that $pt \vee pt = pt$, so

the map $h(pt) \times h(pt) \to h(pt)$ induced by inclusion of the first summand is a bijection, but this map is the projection $(a, b) \mapsto a$, hence $h(pt)$ must have only one element.

(2) Axioms (i), (iii), and the Mayer-Vietoris axiom imply the exactness axiom (ii). Namely, (ii) is equivalent to exactness of $h(A) \leftarrow H(X) \leftarrow h(X \cup CA)$, which in the context of pointed sets just means $Kernel = Image$. The inclusion $Image \subset Kernel$ holds since the composition $A \to X \cup CA$ is nullhomotopic, so the induced map factors through $h(pt) = 0$. To obtain the opposite inclusion, let us decompose $X \cup CA$ into two subspaces Y and Z by cutting along a copy of A halfway up the cone $CA \subset X \cup CA$, so Y is a smaller copy of CA and Z is the mapping cylinder of the inclusion $A \hookrightarrow X$. Given an element $x \in h(X)$, this extends to an element $z \in h(Z)$ since Z deformation retracts to X. If x restricts to the trivial element of $h(A)$, then z restricts to the trivial element of $h(Y \cap Z)$. The latter element extends to the trivial element of $h(Y)$, so the Mayer-Vietoris axiom implies there is an element of $h(X \cup CA)$ restricting to z in $h(Z)$ and hence to x in $h(X)$.

The bulk of the proof of the theorem will be contained in two lemmas. To state the first of these, consider pairs (K, u) with K a basepointed connected CW complex and $u \in h(K)$. Call such a pair $\boldsymbol{n}$**-universal** if the map $T_u : \pi_i(K) \to h(S^i)$, $T_u(f) = f^*(u)$, is surjective for $i \leq n$ and has trivial kernel for $i < n$. Call (K, u) π**-universal** if it is n-universal for all n.

Lemma 4.65. *Given any pair (Z, z), there exists a π-univeral pair (K, u) with Z a subcomplex of K and $u\big|Z = z$.*

Proof: We construct K from Z by an inductive process of attaching cells. To begin, let $K_1 = Z \bigvee_\alpha S^1_\alpha$ where α ranges over the elements of $h(S^1)$. By the wedge axiom there exists $u_1 \in h(K_1)$ restricting to z on Z and α on S^1_α, hence (K_1, u_1) is 1-universal.

For the inductive step, suppose we have constructed $(K_n, u_n) \supset (Z, z)$ with u_n n-universal. Represent each element α in the kernel of $T_{u_n} : \pi_n(K_n) \to h(S^n)$ by a map $f_\alpha : S^n \to K_n$. Let $f = \bigvee_\alpha f_\alpha : \bigvee_\alpha S^n_\alpha \to K_n$. The reduced mapping cylinder M_f deformation retracts to K_n, so we can regard u_n as an element of $h(M_f)$, and this element restricts to the trivial element of $h(\bigvee_\alpha S^n_\alpha)$ by the definition of f. The exactness property of h then implies that for the reduced mapping cone $C_f = M_f / \bigvee_\alpha S^n_\alpha$ there is an element $w \in h(C_f)$ restricting to u_n on K_n. Note that C_f is obtained from K_n by attaching cells e^{n+1}_α by the maps f_α. To finish the construction of K_{n+1}, set $K_{n+1} = C_f \bigvee_\beta S^{n+1}_\beta$ where β ranges over $h(S^{n+1})$. By the wedge axiom again, there exists $u_{n+1} \in h(K_{n+1})$ restricting to w on C_f and β on S^{n+1}_β.

To see that (K_{n+1}, u_{n+1}) is $(n+1)$-universal, consider the commutative diagram

$$\begin{array}{ccc} \pi_i(K_n) & \longrightarrow & \pi_i(K_{n+1}) \\ & {\scriptstyle T_{u_n}}\searrow \quad \swarrow{\scriptstyle T_{u_{n+1}}} & \\ & h(S^i) & \end{array}$$

Since K_{n+1} is obtained from K_n by attaching $(n+1)$-cells, the upper map is an isomorphism for $i < n$ and a surjection for $i = n$. By induction the map T_{u_n} has the same properties, hence so does $T_{u_{n+1}}$. The kernel of $T_{u_{n+1}}$ is trivial for $i = n$ since an element of this kernel pulls back to $Ker(T_n) \subset \pi_n(K_n)$, by surjectivity of the upper map when $i = n$, and we attached cells to K_n by maps representing all elements of $Ker(T_n)$. And lastly, $T_{u_{n+1}}$ is surjective for $i = n+1$ by construction.

Now let $K = \bigcup_n K_n$. We apply a mapping telescope argument as in the proof of Theorem 3.50 to show there is an element $u \in h(K)$ restricting to u_n on K_n, for all n. The mapping telescope of the inclusions $K_1 \hookrightarrow K_2 \hookrightarrow \cdots$ is the subcomplex $T = \bigcup_i K_i \times [i, i+1]$ of $K \times [1, \infty)$. We take '$\times$' to be the reduced product here, with *basepoint* $\times$ *interval* collapsed to a point. The natural projection $T \to K$ is a homotopy equivalence since $K \times [1, \infty)$ deformation retracts onto T, as we showed in Theorem 3.50. Let $A \subset T$ be the union of the subcomplexes $K_i \times [i, i+1]$ for i odd and let B be the corresponding union for i even. Thus $A \cup B = T$, $A \cap B = \bigvee_i K_i$, $A \simeq \bigvee_i K_{2i-1}$, and $B \simeq \bigvee_i K_{2i}$. By the wedge axiom there exist $a \in h(A)$ and $b \in h(B)$ restricting to u_i on each K_i. Then using the fact that $u_{i+1} \big| K_i = u_i$, the Mayer-Vietoris axiom implies that a and b are the restrictions of an element $t \in h(T)$. Under the isomorphism $h(T) \approx h(K)$, t corresponds to an element $u \in h(K)$ restricting to u_n on K_n for all n.

To verify that (K, u) is π-universal, consider the commutative diagram

$$\begin{array}{ccc} \pi_i(K_n) & \longrightarrow & \pi_i(K) \\ & {\scriptstyle T_{u_n}}\searrow \quad \swarrow{\scriptstyle T_u} & \\ & h(S^i) & \end{array}$$

For $n >> i$ the upper map and T_{u_n} are isomorphisms, hence also T_u. □

Lemma 4.66. *Suppose we are given: (i) a π-universal pair (K, u), (ii) a basepointed CW pair (X, A) with an element $x \in h(X)$, and (iii) a map $f : A \to K$ with $f^*(u) = x \big| A$. Then there exists $g : X \to K$ extending f with $g^*(u) = x$.*

Schematically, this is saying that the diagonal arrow in the following diagram always

exists, where the map i is inclusion:

$$\begin{array}{ccc} (A,a) & \xrightarrow{f} & (K,u) \\ {\scriptstyle i}\downarrow & \nearrow_{g} & \\ (X,x) & & \end{array}$$

Proof: Replacing K by the reduced mapping cylinder of f reduces us to the case that f is the inclusion of a subcomplex. Let Z be the union of X and K with the two copies of A identified. By the Mayer-Vietoris axiom, there exists $z \in h(Z)$ with $z\big|X = x$ and $z\big|K = u$. By the previous lemma, we can embed (Z,z) in a π-universal pair (K',u'). The inclusion $(K,u) \hookrightarrow (K',u')$ induces an isomorphism on homotopy groups since both u and u' are π-universal, so K' deformation retracts onto K. This deformation retraction induces a homotopy rel A of the inclusion $X \hookrightarrow K'$ to a map $g : X \to K$. The relation $g^*(u) = x$ holds since $u'\big|K = u$ and $u'\big|X = x$. □

Proof of the Theorem: It suffices to show that a π-universal pair (K,u) is universal. Applying the preceding lemma with A a point shows that $T_u : \langle X,K\rangle \to h(X)$ is surjective. To show injectivity, suppose $T_u(f_0) = T_u(f_1)$. We apply the preceding lemma with $(X \times I, X \times \partial I)$ playing the role of (X,A), using the maps f_0 and f_1 on $X \times \partial I$ and taking 'x" to be $p^*f_0^*(u) = p^*f_1^*(u)$ where p is the projection $X \times I \to X$. (Here $X \times I$ should be the reduced product, with *basepoint* $\times I$ collapsed to a point.) The lemma then gives a homotopy from f_0 to f_1. □

Theorem 4.67. *Every reduced cohomology theory on the category of basepointed CW complexes and basepoint-preserving maps has the form $h^n(X) = \langle X, K_n\rangle$ for some Ω-spectrum $\{K_n\}$, which is unique up to homotopy equivalence.*

Proof: Since suspension is an isomorphism in any reduced cohomology theory, and the suspension of any CW complex is connected, it suffices to restrict attention to connected CW complexes. Each functor h^n satisfies the homotopy, wedge, and Mayer-Vietoris axioms, as we noted earlier, so the Brown representability theorem gives CW complexes K_n with $h^n(X) = \langle X, K_n\rangle$. It remains to show that the natural isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$ correspond to weak homotopy equivalences $K_n \to \Omega K_{n+1}$. The bijection $h^n(X) = \langle X, K_n\rangle$ corresponds to a bijection $\langle X, K_n\rangle \approx \langle \Sigma X, K_{n+1}\rangle = \langle X, \Omega K_{n+1}\rangle$ which we call Φ. The naturality of this bijection gives, for any map $f : X \to K_n$, a

commutative diagram

$$\begin{array}{ccc} \langle K_n, K_n\rangle & \xrightarrow{f^*} & \langle X, K_n\rangle \\ \downarrow\Phi & & \downarrow\Phi \\ \langle K_n, \Omega K_{n+1}\rangle & \xrightarrow{f^*} & \langle X, \Omega K_{n+1}\rangle \end{array}$$

Let $\varepsilon_n = \Phi(\mathbb{1}) : K_n \to \Omega K_{n+1}$. Then $\Phi(f) = \Phi f^*(\mathbb{1}) = f^*\Phi(\mathbb{1}) = f^*(\varepsilon_n) = \varepsilon_n f$. Thus $\Phi : \langle X, K_n\rangle \to \langle X, \Omega K_{n+1}\rangle$ is composition with ε_n. Since Φ is a bijection, if we take X to be S^i, we see that induces an isomorphism on π_i for all i, so ε_n is a weak homotopy equivalence and we have an Ω-spectrum.

There is one final thing to verify, that the bijection $h^n(X) = \langle X, K_n\rangle$ is a group isomorphism, where $\langle X, K_n\rangle$ has the group structure which comes from identifying it with $\langle X, \Omega K_{n+1}\rangle = \langle \Sigma X, K_{n+1}\rangle$. Via the natural isomorphism $h^n(X) \approx h^{n+1}(\Sigma X)$ this is equivalent to showing the bijection $h^{n+1}(\Sigma X) = \langle \Sigma X, K_{n+1}\rangle$ is a homomorphism. This follows from:

Lemma 4.68. *If a contravariant functor h satisfying the homotopy and wedge axioms takes values in the category of groups and homomorphisms, then the transformation $T_u : \langle \Sigma X, K\rangle \to h(\Sigma X)$ is a homomorphism, for all pairs (K, u).*

Proof: Given maps $f, g : \Sigma X \to K$, we wish to show that $T_u(f+g) = T_u(f) + T_u(g)$, or what is the same thing, $(f+g)^*(u) = f^*(u) + g^*(u)$ in $h(\Sigma X)$. Thus we need to show that $(f+g)^* = f^* + g^* : h(K) \to h(\Sigma X)$.

The map $f+g$ is the composition $\Sigma X \xrightarrow{c} \Sigma X \vee \Sigma X \xrightarrow{f\vee g} K$. This induces a diagram

$$\begin{array}{ccccc} h(\Sigma X) & \xleftarrow{c^*} & h(\Sigma X \vee \Sigma X) & \xleftarrow{(f\vee g)^*} & h(K) \\ & & {\scriptstyle i_1^*\oplus i_2^*}\downarrow\approx & & \\ & & h(\Sigma X) \oplus h(\Sigma X) & & \end{array}$$

where i_1 and i_2 are the inclusions of the two wedge summands. Let $q_1, q_2 : \Sigma X \vee \Sigma X \longrightarrow \Sigma X$ be the quotient maps restricting to the identity on the summand indicated by the subscript and collapsing the other summand to a point. Then $h(\Sigma X \vee \Sigma X) = Im(q_1^*) \oplus Im(q_2^*)$ since $q_j i_k$ is the identity map for $j = k$ and the constant map for $j \neq k$. Hence the subgroups $Im(q_1^*)$ and $Im(q_2^*)$ commute and we see that $q_1^* + q_2^*$ is the inverse of the isomorphism $i_1^* \oplus i_2^*$.

Since $(f \vee g)i_1 = f$ and $(f \vee g)i_2 = g$, an element u in the right-hand group $h(K)$ in the diagram maps to the element $(f^*(u), g^*(u))$ in the lower group $h(\Sigma X) \oplus h(\Sigma X)$. Similarly, since $cq_1 \simeq \mathbb{1}$ and $cq_2 \simeq \mathbb{1}$, an element (u, v) in the lower group in the diagram maps to the element $u + v$ in the left-hand group $h(\Sigma X)$. Putting these facts together,

we conclude that $u \in h(K)$ maps to $f^*(u) + g^*(u)$. On the other hand, the composition across the top of the diagram is induced by $f + g$, so it sends u to $(f + g)^*(u)$. Thus we obtain the desired relation $(f + g)^*(u) = f^*(u) + g^*(u)$. □

A space Z_i in an Ω-spectrum can be regarded as an **infinite loopspace** since there are weak homotopy equivalences $Z_i \simeq_w \Omega Z_{i+1} \simeq_w \Omega^2 Z_{i+2} \simeq_w \cdots$. A number of important spaces in algebraic topology turn out to be infinite loopspaces. Besides Eilenberg-MacLane spaces, two other examples are the infinite-dimensional orthogonal and unitary groups O and U, which satisfy $O \simeq_w \Omega^8 O$ and $U \simeq_w \Omega^2 U$ by a strong form of the Bott periodicity theorem, as we will show in [VBKT]. So O and U give periodic Ω-spectra, hence periodic cohomology theories known as real and complex K-theory.

Exercises

1. Fill in the details of the alternative proof that $H^n(X; G) \approx \langle X, K(G, n)\rangle$ for CW complexes X, sketched above.

f. Spectra and Homology Theories

We saw in the preceding section how cohomology theories have a homotopy-theoretic interpretation in terms of Ω-spectra, and it is natural to inquire whether there is a corresponding description of homology theories. In this case we do not already have a homotopy-theoretic description of ordinary homology to serve as a starting point. But there is another homology theory we have encountered which does have a very homotopy-theoretic flavor:

Proposition 4.69. *Stable homotopy groups $\pi_n^s(X)$ define a reduced homology theory on the category of basepointed CW complexes and basepoint-preserving maps.*

Proof: In the preceding section we reformulated the axioms for a cohomology theory so that the exactness axiom asserts just the exactness of $h^n(X/A) \to h^n(X) \to h^n(A)$ for CW pairs (X, A). In order to derive long exact sequences, the reformulated axioms require also that natural suspension isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$ be specified as part of the cohomology theory. The analogous reformulation of the axioms for a homology theory is valid as well, by the same argument, and we shall use this in what follows.

For stable homotopy groups, suspension isomorphisms $\pi_n^s(X) \approx \pi_{n+1}^s(\Sigma X)$ are automatic, so it remains to verify the three axioms. The homotopy axiom is apparent. The exactness of a sequence $\pi_n^s(A) \to \pi_n^s(X) \to \pi_n^s(X/A)$ follows from exactness of $\pi_n(A) \to \pi_n(X) \to \pi_n(X, A)$ together with the isomorphism $\pi_n(X, A) \approx {}_n(X/A)$ which holds under connectivity assumptions which are achieved after sufficiently many suspensions. The wedge sum axiom $\pi_n^s(\bigvee_\alpha X_\alpha) \approx \bigoplus_\alpha \pi_n^s(X_\alpha)$ reduces to the case of finitely many summands by the usual compactness argument, and the case of finitely many summands reduces to the case of two summands by induction. Then we have isomorphisms $\pi_{n+i}(\Sigma^i X \vee \Sigma^i Y) \approx \pi_{n+i}(\Sigma^i X \times \Sigma^i Y) \approx \pi_{n+i}(\Sigma^i X) \bigoplus \pi_{n+i}(\Sigma^i Y)$, the first of these isomorphisms holding when $n+i < 2i-1$, or $i > n+1$, since $\Sigma^i X \vee \Sigma^i Y$ is the $(2i-1)$-skeleton of $\Sigma^i X \times \Sigma^i Y$. Passing to the limit over increasing i, we get the desired isomorphism $\pi_n^s(X \vee Y) \approx \pi_n^s(X) \bigoplus \pi_n^s(Y)$. □

A modest generalization of the homology theory defined by stable homotopy groups is to define $h_n(X) = \pi_n^s(X \wedge K)$ for a fixed complex K. That this is a homology theory reduces to the case of stable homotopy groups themselves by basic properties of smash product:

- $h_n(X) \approx h_{n+1}(\Sigma X)$ since $\Sigma(X \wedge K) = (\Sigma X) \wedge K$, both spaces being $S^1 \wedge X \wedge K$.
- The exactness axiom holds since $(X \wedge K)/(A \wedge K) = (X/A) \wedge K$, both spaces being quotients of $X \times K$ with $A \times K \cup X \times \{k_0\}$ collapsed to a point.
- The wedge axiom follows from the distributivity property $(\bigvee_\alpha X_\alpha) \wedge K = \bigvee_\alpha (X_\alpha \wedge K)$.

The coefficients of this homology theory are $h_n(S^0) = \pi^s_n(S^0 \wedge K) = \pi^s_n(K)$ since $S^0 \wedge K = K$.

Suppose for example that K is an Eilenberg-MacLane space $K(G,n)$. Since $K(G,n)$ is $(n-1)$-connected, its stable homotopy groups are the same as its unstable homotopy groups below dimension $2n$. Thus if we shift dimensions by defining $h_i(X) = \pi^s_{i+n}\big(X \wedge K(G,n)\big)$ we obtain a homology theory whose coefficient groups below dimension n are the same as ordinary homology with coefficients in G. It follows as in Proposition 2.29 that this homology theory agrees with ordinary homology for CW complexes of dimension less than n (or perhaps $n-1$).

This dimension restriction could be removed if there were a 'stable Eilenberg-MacLane space' whose stable homotopy groups were zero except in one dimension. However, this is a lot to ask for, so instead one seeks to form a limit of the groups $\pi^s_{i+n}\big(X \wedge K(G,n)\big)$ as n goes to infinity. The spaces $K(G,n)$ for varying n are related by weak homotopy equivalences $K(G,n) \to \Omega K(G,n+1)$. Since suspension plays such a large role in the current discussion, let us consider instead the corresponding map $\Sigma K(G,n) \to K(G,n+1)$, or to write this more concisely, $\Sigma K_n \to K_{n+1}$ This induces a map $\pi^s_{i+n}(X \wedge K_n) = \pi^s_{i+n+1}(X \wedge \Sigma K_n) \to \pi^s_{i+n+1}(X \wedge K_{n+1})$. Via these maps, it then makes sense to consider the direct limit as n goes to infinity, the group $h_i(X) = \varinjlim \pi^s_{i+n}(X \wedge K_n)$. This gives a homology theory since direct limits preserve exact sequences so the exactness axiom holds, and direct limits preserve isomorphisms so the suspension isomorphism and the wedge axiom hold. The coefficient groups of this homology theory are the same as for ordinary homology with G coefficients since $h_i(S^0) = \varinjlim \pi^s_{i+n}(K_n)$ is zero unless $i = 0$, when it is G. Hence this homology theory coincides with ordinary homology by Theorem 2.29.

To place this result in its natural generality, define a **spectrum** to be a sequence of CW complexes K_n together with basepoint-preserving maps $\Sigma K_n \to K_{n+1}$. This generalizes the notion of an Ω-spectrum, where the maps $\Sigma K_n \to K_{n+1}$ come from weak homotopy equivalences $K_n \to \Omega K_{n+1}$. Another obvious family of examples is **suspension spectra**, where one starts with an arbitrary CW complex X and defines $K_n = \Sigma^n X$ with $\Sigma K_n \to K_{n+1}$ the identity map.

The homotopy groups of a spectrum K are defined to be $\pi_i(K) = \varinjlim \pi_{i+n}(K_n)$ where the direct limit is computed using the compositions $\pi_{i+n}(K_n) \xrightarrow{\Sigma} \pi_{i+n+1}(\Sigma K_n) \to \pi_{i+n+1}(K_{n+1})$, the latter map being induced by the given map $\Sigma K_n \to K_{n+1}$. Thus in the case of the suspension spectrum of a space X, the homotopy groups of the spectrum are the same as the stable homotopy groups of X. For a general spectrum K we could also describe $\pi_i(K)$ as $\varinjlim \pi^s_{i+n}(K_n)$ since the composition $\pi_{i+n}(K_n) \to \pi_{i+n+j}(K_{n+j})$ factors

through $\pi_{i+n+j}(\Sigma^j K_n)$. So the homotopy groups of a spectrum are 'stable homotopy groups' essentially by definition.

Returning now to the context of homology theories, if we are given a spectrum K and a CW complex X, then we have a spectrum $X \wedge K$ with $(X \wedge K)_n = X \wedge K_n$, and taking the obvious maps $\Sigma(X \wedge K_n) = X \wedge \Sigma K_n \to X \wedge K_{n+1}$. The groups $\pi_i(X \wedge K)$ are the groups $\varinjlim \pi^s_{i+n}(K_n)$ considered earlier in the case of an Eilenberg-MacLane spectrum, and the arguments given there show:

Proposition 4.70. *For a spectrum K, the groups $h_i(X) = \pi_i(X \wedge K)$ form a reduced homology theory. When K is the Eilenberg-MacLane sprectrum $K_n = K(G, n)$ this homology theory is ordinary homology, so $\pi_i(X \wedge K) \approx \widetilde{H}_i(X; G)$.* □

It is also possible to associate a cohomology theory to an arbitrary spectrum K by setting $h^i(X) = \varinjlim \langle \Sigma^n X, K_{n+i} \rangle$, the direct limit with respect to the compositions $\langle \Sigma^n X, K_{n+i} \rangle \xrightarrow{\Sigma} \langle \Sigma^{n+1} X, \Sigma K_{n+i} \rangle \to \langle \Sigma^{n+1} X, K_{n+i+1} \rangle$. Verification of the axioms is an exercise using the techniques employed above. Some of the more important cohomology theories such as the various cobordism theories are usually defined in this way, as cohomology theories associated to spectra. Another example is **stable cohomotopy groups** $\pi^i_s(X) = \varinjlim \langle \Sigma^n X, S^{n+i} \rangle$, the cohomology theory associated to the sphere spectrum $S = \{S^n\}$.

By the results in the preceding section, each cohomology theory corresponds to an Ω-spectrum, so it follows that there is a natural way of associating an Ω-spectrum to each spectrum, and this is not hard to describe explicitly. Given a spectrum K, the associated Ω-spectrum K' has $K'_n = \varinjlim \Omega^i K_{n+i}$, the mapping telescope of the sequence $K_n \to \Omega K_{n+1} \to \Omega^2 K_{n+2} \to \cdots$. This is an Ω-spectrum via weak equivalences

$$K'_n = \varinjlim \Omega^i K_{n+i} \simeq \varinjlim \Omega^{i+1} K_{n+i+1} \longrightarrow \Omega \varinjlim \Omega^i K_{n+i+1} = \Omega K'_{n+1}$$

Here the first homotopy equivalence comes from deleting the first term of the sequence $K_n \to \Omega K_{n+1} \to \Omega^2 K_{n+2} \to \cdots$, and the other displayed map is a weak equivalence as a special case of an exercise at the end of the section. (Strictly speaking, we should let K'_n be a CW approximation to the mapping telescope $\varinjlim \Omega^i K_{n+i}$ in order to obtain a spectrum consisting of CW complexes, in accordance with our definition of spectrum.)

In case one starts with a suspension spectrum $K_n = \Sigma^n K$ it is not necessary to take mapping telescopes, one can just set $K'_n = \bigcup_i \Omega^i \Sigma^{i+n} K$, the union with respect to the natural inclusions $\Omega^i \Sigma^i X \subset \Omega^{i+1} \Sigma^{i+1} X$ for $X = \Sigma^n K$. This union is often denoted $\Omega^\infty \Sigma^\infty X$, or more briefly QX. Thus Q is a functor converting stable homotopy groups into ordinary homotopy groups: $\pi_i(QX) = \pi^s_i(X)$.

It follows routinely from the definitions that the homology theory defined by a spectrum is the same as the homology theory defined by the associated Ω-spectrum, and likewise for cohomology.

One may ask whether every homology theory is defined by a spectrum, as we showed for cohomology. The answer is yes if one replaces the wedge axiom by a stronger **direct limit axiom**: $h_i(X) = \varinjlim h_i(X_\alpha)$, the direct limit over the finite subcomplexes X_α of X. The homology theory defined by a spectrum satisfies this axiom, and for the converse see [Adams71].

Spectra have become the preferred language for describing many stable phenomena in algebraic topology. The increased flexibility of spectra is not without its price, however, since a number of concepts which are elementary for spaces become quite a bit more subtle for spectra, such as the proper definition of a map between spectra, or the smash product of two spectra. For the reader who wants to learn more about this language a good starting point is [Adams74].

Exercises

1. Assuming the first two axioms for a homology theory on the CW category, show that the direct limit axiom implies the wedge sum axiom. Show that the converse also holds for countable CW complexes.

2. For CW complexes X and Y consider the suspension sequence

$$\langle X, Y\rangle \xrightarrow{\Sigma} \langle \Sigma X, \Sigma Y\rangle \xrightarrow{\Sigma} \langle \Sigma^2 X, \Sigma^2 Y\rangle \longrightarrow \cdots$$

Show that if X is a finite complex, these maps eventually become isomorphisms. [Use induction on the number of cells of X and the five-lemma.]

3. Show that for any sequence $Z_1 \to Z_2 \to \cdots$, the natural map $\varinjlim \Omega Z_n \to \Omega \varinjlim Z_n$ is a weak homotopy equivalence, where the direct limits mean mapping telescopes.

g. Eckmann-Hilton Duality

There is a very nice duality principle in homotopy theory, called Eckmann-Hilton duality in its more refined and systematic aspects, but which in its most basic form involves the simple idea of reversing the direction of all arrows in a given construction. For example, if in the definition of a fibration as a map satisfying the homotopy lifting property we reverse the direction of all the arrows, we obtain the dual notion of a **cofibration**. This is a map $i : A \to B$ satisfying the following property: Given $\widetilde{g}_0 : B \to X$ and a homotopy $g_t : A \to X$ such that $\widetilde{g}_0 i = g_0$, there exists a homotopy $\widetilde{g}_t : B \to X$ such that $\widetilde{g}_t i = g_t$.

$$\begin{array}{ccc} A & \xrightarrow{g_t} & X \\ {\scriptstyle i}\downarrow & \nearrow_{\widetilde{g}_0} & \\ B & & \end{array}$$

In the special case that i is the inclusion of a subspace, this is the homotopy extension property, and the next proposition says that this is indeed the general case. So a cofibration is the same as an inclusion satisfying the homotopy extension property.

Proposition 4.71. *If $i : A \to B$ is a cofibration, then i is injective, and in fact a homeomorphism onto its image.*

Proof: Consider the mapping cylinder M_i, the quotient of $A \times I \amalg B$ in which $(a, 1)$ is identified with $i(a)$. Let $g_t : A \to M_i$ be the homotopy mapping $a \in A$ to the image of $(a, 1-t) \in A \times I$ in M_i, and let $\widetilde{g}_0$ be the inclusion $B \hookrightarrow M_i$. The cofibration property gives $\widetilde{g}_t : B \to M_i$ with $\widetilde{g}_t i = g_t$. Restricting to a fixed $t > 0$, this implies i is injective since g_t is. Furthermore, since g_t is a homeomorphism onto its image $A \times \{1-t\}$, the relation $\widetilde{g}_t i = g_t$ implies that the map $g_t^{-1}\widetilde{g}_t : i(A) \to A$ is a continuous inverse of $i : A \to i(A)$. □

Many constructions for fibrations have analogs for cofibrations, and vice versa. For example, for an arbitrary map $f : A \to B$ the inclusion $A \hookrightarrow M_f$ is readily seen to be a cofibration, so the analog of the factorization $A \hookrightarrow E_f \to B$ of f into a homotopy equivalence followed by a fibration is the factorization $A \hookrightarrow M_f \to B$ into a cofibration followed by a homotopy equivalence. Even the definition of M_f is in some way dual to the definition of E_f, since E_f can be defined as a pullback and M_f can be defined as a dual 'pushout:'

$$\begin{array}{ccc} E_f & \longrightarrow & B^I \\ \downarrow & & \downarrow \\ A & \longrightarrow & B \end{array} \qquad\qquad \begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ A \times I & \longrightarrow & M_f \end{array}$$

In general, the **pushout** of maps $f: Z \to X$ and $g: Z \to Y$ is defined as the quotient of $X \amalg Y$ under the identifications $f(z) \sim g(z)$. Thus the pushout is a quotient of the 'sum' $X \amalg Y$, while the pullback of maps $X \to Z$ and $Y \to Z$ is a subobject of the product $X \times Y$. Since the dual of a subobject is a quotient object, the **cofiber** of a cofibration $A \hookrightarrow B$ is the quotient space B/A. In particular, making an arbitrary map $f: A \to B$ into a cofibration $A \hookrightarrow M_f$, the cofiber is the mapping cone $C_f = M_f/(A \times \{0\})$.

In a similar vein, the dual of the fibration sequence

$$\cdots \to \Omega F \to \Omega E \to \Omega B \to F \to E \to B$$

is the cofibration sequence

$$A \to X \to X/A \to \Sigma A \to \Sigma X \to \Sigma(X/A) \to \cdots$$

considered in §4.4e. So Σ and Ω are in some way dual operations, and the fact that maps $\Sigma X \to Y$ are the same as maps $X \to \Omega Y$ is expressed by saying Σ and Ω are **adjoint functors**.

It is interesting to think about limits from the viewpoint of duality. The algebraic notions of direct and inverse limits of groups are clearly dual. Topologically, we have encountered direct limits of spaces in §3.4f and inverse limits of spaces in the construction of Postnikov towers in §4.3. The direct limit of an arbitrary sequence of maps $X_0 \to X_1 \to \cdots$ is rarely encountered in algebraic topology and would not be easy to handle, but things are much nicer if the maps are cofibrations, for example when one expresses a CW complex as the union of an increasing sequence of subcomplexes. The mapping telescope construction serves to turn an arbitrary sequence of maps into a sequence of cofibrations. Dually, the inverse limit of a sequence of maps $\cdots \to X_1 \to X_0$ is really only useful when the maps are fibrations. We encountered this issue in §4.3, where in reconstructing a space from its Postnikov tower we first turned all the maps in the tower into fibrations by iterating the usual pathspace construction. When this iterated pathspace construction is applied to an arbitrary sequence $\cdots \to X_1 \to X_0$, the resulting inverse limit space is often denoted $holim\, X_i$, for 'homotopy inverse limit.' The dual construction of turning a sequence of maps $X_0 \to X_1 \to \cdots$ into cofibrations and then taking the direct limit produces a space called, unfortunately, $hocolim\, X_i$, the 'homotopy direct limit.' Presumably there is some categorical justification for direct limit being 'colim' and inverse limit 'lim,' rather than vice versa.

Strictly speaking, if one iterates the mapping cylinder construction for a sequence of maps $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots$, what one gets is not the mapping telescope, but a larger

space which might be called an iterated mapping cylinder $M(f_1, f_2, \cdots)$. This works for a finite sequence of maps just as well as for an infinite sequence. Namely, $M(f_1)$ is just the usual mapping cylinder of f_1, $M(f_1, f_2)$ is the mapping cylinder of the composition $M(f_1) \to X_1 \to X_2$, $M(f_1, f_2, f_3)$ is the mapping cylinder of the composition $M(f_1, f_2) \to X_2 \to X_3$, etc., and $M(f_1, f_2, \cdots)$ is the union of this increasing sequence of iterated mapping cylinders $M(f_1, \cdots, f_n)$.

It is possible to form homotopy direct limits of more complicated diagrams of mappings, not just linear sequences. Suppose one has a Δ-complex K and a function assigning a space to each vertex of K and a map to each directed edge of K, such that the maps corresponding to the edges of each simplex of K form a commutative diagram. Then to each n-simplex of K there is an associated iterated mapping cylinder $M(f_1, \cdots, f_n)$, and these iterated mapping cylinders all fit together naturally, modelled on the way the simplices of K fit together, to form a big space which could be called 'hocolim' of the given complex of spaces. For example, for a linear sequence $X_0 \to X_1 \to \cdots$ this produces the mapping telescope, and if we include as well all the compositions of the mappings in this sequence then we get the infinite iterated mapping cylinder $M(f_1, f_2, \cdots)$. When K is one-dimensional this construction arose also in §1.4c in the context of graphs of groups.

Homotopy Groups with Coefficients

At a somewhat deeper level there is a certain duality between homotopy groups and homology/cohomology groups, since $\pi_n(X)$ is homotopy classes of maps from a space with a single nonzero homology group to X, while $H^n(X; G)$ is homotopy classes of maps from X to a space with a single nonzero homotopy group. This duality becomes more complete if one defines **homotopy groups with coefficients** by $\pi_n(X; G) = \langle M(G, n), X \rangle$, where $M(G, n)$ is a Moore space, so that $\pi_n(X; \mathbb{Z}) = \pi_n(X)$ since S^n is an $M(\mathbb{Z}, n)$. We should assume $n > 1$ to guarantee that the homotopy type of $M(G, n)$ is well-defined by Proposition 4.15. For $n > 1$, $\pi_n(X; G)$ is a group since we can choose $M(G, n)$ to be the suspension of an $M(G, n-1)$. Moreover, if $n > 2$ then $\pi_n(X; G)$ is abelian since we can choose $M(G, n)$ to be a double suspension.

There is a universal coefficient theorem for homotopy groups with coefficients:

Proposition 4.72. *For $n > 1$ there are natural short exact sequences*

$$0 \longrightarrow Ext(G, \pi_{n+1}(X)) \longrightarrow \pi_n(X; G) \longrightarrow Hom(G, \pi_n(X)) \longrightarrow 0.$$

This may look a bit strange at first glance, but notice that $\pi_n(X; G)$ is covariant as a functor of X and contravariant as a functor of G, just like the Ext and Hom terms in the exact sequence.

Proof: Let $0 \to R \xrightarrow{i} F \longrightarrow G \to 0$ be a free resolution of G. The inclusion map i is realized by a map $M(R,n) \to M(F,n)$, where $M(R,n)$ and $M(F,n)$ are wedges of S^n's corresponding to bases of F and R. Converting this into a cofibration via the mapping cylinder, the cofiber is an $M(G,n)$, as one sees from the long exact sequence of homology groups. As in §4.4e, if we map the cofibration sequence $M(R,n) \to M(F,n) \to M(G,n) \to M(R,n+1) \to M(F,n+1)$ into X we get an exact sequence

$$\begin{array}{ccccccccc}
\pi_{n+1}(X;F) & \longrightarrow & \pi_{n+1}(X;R) & \longrightarrow & \pi_n(X;G) & \longrightarrow & \pi_n(X;F) & \longrightarrow & \pi_n(X;R) \\
\| & & \| & & & & \| & & \| \\
Hom(F,\pi_{n+1}(X)) & \xrightarrow{i^*} & Hom(R,\pi_{n+1}(X)) & & & & Hom(F,\pi_n(X)) & \xrightarrow{i^*} & Hom(R,\pi_n(X))
\end{array}$$

Mapping a wedge of S^n's into X amounts to assigning an element of $\pi_n(X)$ to each wedge summand, which gives the identifications with the Hom terms in the diagram. The kernel and cokernel of i^* are $Hom(G,-)$ and $Ext(G,-)$ by definition, and the result follows. □

Unlike in the universal coefficient theorems for homology and cohomology, the short exact sequence in this proposition does not split in general. An example is $G = \mathbb{Z}_2$, $X = M(\mathbb{Z}_2,n)$ for $n \geq 2$, where $\pi_n(M(\mathbb{Z}_2,n);\mathbb{Z}_2) = \langle M(\mathbb{Z}_2,n), M(\mathbb{Z}_2,n)\rangle$ is $\mathbb{Z}_4$ generated by the identity map of $M(\mathbb{Z}_2,n)$ (see [Neisendorfer]) while the two outer terms in the short exact sequence can only contain elements of order 2 since $G = \mathbb{Z}_2$. This example also shows that $\pi_n(X;\mathbb{Z}_m)$ need not be a module over $\mathbb{Z}_m$, as $H_n(X;\mathbb{Z}_m)$ is.

Homology Decompositions

As an organizing principle, Eckmann-Hilton duality can be extremely helpful, reducing significantly what one has to remember and giving valuable hints on how to proceed in various situations. To illustrate, let us consider what would happen if we dualized the notion of a Postnikov tower of principal fibrations, where a space is represented as an inverse limit of a sequence of fibers of maps to Eilenberg-MacLane spaces. In the dual representation, a space would be realized as a direct limit of a sequence of cofibers of maps from Moore spaces. To describe this in more detail, suppose we are given a sequence of abelian groups G_n, $n \geq 1$, and we build a CW complex X with $H_n(X) \approx G_n$ for all n by constructing inductively an increasing sequence of subcomplexes $X_1 \subset X_2 \subset \cdots$ with $H_i(X_n) \approx G_i$ for $i \leq n$ and $H_i(X_n) = 0$ for $i > n$, where

(1) X_1 is a Moore space $M(G_1,1)$.

(2) X_{n+1} is the mapping cone of a cellular map $h_n : M(G_{n+1},n) \to X_n$ such that the induced map $h_{n*} : H_n(M(G_{n+1},n)) \to H_n(X_n)$ is trivial.

(3) $X = \bigcup_n X_n$.

One sees inductively that X_{n+1} has the desired homology groups by comparing the long exact sequences of the pairs (X_{n+1}, X_n) and (CM, M) where $M = M(G_{n+1}, n)$ and CM is the cone $M \times I/M \times \{0\}$:

$$\begin{array}{ccccccccc} 0 \longrightarrow & H_{n+1}(X_{n+1}) & \longrightarrow & H_{n+1}(X_{n+1}, X_n) & \xrightarrow{\partial} & H_n(X_n) & \longrightarrow & H_n(X_{n+1}) & \longrightarrow 0 \\ & & & \uparrow\approx & & \uparrow h_{n*} & & & \\ & & & H_{n+1}(CM, M) & \xrightarrow[\approx]{\partial} & H_n(M) \approx G_{n+1} & & & \end{array}$$

The assumption that h_{n*} is trivial means that the boundary map in the upper row is zero, hence $H_{n+1}(X_{n+1}) \approx G_{n+1}$. The other homology groups of X_{n+1} are the same as those of X_n since $H_i(X_{n+1}, X_n) \approx H_i(CM, M)$ for all i by excision, and $H_i(CM, M) \approx \widetilde{H}_{i-1}(M)$ since CM is contractible.

In case all the maps h_n are trivial, X is the wedge sum of the Moore spaces $M(G_n, n)$ since in this case the mapping cone construction in (2) produces a wedge sum with the suspension of $M(G_{n+1}, n)$, a Moore space $M(G_{n+1}, n+1)$.

For a space Y, a homotopy equivalence $f : X \to Y$ where X is constructed as in (1)-(3) is called a **homology decomposition** of Y.

Theorem 4.73. *Every simply-connected CW complex has a homology decomposition.*

Proof: Given a simply-connected CW complex Y, let $G_n = H_n(Y)$. Suppose inductively that we have constructed X_n via maps h_i as in (2), together with a map $f : X_n \to Y$ inducing an isomorphism on H_i for $i \le n$. The induction can start with X_1 a point since Y is simply-connected. To construct X_{n+1} we first replace Y by the mapping cylinder of $f : X_n \to Y$, converting f into an inclusion. By the Hurewicz theorem and the exact homology sequence of the pair (Y, X_n) we have $\pi_{n+1}(Y, X_n) \approx H_{n+1}(Y, X_n) \approx H_{n+1}(Y) = G_{n+1}$. We use this isomorphism to construct $h_n : M(G_{n+1}, n) \to X_n$ and an extension $f : X_{n+1} \to Y$ as follows.

The standard construction of an $M(G_{n+1}, n)$ consists of a wedge of spheres S^n_α corresponding to generators g_α of G_{n+1}, with cells e^{n+1}_β attached according to certain linear combinations $r_\beta = \sum_\alpha n_{\alpha\beta} g_\alpha$ which are zero in G_{n+1}. Under the isomorphism $G_{n+1} \approx \pi_{n+1}(Y, X_n)$ each g_α corresponds to a basepoint-preserving map $f_\alpha : (CS^n, S^n) \to (Y, X_n)$ where CS^n is the cone on S^n. The restrictions of these f_α's to S^n define $h_n : \bigvee_\alpha S^n_\alpha \to X_n$, and the maps $f_\alpha : CS^n \to Y$ themselves give an extension of $f : X_n \to Y$ to the mapping cone of $h_n : \bigvee_\alpha S^n_\alpha \to X_n$. Each relation r_β gives a homotopy $F_\beta : (CS^n, S^n) \times I \to (Y, X_n)$ from $\sum_\alpha n_{\alpha\beta} f_\alpha$ to the constant map. We use $F_\beta \big| S^n \times \{0\}$ to attach e^{n+1}_β, and then $F_\beta \big| S^n \times I$ gives h_n on e^{n+1}_β and F_β gives an extension of f over the cone on e^{n+1}_β.

This construction assures that $f_*: H_{n+1}(X_{n+1}, X_n) \to H_{n+1}(Y, X_n)$ is an isomorphism, so from the five-lemma applied to the long exact sequences of these pairs we deduce that $f_*: H_i(X_{n+1}) \to H_i(Y)$ is an isomorphism for $i \le n+1$. This finishes the induction step. We may assume the maps f_α and F_β are cellular, so $X = \bigcup_n X_n$ is a CW complex with subcomplexes X_n. Since $f: X \to Y$ is a homology isomorphism between simply-connected CW complexes, it is a homotopy equivalence. □

As an example, suppose that Y is a simply-connected CW complex having all its homology groups free. Then the Moore spaces used in the construction of X can be taken to be wedges of spheres, and so X_n is obtained from X_{n-1} by attaching an n-cell for each $\mathbb{Z}$ summand of $H_n(Y)$. The attaching maps may be taken to be cellular, making X into a CW complex whose cellular chain complex has trivial boundary maps. Similarly, finite cyclic summands of $H_n(Y)$ can be realized by wedge summands $S^{n-1} \cup e^n$ in $M(H_n(Y), n-1)$, contributing an n-cell and an $(n+1)$-cell to X. This is Proposition 4.52, but the present result is stronger because it tells us that a finite cyclic summand of H_n can be realized in one step by attaching the cone on a Moore space $M(\mathbb{Z}_k, n-1)$, rather than in two steps of attaching an n-cell and then an $(n+1)$-cell.

Exercises

1. Show that if $A \hookrightarrow X$ is a cofibration of compact Hausdorff spaces, then for any space Y, the map $Y^X \to Y^A$ obtained by restriction of functions is a fibration. [If $A \hookrightarrow X$ is a cofibration, so is $A \times Y \hookrightarrow X \times Y$ for any space Y.]

2. Consider a pushout diagram

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \downarrow & & \downarrow \\ X & \longrightarrow & B \cup_f X \end{array}$$

where $B \cup_f X$ is B with X attached along A via f. Show that if $A \hookrightarrow X$ is a cofibration, so is $B \hookrightarrow B \cup_f X$.

3. What is the dual of a mapping telescope, precisely?

4. Show that the iterated mapping cylinder $M(f_1, f_2, \cdots)$ of a sequence of maps $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots$, deformation retracts onto the mapping telescope of this sequence.

h. Stable Splittings of Spaces

In this section we describe some examples of an interesting phenomenon: When one suspends a space, it can become much simpler, becoming homotopy equivalent to a wedge sum of smaller spaces. The simplest example of this is the torus $S^1 \times S^1$, whose reduced suspension $\Sigma(S^1 \times S^1)$ is homotopy equivalent to $S^2 \vee S^2 \vee S^3$ since $\Sigma(S^1 \times S^1)$ is $S^2 \vee S^2$ with a 3-cell attached by the suspension of the attaching map of the 2-cell of the torus, but the latter attaching map is the commutator of the two inclusions $S^1 \hookrightarrow S^1$, and the suspension of this commutator is trivial since it lies in the abelian group $\pi_2(S^2 \vee S^2)$.

More generally:

Proposition 4.74. *If X and Y are connected CW complexes, then $\Sigma(X \times Y) \simeq \Sigma X \vee \Sigma Y \vee \Sigma(X \wedge Y)$.*

For example, $\Sigma(S^m \times S^n) \simeq S^{m+1} \vee S^{n+1} \vee S^{m+n+1}$. In view of the cup product structure on $H^*(S^m \times S^n)$ there can be no such splitting of $S^m \times S^n$ before suspension.

The proof of the proposition will use the natural group structure in $\langle \Sigma Z, W \rangle$, defined in the paragraph preceding Proposition 4.62. In particular, we will use the fact that for basepoint-preserving maps $f, g : \Sigma Z \to W$, the induced maps on homology satisfy $(f+g)_* = f_* + g_*$. This is proved by the same argument as in Proposition 4.6, replacing the disk D^n there by the reduced suspension ΣZ.

Proof: From the long exact sequence of the pair $(X \times Y, X \vee Y)$ we obtain splittings $H_i(X \times Y) \approx H_i(X) \bigoplus H_i(Y) \bigoplus H_i(X \wedge Y)$ for $i > 0$, with the projections on the three summands induced by the projections of $X \times Y$ onto X, Y, and $X \wedge Y$. We can add the suspensions of these three projections to get a map $f : \Sigma(X \times Y) \to \Sigma X \vee \Sigma Y \vee \Sigma(X \wedge Y)$. By the remark preceding the proof, this induces an isomorphism on homology. Assuming X and Y are connected, the domain and range spaces of f are simply-connected, so Whitehead's theorem implies that f is a homotopy equivalence. □

This result can be applied inductively to obtain splittings for products of more that two spaces, using the fact that reduced suspension is smash product with S^1, and smash product is associative and commutative. For example,

$$\Sigma(X \times Y \times Z) \simeq \Sigma X \vee \Sigma Y \vee \Sigma Z \vee \Sigma(X \wedge Y) \vee \Sigma(X \wedge Z) \vee \Sigma(Y \wedge Z) \vee \Sigma(X \wedge Y \wedge Z)$$

Our next example involves the reduced product $J(X)$ defined in §3.2. An interesting case is the space $J(S^n)$ which has a CW structure of the form $S^n \cup e^{2n} \cup e^{3n} \cup \cdots$. The cells e^{in} are all attached nontrivially since $H^*(J(S^n); \mathbb{Q})$ is a polynomial ring $\mathbb{Q}[x]$ for n

even and a tensor product $\mathbb{Q}[x]\otimes\Lambda_{\mathbb{Q}}[y]$ for n odd. However, after we suspend to $\Sigma J(S^n)$, it is a rather surprising fact that all the attaching maps become trivial:

Proposition 4.75. *$\Sigma J(S^n) \simeq S^{n+1} \vee S^{2n+1} \vee S^{3n+1} \vee \cdots$. More generally, if X is a connected CW complex then $\Sigma J(X) \simeq \bigvee_n \Sigma X^{\wedge n}$ where $X^{\wedge n}$ denotes the smash product of n copies of X.*

Proof: The space $J(X)$ is the union of an increasing sequence of subcomplexes $J_k(X)$ with $J_k(X)$ a quotient of the k-fold product $X^{\times k}$. The quotient $J_k(X)/J_{k-1}(X)$ is $X^{\wedge k}$.

$$\begin{array}{ccccc} & & X^{\times k} & & \\ & & \downarrow & \searrow & \\ J_{k-1}(X) & \longrightarrow & J_k(X) & \longrightarrow & X^{\wedge k} \end{array}$$

By repeated application of the preceding proposition, $\Sigma X^{\wedge k}$ is a wedge summand of $\Sigma X^{\times k}$, up to homotopy equivalence. The proof shows moreover that there is a map $\Sigma X^{\wedge k} \to \Sigma X^{\times k}$ such that the composition $\Sigma X^{\wedge k} \to \Sigma X^{\times k} \to \Sigma X^{\wedge k}$ is homotopic to the identity. By composition with the quotient map $X^{\times k} \to J_k(X)$ we obtain a map $s_k : X^{\wedge k} \to J_k(X)$ such that $X^{\wedge k} \xrightarrow{s_k} J_k(X) \to X^{\wedge k}$ is homotopic to the identity.

From the long exact sequence of homology groups for the pair $(\Sigma J_k(X), \Sigma J_{k-1}(X))$ we see that the map $i \vee s_k : \Sigma J_{k-1}(X) \vee \Sigma X^{\wedge k} \to \Sigma J_k(X)$ induces an isomorphism on homology, where i denotes the inclusion map. It follows by induction that the map $\bigvee_{k=1}^n s_k : \bigvee_{k=1}^n \Sigma X^{\wedge k} \to J_n(X)$ induces an isomorphism on homology, for all finite n. This implies the corresponding statement for $n = \infty$ since $X^{\wedge n}$ is $(n-1)$-connected if X is connected. Thus we have a map $\bigvee_k \Sigma X^{\wedge k} \to \Sigma J(X)$ inducing an isomorphism on homology, so this map is a homotopy equivalence since the spaces involved are simply-connected CW complexes. □

Here is another stable splitting result of a different sort:

Proposition 4.76. *For any prime power p^n there is a splitting $\Sigma K(\mathbb{Z}_{p^n}, 1) \simeq X_1 \vee \cdots \vee X_{p-1}$ where X_i is a CW complex having $\widetilde{H}_*(X_i;\mathbb{Z})$ nonzero only in dimensions congruent to $2i$ mod $2p-2$.*

This result is best possible in a strong sense: No matter how many times any one of the spaces X_i is suspended, it never becomes homotopy equivalent to a nontrivial wedge sum. This will be shown in §4.4k by studying cohomology operations in $H^*(K(\mathbb{Z}_{p^n}, 1); \mathbb{Z}_p)$. There is also a somewhat simpler K-theoretic explanation for this; see [VBKT].

Proof: Let $K = K(\mathbb{Z}_{p^n}, 1)$. The multiplicative group of nonzero elements in the field $\mathbb{Z}_p$ is cyclic, so let the integer r represent a generator. As we saw in Example 3 in §3.4h, there

is a map $f:K\to K$ inducing multiplication by r on $\pi_1(K)$, and hence multiplication by r^i on $H^{2i}(K;\mathbb{Z})\approx\mathbb{Z}_{p^n}$. For each integer $j\ge 0$ let $h_j:\Sigma K\to\Sigma K$ be the difference $\Sigma f - r^j\mathbb{1}$, $\mathbb{1}$ being the identity map of ΣK. Then h_j induces multiplication by r^i-r^j on $H^{2i+1}(\Sigma K;\mathbb{Z})\approx\mathbb{Z}_{p^n}$, hence also on $H_{2i}(\Sigma K;\mathbb{Z})\approx\mathbb{Z}_{p^n}$ by the universal coefficient theorem. By the choice of r we know that p divides r^i-r^j iff $i\equiv j \mod p-1$. This means that the map induced by h_j on $\widetilde{H}_{2i}(\Sigma K;\mathbb{Z})$ has nontrivial kernel iff $i\equiv j \mod p-1$. Therefore the composition $m_i = h_1\circ\cdots h_{i-1}\circ h_{i+1}\cdots h_{p-1}$ induces an isomorphism on $\widetilde{H}_*(\Sigma K;\mathbb{Z})$ in dimensions congruent to $2i \mod 2p-2$ and has a nontrivial kernel in other dimensions. When there is a nontrivial kernel, some power of m_i will induce the zero map since we are dealing with homomorphisms $\mathbb{Z}_{p^n}\to\mathbb{Z}_{p^n}$.

Let X_i be the mapping telescope of the sequence $\Sigma K\to\Sigma K\to\cdots$ where each map is m_i. Since homology commutes with direct limits, the inclusion of the first factor $\Sigma K\subset X_i$ induces an isomorphism on $\widetilde{H}_*$ in dimensions congruent to $2i \mod 2p-2$, and $\widetilde{H}_*(X_i;\mathbb{Z})=0$ in all other dimensions. The sum of these inclusions is a map $\Sigma K\to X_1\vee\cdots\vee X_{p-1}$ inducing an isomorphism on all homology groups. Since these complexes are simply-connected, the result follows by Whitehead's theorem. □

The construction of the spaces X_i as mapping telescopes produces rather large spaces, with infinitely many cells in each dimension. However, by Proposition 4.52 each X_i is homotopy equivalent to a CW complex with the minimum configuration of cells consistent with its homology, namely, a 0-cell and a k-cell for each k congruent to $2i$ or $2i+1 \mod 2p-2$.

Stable splittings of $K(G,1)$'s for finite groups G have been much studied in recent years, and are a complicated and subtle business. To take a simple example, Proposition 4.74 implies that $\Sigma K(\mathbb{Z}_2\times\mathbb{Z}_2,1)$ splits as the wedge sum of two copies of $\Sigma K(\mathbb{Z}_2,1)$ and $\Sigma(K(\mathbb{Z}_2,1)\wedge K(\mathbb{Z}_2,1))$, but the latter summand can be split further, according to a result in [Harris-Kuhn] which says that for G the direct sum of k copies of $\mathbb{Z}_{p^n}$, $\Sigma K(G,1)$ splits canonically as the wedge sum of pieces having exactly p^k-1 distinct homotopy types. Some of these summands occur more than once, for example in the case of $\mathbb{Z}_2\times\mathbb{Z}_2$.

Exercises

1. If a CW complex X retracts onto a subcomplex A, show that $\Sigma X\simeq\Sigma A\vee\Sigma(X/A)$.

2. Extending Proposition 4.76, show that the $(2k+1)$-skeleton of the suspension of a lens space with fundamental group of order p^n is homotopy equivalent to the wedge sum of the $(2k+1)$-skeleta of the spaces X_i, if these X_i's are chosen to have the minimum number of cells in each dimension, as described in the remarks following the proof.

i. The Loopspace of a Suspension

Loopspaces appear at first glance to be hopelessly complicated objects, but if one is only interested in homotopy type, there are many cases when almost miraculous simplifications occur. One of the nicest of these cases is the loopspace of a sphere. We show in this section that ΩS^{n+1} has the weak homotopy type of the James reduced product $J(S^n)$ introduced in §3.2, and more generally, that $\Omega\Sigma X$ has the weak homotopy type of $J(X)$ for every connected CW complex X. If one wants, one can strengthen 'weak homotopy type' to 'homotopy type' by quoting Milnor's theorem, mentioned in §4.3, that the loopspace of a CW complex has the homotopy type of a CW complex.

Part of the interest in $\Omega\Sigma X$ is its close connection with the suspension homomorphism $\pi_i(X) \to \pi_{i+1}(\Sigma X)$, and we will use the weak homotopy equivalence $\Omega S^{n+1} \simeq_w J(S^n)$ to get another proof that this suspension homomorphism is an isomorphism in dimensions up to approximately double the connectivity of X. In addition, we will obtain an exact sequence which measures the failure of the suspension map to be an isomorphism in dimensions between approximately double and triple the connectivity of X. An easy application of this, together with results proved elsewhere in the book, will be to compute $\pi_{n+1}(S^n)$ and $\pi_{n+2}(S^n)$ for all n.

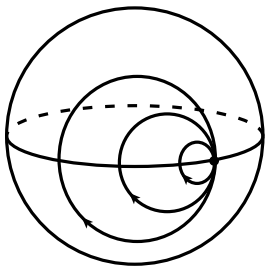

As a first approximation to ΩS^{n+1}, there is a natural inclusion $S^n \hookrightarrow \Omega S^{n+1}$ obtained by regarding S^{n+1} as the reduced suspension $\Sigma S^n = (S^n \times I)/(S^n \times \partial I \cup \{e\} \times I)$ where e is the basepoint of S^n, then associating to each point $x \in S^n$ the loop $\lambda(x)$ in ΣS^n given by $t \mapsto (x, t)$. The figure shows what a few such loops look like. However, we cannot expect this inclusion $S^n \hookrightarrow \Omega S^{n+1}$ to be a homotopy equivalence since ΩS^{n+1} is an H-space but S^n is only an H-space when $n = 1, 3, 7$ by Adams' theorem. The simplest way to correct this deficiency in S^n would be to replace it by the free H-space which it generates, i.e., the reduced product $J(S^n)$. Recall from §3.2 that a point in $J(S^n)$ is a formal product $x_1 \cdots x_k$ of points $x_i \in S^n$, with the basepoint e acting as an identity element for the multiplication obtained by juxtaposition of formal products. We would like to define a map $\lambda : J(S^n) \to \Omega S^{n+1}$ by setting $\lambda(x_1 \cdots x_k) = \lambda(x_1) \cdots \lambda(x_k)$, the product of the loops $\lambda(x_i)$. The only difficulty is in the parametrization of this product, which needs to be adjusted so that λ is continuous. The problem is that when some x_i approaches the basepoint $e \in S^n$, one wants the loop $\lambda(x_i)$ to disappear gradually from the product $\lambda(x_1) \cdots \lambda(x_k)$, without disrupting the parametrization as simply deleting $\lambda(e)$ would do. This can be achieved by first making the time it takes to traverse each loop $\lambda(x_i)$ equal to the distance from x_i

to the basepoint of S^n, then normalizing the resulting product of loops to take unit time, giving a map $I \to \Sigma S^n$.

More generally, this same procedure defines a map $\lambda : J(X) \to \Omega\Sigma X$ for any connected CW complex X, where 'distance to the basepoint' is replaced by any map $d : X \to [0,1]$ with $d^{-1}(0) = e$, the basepoint of X.

Theorem 4.77. *The map $\lambda : J(X) \to \Omega\Sigma X$ is a weak homotopy equivalence for every connected CW complex X.*

Proof: The main step will be to compute the homology of $\Omega\Sigma X$. After this is done, it will be easy to deduce that λ induces an isomorphism on homology using the calculation of the homology of $J(X)$ in §3.4c, and from this conclude that λ is a weak homotopy equivalence.

Let us begin the main step, which is to show that $H_*(\Omega\Sigma X; F)$ is the tensor algebra $T\widetilde{H}_*(X; F)$ if X is a connected CW complex and F is a field. Write the reduced suspension $Y = \Sigma X$ as the union of two reduced cones $Y_+ = C_+X$ and $Y_- = C_-X$ intersecting in the equatorial $X \subset \Sigma X$. Consider the path fibration $p : PY \to Y$ with fiber ΩY. Let $P_+Y = p^{-1}(Y_+)$ and $P_-Y = p^{-1}(Y_-)$, so P_+Y consists of paths in Y starting at the basepoint and ending in Y_+, and similarly for P_-Y. Then $P_+Y \cap P_-Y$ is $p^{-1}(X)$, the paths from the basepoint to X. Since Y_+ and Y_- are deformation retracts of open neighborhoods U_+ and U_- in Y such that $U_+ \cap U_-$ deformation retracts onto $Y_+ \cap Y_- = X$, the homotopy lifting property implies that P_+Y, P_-Y, and $p_+Y \cap P_-Y$ are deformation retracts, in the weak sense, of open neighborhoods $p^{-1}(U_+)$, $p^{-1}(U_-)$, and $p^{-1}(U_+) \cap p^{-1}(U_-)$, respectively. Therefore there is a Mayer-Vietoris sequence in homology for the decomposition $PY = P_+Y \cup P_-Y$. Taking reduced homology and using the fact that PY is contractible, this gives an isomorphism

$$(i) \qquad \Phi : \widetilde{H}_*(P_+Y \cap P_-Y; F) \xrightarrow{\approx} \widetilde{H}_*(P_+Y; F) \bigoplus \widetilde{H}_*(P_-Y; F)$$

The two coordinates of Φ are induced by the inclusions, with a minus sign in one case, but Φ will still be an isomorphism if this minus sign is eliminated, so we assume this has been done.

We claim that the isomorphism Φ can be rewritten as an isomorphism

$$(ii) \qquad \Theta : \widetilde{H}_*(\Omega Y \times X; F) \xrightarrow{\approx} \widetilde{H}_*(\Omega Y; F) \bigoplus \widetilde{H}_*(\Omega Y; F)$$

To see this, we observe that the fibration $P_+Y \to Y_+$ is fiber-homotopically trivial. This is true since the cone Y_+ is contractible, but we shall need an explicit fiber homotopy

equivalence $P_+Y \simeq \Omega Y \times Y_+$ and this is easily constructed as follows. Define $f_+ : P_+Y \to \Omega Y \times Y_+$ by $f_+(\gamma) = (\gamma \cdot \gamma_y^+, y)$ where $y = \gamma(1)$ and γ_y^+ is the path in Y_+ from $y = (x,t)$ to the basepoint along a portion of the segment $\{x\} \times I$. In the other direction, define $g_+ : \Omega Y \times Y_+ \to P_+Y$ by $g_+(\gamma, y) = \gamma \cdot \overline{\gamma}_y^+$ where 'bar' denotes the inverse path. Then fg and gf are fiber-homotopic to the respective identity maps since $\overline{\gamma}_y^+ \cdot \gamma_y^+$ and $\gamma_y^+ \cdot \overline{\gamma}_y^+$ are homotopic to the constant paths.

In similar fashion the fibration $P_-Y \to Y_-$ is fiber-homotopically trivial via maps f_- and g_-. By restricting a fiber-homotopy trivialization of either P_+Y or P_-Y to $P_+Y \cap P_-Y$, we see that the fibration $P_+Y \cap P_-Y$ is fiber-homotopy equivalent to the product $\Omega Y \times X$. Let us do this using the fiber-homotopy trivialization of P_-Y. The groups in (i) can now be replaced by those in (ii). The map Φ has coordinates induced by inclusion, and it follows that the corresponding map Θ in (ii) has coordinates induced by the two maps $\Omega Y \times X \to \Omega Y$, $(\gamma, x) \mapsto \gamma \cdot \lambda(x)$ and $(\gamma, x) \mapsto \gamma$. Namely, the first coordinate of Θ is induced by $f_+g_- \,\big|\, \Omega Y \times X$ followed by projection to ΩY, and the second coordinate is the same but with f_-g_- in place of f_+g_-.

Writing the two coordinates of Θ as Θ_1 and Θ_2, the fact that Θ is an isomorphism means that the restriction of Θ_1 to the kernel of Θ_2 is an isomorphism. Via the Künneth isomorphism we can write $\widetilde{H}_*(\Omega Y \times X; F) = [H_*(\Omega Y; F) \otimes \widetilde{H}_*(X; F)] \bigoplus \widetilde{H}_*(\Omega Y; F)$ where projection onto the latter summand is Θ_2. Hence Θ_1 gives an isomorphism from the first summand $H_*(\Omega Y; F) \otimes \widetilde{H}_*(X; F)$ onto $\widetilde{H}_*(\Omega Y; F)$. Since $\Theta_1(\gamma, x) = (\gamma \cdot \lambda(x))$, this means that the composed map

$$H_*(\Omega Y; F) \otimes \widetilde{H}_*(X; F) \xrightarrow{1\!\!1 \,\otimes\, \lambda_*} H_*(\Omega Y; F) \otimes \widetilde{H}_*(\Omega Y; F) \longrightarrow \widetilde{H}_*(\Omega Y; F)$$

where the second map is Pontryagin product, is an isomorphism. Now to finish the calculation of $H_*(\Omega Y; F)$ as the tensor algebra $T\widetilde{H}_*(X; F)$, we apply the following algebraic lemma, with $A = H_*(\Omega Y; F)$, $V = \widetilde{H}_*(X; F)$, and $i = \lambda_*$.

Lemma 4.78. *Let A be a graded algebra over a field F with $A_0 = F$ and let V be a graded vector space over F with $V_0 = 0$. Suppose we have a linear map $i : V \to A$ preserving grading, such that the multiplication map $\mu : A \otimes V \to \widetilde{A}$, $\mu(a \otimes v) = ai(v)$, is an isomorphism. Then the canonical algebra homomorphism $i : TV \to A$ extending the previous i is an isomorphism.*

For example, if V is a one-dimensional vector space over F, as in the case $X = S^n$, then this says that if the map $A \to \widetilde{A}$ given by right-multiplication by an element $a = i(v)$ is an isomorphism, then A is the polynomial algebra $F[a]$. The general case can be

viewed as the natural generalization of this to polynomials in any number of noncommuting variables.

Proof: Since μ is an isomorphism, each element $a \in A_n$ with $n > 0$ can be written uniquely in the form $\mu(\sum_j a_j \otimes v_j) = \sum_j a_j i(v_j)$ for $v_j \in V$ and $a_j \in A_{n(j)}$, with $n(j) < n$ since $V_0 = 0$. By induction on n, each a_j is $i(\alpha_j)$ for a unique $\alpha_j \in (TV)_{n(j)}$. Thus $a = i(\sum_j \alpha_j \otimes v_j)$ so i is surjective, and since these representations are unique, i is also injective. The induction starts with the hypothesis that $A_0 = F$, the scalars in TV. □

Returning now to the proof of the theorem, we observe that the map λ is an H-map, i.e., the two maps $J(X) \times J(X) \to \Omega\Sigma X$, $(x, y) \mapsto \lambda(xy)$ and $(x, y) \mapsto \lambda(x)\lambda(y)$, are homotopic since the loops $\lambda(xy)$ and $\lambda(x)\lambda(y)$ differ only in their parametrizations. Since λ is an H-map, the maps $X \hookrightarrow J(X) \xrightarrow{\lambda} \Omega\Sigma X$ induce a commutative diagram

$$\begin{array}{ccc} & T\widetilde{H}_*(X;F) & \\ \swarrow & & \searrow \\ H_*(J(X);F) & \xrightarrow{\lambda_*} & H_*(\Omega\Sigma X;F) \end{array}$$

We have shown that the downward map on the right is an isomorphism, and the same is true of the one on the left by the calculation of $H_*(J(X);F)$ in §3.4c. Hence λ_* is an isomorphism. By Corollary 3.34 this is also true for $\mathbb{Z}$ coefficients. When X is simply-connected, so are $J(X)$ and $\Omega\Sigma X$, so after taking a CW approximation to $\Omega\Sigma X$, Whitehead's theorem implies that λ is a weak homotopy equivalence. In the general case we obtain the same conclusion by applying the generalization of Whitehead's theorem to abelian spaces, Proposition 4.48, since $J(X)$ and $\Omega\Sigma X$ are H-spaces, hence have trivial action of π_1 on all homotopy groups. □

Using the natural identification $\pi_i(\Omega\Sigma X) = \pi_{i+1}(\Sigma X)$, the inclusion $X \hookrightarrow \Omega\Sigma X$ induces the suspension map $\pi_i(X) \to \pi_{i+1}(\Sigma X)$. Since this inclusion factors through $J(X)$, we can identify the relative groups $\pi_i(\Omega\Sigma X, X)$ with $\pi_i(J(X), X)$. If X is n-connected then the pair $(J(X), X)$ is $(2n+1)$-connected since we can replace X by a complex with n-skeleton a point, and then the $(2n+1)$-skeleton of $J(X)$ is contained in X. Thus we have:

Corollary 4.79. *If X is an n-connected CW complex, the suspension map $\pi_i(X) \to \pi_{i+1}(\Sigma X)$ is an isomorphism for $i \leq 2n$ and a surjection for $i = 2n+1$.* □

In the case of a sphere we can describe what happens in the lowest dimension when suspension is not an isomorphism, namely the suspension $\pi_{2n-1}(S^n) \to \pi_{2n}(S^{n+1})$ which

the corollary guarantees only to be a surjection. The CW structure on $J(S^n)$ consists of a single cell in each dimension a multiple of n, so from exactness of $\pi_{2n}(J(S^n), S^n) \xrightarrow{\partial} \pi_{2n-1}(S^n) \xrightarrow{\Sigma} \pi_{2n}(S^{n+1})$ we see that the kernel of the suspension $\pi_{2n-1}(S^n) \to \pi_{2n}(S^{n+1})$ is generated by the attaching map of the $2n$-cell of $J(S^n)$. This attaching map is the Whitehead product $[\mathbb{1}, \mathbb{1}]$, as we noted in §4.3 when we defined Whitehead products. When n is even, the Hopf invariant homomorphism $\pi_{2n-1}(S^n) \to \mathbb{Z}$ has the value ± 2 on $[\mathbb{1}, \mathbb{1}]$, as we saw in §4.4b. If there is no map of Hopf invariant ± 1, it follows that $[\mathbb{1}, \mathbb{1}]$ generates a $\mathbb{Z}$ summand of $\pi_{2n-1}(S^n)$, and so the suspension homomorphism simply cancels this summand from $\pi_{2n-1}(S^n)$. By Adams' theorem, this is the situation for all even n except 2, 4, and 8.

When $n = 2$ we have $\pi_3(S^2) \approx \mathbb{Z}$ generated by the Hopf map η with Hopf invariant 1, so $2\eta = \pm[\mathbb{1}, \mathbb{1}]$, generating the kernel of the suspension, hence:

Corollary 4.80. *$\pi_{n+1}(S^n)$ is $\mathbb{Z}_2$ for $n \geq 3$, generated by the suspension or iterated suspension of the Hopf map.* □

The situation for $n = 4$ and 8 is more subtle. We don't have the tools available here to do the actual calculation, but if we consult the table in §4.1 we see that the suspension $\pi_7(S^4) \to \pi_8(S^5)$ is a map $\mathbb{Z} \oplus \mathbb{Z}_{12} \to \mathbb{Z}_{24}$. By our preceding remarks we know this map is surjective with kernel generated by the single element $[\mathbb{1}, \mathbb{1}]$. Algebraically, what must be happening is that the coordinate of $[\mathbb{1}, \mathbb{1}]$ in the $\mathbb{Z}$ summand is twice a generator, while the coordinate in the $\mathbb{Z}_{12}$ summand is a generator. Thus a generator of the $\mathbb{Z}$ summand, which we may take to be the Hopf map $S^7 \to S^4$, suspends to a generator of the $\mathbb{Z}_{24}$. For $n = 8$ the situation is entirely similar, with the suspension $\pi_{15}(S^8) \to \pi_{16}(S^9)$ a homomorphism $\mathbb{Z} \oplus \mathbb{Z}_{120} \to \mathbb{Z}_{240}$.

We can also get information about suspension somewhat beyond the edge of the stable dimension range. Since S^n is $(n-1)$-connected and $(J(S^n), S^n)$ is $(2n-1)$-connected, we have isomorphisms $\pi_i(J(S^n), S^n) \approx \pi_i(J(S^n)/S^n)$ for $i \leq 3n-2$ by Proposition 4.27. The group $\pi_i(J(S^n)/S^n)$ is isomorphic to $\pi_i(S^{2n})$ in the same range $i \leq 3n-2$ since $J(S^n)/S^n$ has S^{2n} as its $(3n-1)$-skeleton. Thus the terminal portion of the long exact sequence of the pair $(J(S^n), S^n)$ starting with the term $\pi_{3n-2}(S^n)$ can be written in the form

$$\pi_{3n-2}(S^n) \xrightarrow{\Sigma} \pi_{3n-1}(S^{n+1}) \to \pi_{3n-2}(S^{2n}) \to \pi_{3n-3}(S^n) \xrightarrow{\Sigma} \pi_{3n-2}(S^{n+1}) \to \cdots$$

This is known as the **EHP sequence** since its three maps were originally called E, H, and P. (The German word for 'suspension' begins with 'E,' the 'H' refers to a generalization

of the Hopf invariant, and the 'P' denotes a connection with Whitehead products; see [G.Whitehead] for more details.) Note that the terms $\pi_i(S^{2n})$ in the EHP sequence are stable homotopy groups since $i \leq 3n-2$. Thus we have the curious situation that stable homotopy groups are measuring the instability of the groups $\pi_i(S^n)$ in the range $2i-1 \leq i \leq 3n-2$, the so-called *metastable* range.

Specializing to the first interesting case $n = 2$, the sequence becomes

$$\begin{array}{ccccccccc} \pi_4(S^2) & \xrightarrow{\Sigma} & \pi_5(S^3) & \longrightarrow & \pi_4(S^4) & \longrightarrow & \pi_3(S^2) & \xrightarrow{\Sigma} & \pi_4(S^3) & \longrightarrow & 0 \\ \wr & & & & \wr & & \wr & & \wr & & \\ \mathbb{Z}_2 & & & & \mathbb{Z} & & \mathbb{Z} & & \mathbb{Z}_2 & & \end{array}$$

From the Hopf bundle $S^1 \to S^3 \to S^2$ we have $\pi_4(S^2) \approx \pi_4(S^3) \approx \mathbb{Z}_2$, with $\pi_4(S^2)$ generated by the composition $(\Sigma\eta)\eta$ where η is the Hopf map $S^3 \to S^2$. From exactness of the latter part of the sequence we deduce that the map $\pi_4(S^4) \to \pi_3(S^2)$ is injective, hence that the suspension $\pi_4(S^2) \to \pi_5(S^3)$ is surjective, so $\pi_5(S^3)$ is either $\mathbb{Z}_2$ or 0. From the general suspension theorem, the suspension $\pi_5(S^3) \to \pi_6(S^4)$ is surjective as well, and the latter group is in the stable range. We show in Proposition 4.91 that the stable group π_2^s is nonzero, and so we conclude that $\pi_{n+2}(S^n) \approx \mathbb{Z}_2$ for all $n \geq 2$, generated by the composition $(\Sigma^{n-1}\eta)(\Sigma^{n-2}\eta)$.

We will see in Volume II that the EHP sequence extends all the way to the left to form an infinite exact sequence when n is odd, and when n is even a weaker statement holds: The sequence extends after factoring out all odd torsion.

Replacing S^n by any $(n-1)$-connected CW complex X, our derivation of the finite EHP sequence generalizes immediately to give an exact sequence

$$\pi_{3n-2}(X) \xrightarrow{\Sigma} \pi_{3n-1}(\Sigma X) \to \pi_{3n-2}(X \wedge X) \to \pi_{3n-3}(X) \xrightarrow{\Sigma} \pi_{3n-2}(\Sigma X) \to \cdots$$

using the fact that $J_2(X)/X = X \wedge X$.

The generalization of the results of this section to $\Omega^n\Sigma^n X$ turns out to be of some importance in homotopy theory. In case we don't get to this topic in Volume II, the reader can begin to learn about it by looking at [Carlsson-Milgram].

j. Symmetric Products and the Dold-Thom Theorem

In the preceding section we studied the free monoid $J(X)$ generated by a space X, and in this section we take up its commutative analog, the free abelian monoid generated by X. This is the infinite symmetric product $SP(X)$ introduced briefly in §3.4c. The main result will be a theorem of [Dold-Thom] asserting that $\pi_* SP(X) \approx \widetilde{H}_*(X;\mathbb{Z})$ for all connected CW complexes X. In particular this yields the quite striking fact that $SP(S^n)$ is a $K(\mathbb{Z},n)$, and more generally that the functor SP takes Moore spaces $M(G,n)$ to Eilenberg-MacLane spaces $K(G,n)$. This leads to the general result that for all connected CW complexes X, $SP(X)$ has the homotopy type of a product of Eilenberg-MacLane spaces, or in other words that the k-invariants of $SP(X)$ are all trivial.

The Dold-Thom theorem is proved by showing that the homotopy groups $\pi_* SP(X)$ define a homology theory, then deducing that this must be ordinary homology with $\mathbb{Z}$ coefficients by computing the coefficient groups $\pi_* SP(S^n)$. A key new ingredient in the verification of the axioms for a homology theory is the notion of a quasifibration, which is useful in other contexts as well. A crucial property of quasifibrations which we need will follow from a certain fact about relative homotopy groups. Though this fact is essentially elementary, we have not encountered it previously in the book, so let us begin with this elementary topic.

A Mayer-Vietoris Property of Homotopy Groups

Throughout this section relative homotopy groups will be ubiquitous, and it will be impossible to avoid the awkward fact that there is not a really good way to define the relative π_0. What we will do as a compromise is to take $\pi_0(X,A,x_0)$ to be the quotient set $\pi_0(X,x_0)/\pi_0(A,x_0)$. This at least allows the long exact sequence of homotopy groups for (X,A) to end with the terms $\pi_0(A,x_0) \to \pi_0(X,x_0) \to \pi_0(X,A,x_0) \to 0$. An exercise for §4.1 shows that the five-lemma can be applied to the map of long exact sequences induced by a map $(X,A) \to (Y,B)$, provided the basepoint is allowed to vary. However, the long exact sequence of a triple cannot be extended through the π_0 terms with this definition, so one must proceed with some caution.

The excision theorem for homology involves a space X with subspaces A and B such that X is the union of the interiors of A and B. In this situation we call $(X;A,B)$ an **excisive triad**. By a map $f:(X;A,B) \to (Y;C,D)$ we mean $f:X \to Y$ with $f(A) \subset C$ and $f(B) \subset D$.

Proposition 4.81. *Given a map of excisive triads $f:(X;A,B) \to (Y;C,D)$, if the induced maps $\pi_i(A, A\cap B) \to \pi_i(C, C\cap D)$ and $\pi_i(B, A\cap B) \to \pi_i(D, C\cap D)$ are bijections*

for $i < n$ and surjections for $i = n$, for all choices of basepoints, then the same is true of the induced maps $\pi_i(X, A) \to \pi_i(Y, C)$. By symmetry the conclusion holds also for the maps $\pi_i(X, B) \to \pi_i(Y, D)$.

The corresponding statement for homology is a trivial consequence of excision which says simply that $H_i(X, A) \approx H_i(B, A \cap B)$ and $H_i(Y, C) \approx H_i(D, C \cap D)$, so it is not necessary to assume anything about $H_i(A, A \cap B)$ and $H_i(C, C \cap D)$. With the failure of excision for homotopy groups, however, it is not surprising that the assumption on $\pi_i(A, A \cap B) \to \pi_i(C, C \cap D)$ cannot be dropped. An example is provided by the quotient map $f : D^2 \to S^2$ collapsing ∂D^2 to the north pole of S^2, with C and D the northern and southern hemispheres of S^2, and A and B their pre-images under f.

Proof: We may assume f is an inclusion by replacing Y by the mapping cylinder of f, with C and D then being replaced by the mapping cylinders of the restrictions of f to A and B.

Consider an inclusion $(X, A) \subset (Y, C)$. As a substitute for a hypothetical long exact sequence $\cdots \to \pi_i(X, A) \to \pi_i(Y, C) \to \pi_i((Y, C), (X, A)) \to \cdots$ we will show the following three conditions are equivalent for each $n \geq 1$:

(*i*) For all choices of basepoints the map $\pi_i(X, A) \to \pi_i(Y, C)$ induced by inclusion is surjective for $i = n$ and has trivial kernel for $i = n - 1$.

(*ii*) Let ∂D^n be written as the union of hemispheres D^{n-1}_+ and D^{n-1}_- intersecting in S^{n-2}. Then every map $(D^n \times \{0\} \cup D^{n-1}_+ \times I, D^{n-1}_- \times \{0\} \cup S^{n-2} \times I) \to (Y, C)$ with $(D^{n-1}_+ \times \{1\}, S^{n-2} \times \{1\}) \to (X, A)$ extends to a map $(D^n \times I, D^{n-1}_- \times I) \to (Y, C)$ with $(D^n \times \{1\}, D^{n-1}_- \times \{1\}) \to (X, A)$.

(*iii*) Condition (*ii*) with the added hypothesis that the restriction of the given map to $D^{n-1}_+ \times I$ is a constant homotopy.

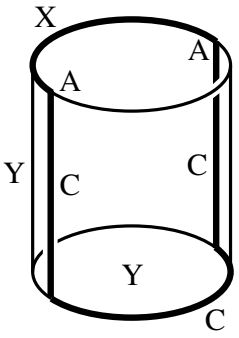

It is obvious that (*ii*) and (*iii*) are equivalent since the stronger hypothesis in (*iii*) can always be achieved by composing with a homotopy of $D^n \times I$ which shrinks $D^{n-1}_+ \times I$ to $D^{n-1}_+ \times \{1\}$.

To see that (*iii*) implies (*i*), let $f : (D^{n-1}_+ \times \{1\}, S^{n-2} \times \{1\}) \to (X, A)$ represent an element of $\pi_{n-1}(X, A)$. If this is in the kernel of the map to $\pi_{n-1}(Y, C)$, then we get an extension of f over $D^n \times \{0\} \cup D^{n-1}_+ \times I$, with the constant homotopy on $D^{n-1}_+ \times I$ and $(D^n \times \{0\}, D^{n-1}_- \times \{0\})$ mapping to (Y, C). Condition (*iii*) then gives an extension over $D^n \times I$, whose restriction to $D^n \times \{1\}$ shows that f is zero in $\pi_{n-1}(X, A)$, so the kernel of $\pi_{n-1}(X, A) \to \pi_{n-1}(Y, C)$ is trivial. To check surjectivity of the map $\pi_n(X, A) \to \pi_n(Y, C)$, represent an element of $\pi_n(Y, C)$ by a map $f : D^n \times \{0\} \to Y$ taking $D^{n-1}_- \times \{0\}$

to C and $D^{n-1}_+ \times \{0\}$ to a chosen basepoint. Extend f over $D^{n-1}_+ \times I$ via the constant homotopy, then extend over $D^n \times I$ by applying (*iii*). The result is a homotopy of the given f to a map representing an element of the image of $\pi_n(X, A) \to \pi_n(Y, C)$.

Now we show that (*i*) implies (*ii*). Given a map f as in the hypothesis of (*ii*), the injectivity part of (*i*) gives an extension of f over $D^n \times \{1\}$. Choose a small disk $E^n \subset D^{n-1}_- \times I$ intersecting $D^{n-1}_- \times \{1\}$ in a hemisphere E^{n-1}_+ of its boundary. We may assume the extended f has a constant value $x_0 \in A$ on E^{n-1}_+. Viewing the extended f as representing an element of $\pi_n(Y, C, x_0)$, the surjectivity part of (*i*) then gives an extension of f over $D^n \times I$ taking (E^n, E^{n-1}_-) to (X, A) and the rest of $D^{n-1}_- \times I$ to C. The argument is finished by composing this extended f with a deformation of $D^n \times I$ pushing E^n into $D^n \times \{1\}$.

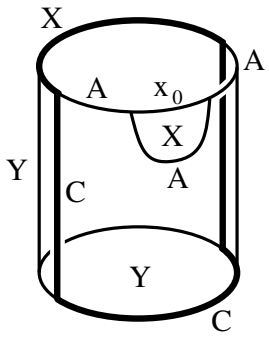

Having shown the equivalence of (*i*)-(*iii*), let us prove the proposition. The case $n = 0$ being trivial from the definitions, let us assume $n \geq 1$. After what we have shown so far, it suffices to show that condition (*ii*) for the inclusions $(A, A \cap B) \hookrightarrow (C, C \cap D)$ and $(B, A \cap B) \hookrightarrow (D, C \cap D)$ implies (*iii*) for the inclusion $(X, A) \hookrightarrow (Y, C)$. Let a map $f : D^n \times \{0\} \cup D^{n-1}_+ \times I \to Y$ as in the hypothesis of (*iii*) be given. The argument will involve subdivision of D^n into smaller disks, and for this it is more convenient to use the cube I^n instead of D^n, so let us identify I^n with D^n in such a way that D^{n-1}_- corresponds to the face $I^{n-1} \times \{1\}$ and D^{n-1}_+ corresponds to the rest of ∂I^n, which we denote by $J(I^n)$. Thus we are given f on $I^n \times \{0\}$ taking $J(I^n) \times \{0\}$ to X and $I^{n-1} \times \{1\} \times \{0\}$ to C, and on $J(I^n) \times I$ we have the constant homotopy.

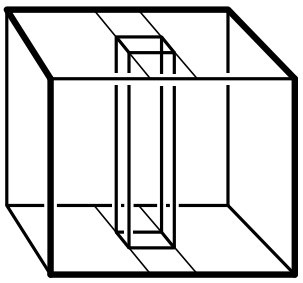

Since $(Y; C, D)$ is an excisive triad, we can subdivide each of the I factors of $I^n \times \{0\}$ into subintervals so that f takes each of the resulting n-dimensional subcubes σ of $I^n \times \{0\}$ into either C or D. Consider the problem of extending f over $\sigma \times I$ for σ one of the subcubes. The extension we construct will have the property that it takes $(\sigma \times I, \sigma \times \{1\})$ to (C, A) or (D, B) according to whether the given f takes $\sigma \times \{0\}$ to C or D. We may assume that f has already been extended to $J(\sigma) \times I$, by induction on n and on the sequence of subintervals of the last coordinate of $I^n \times \{0\}$. To extend f over $\sigma \times I$ there are now four subcases according to whether the given f takes $(\sigma, \partial\sigma - J(\sigma)) \times \{0\}$ to $(C, C \cap D)$, $(D, C \cap D)$, (C, C), or (D, D). In the first two cases we apply (*ii*) for the inclusion $(A, A \cap B) \hookrightarrow (C, C \cap D)$ or $(B, A \cap B) \hookrightarrow (D, C \cap D)$. In the last two cases we can apply (*ii*) trivially. □

Corollary 4.82. *Suppose given a map* $f:X\to Y$ *and open covers* $\{U_i\}$ *of* X *and* $\{V_i\}$ *of* Y *with* $f(U_i)\subset V_i$ *for all* i. *If each restriction* $f:U_i\to V_i$ *and more generally each* $f:U_{i_1}\cap\cdots\cap U_{i_n}\to V_{i_1}\cap\cdots\cap V_{i_n}$ *is a weak homotopy equivalence, then so is* $f:X\to Y$.

Proof: First let us do the case of finite covers by induction on the number of sets in the cover, starting with the case of covers by two sets. By the five-lemma, the hypotheses imply that $\pi_n(U_i,U_1\cap U_2)\to\pi_n(V_i,V_1\cap V_2)$ is bijective for $i=1,2$, $n\geq 0$, and all choices of basepoints. The preceding proposition then implies that the maps $\pi_n(X,U_1)\to\pi_n(Y,V_1)$ are isomorphisms. Hence by the five-lemma, so are the maps $\pi_n(X)\to\pi_n(Y)$.

The inductive step for finite covers by $k>2$ sets reduces to the case of covers by two sets by letting one of the two sets be the union of the first $k-1$ of the given sets and the other be the k^{th} set. The case of infinite covers reduces to the finite case since for surjectivity of $\pi_n(X)\to\pi_n(Y)$, a map $S^n\to Y$ has compact image covered by finitely many V_i's, and similarly for injectivity. □

Symmetric Products

Let us recall the definition from §3.4c. For a space X the n-fold symmetric product $SP_n(X)$ is the quotient space of the product of n copies of X obtained by factoring out the action of the symmetric group permuting the factors. If we choose a basepoint $e\in X$, then there are inclusions $SP_n(X)\hookrightarrow SP_{n+1}(X)$ induced by $(x_1,\cdots,x_n)\mapsto(x_1,\cdots,x_n,e)$, and $SP(X)$ is defined to be the union of this increasing sequence of spaces, with the direct limit topology. Note that SP_n is a homotopy functor: A map $f:X\to Y$ induces $f_*:SP_n(X)\to SP_n(Y)$, and $f\simeq g$ implies $f_*\simeq g_*$. Hence $X\simeq Y$ implies $SP_n(X)\simeq SP_n(Y)$. In similar fashion SP is a homotopy functor on the category of basepointed spaces and basepoint-preserving homotopy classes of maps. It follows that $X\simeq Y$ implies $SP(X)\simeq SP(Y)$ for connected CW complexes X and Y since in this case requiring maps and homotopies to preserve basepoints does not affect the relation of homotopy equivalence.

We saw in §3.4c that $SP_n(S^2)$ is homeomorphic to $\mathbb{C}\mathrm{P}^n$ and $SP(S^2)$ is homeomorphic to $\mathbb{C}\mathrm{P}^\infty$. The same argument can be used to show that $SP_n(S^1)\simeq S^1\simeq SP(S^1)$. Namely, the argument shows that $SP_n(\mathbb{C}-\{0\})$ can be identified with the polynomials $z^n+a_{n-1}z^{n-1}+,,,+a_0$ with $a_0\neq 0$, or in other words with n-tuples $(a_0,\cdots,a_{n-1})\in\mathbb{C}^n$ with $a_0\neq 0$, and this subspace of $\mathbb{C}^n$ deformation retracts onto a circle.

The symmetric products of higher-dimensional spheres are more complicated, though things are not so bad for the 2-fold symmetric product:

Example. Let us show that $SP_2(S^n)$ is homeomorphic to the mapping cone of a map $\Sigma^n\mathbb{R}\mathrm{P}^{n-1}\to S^n$. In particular, this gives $H_*(SP_2(S^n))\approx H_*(S^n)\oplus\widetilde{H}_*(\Sigma^{n+1}\mathbb{R}\mathrm{P}^{n-1})$ by

looking at the long exact sequence of homology groups for the pair $(SP_2(S^n), S^n)$, since $SP_2(S^n)/S^n$ is $\Sigma^{n+1}\mathbb{R}\mathrm{P}^{n-1}$ with no reduced homology below dimension $n+2$.

If we view S^n as $D^n/\partial D^n$, then $SP_2(S^n)$ becomes a certain quotient of $D^n \times D^n$. Viewing $D^n \times D^n$ as the cone on its boundary $D^n \times \partial D^n \cup \partial D^n \times D^n$, the identifications which produce $SP_2(S^n)$ respect the various concentric copies of this boundary which fill up the interior of $D^n \times D^n$, so it suffices to analyze the identifications in all these copies of the boundary. The identifications on the boundary of $D^n \times D^n$ itself yield S^n. This is clear since the identification $(x, y) \sim (y, x)$ converts $D^n \times \partial D^n \cup \partial D^n \times D^n$ to $D^n \times \partial D^n$, and all points of ∂D^n are identified in S^n.

It remains to see that the identifications on each concentric copy of the boundary in the interior of $D^n \times D^n$ produce $\Sigma^n\mathbb{R}\mathrm{P}^{n-1}$. These identifications are just $(x, y) \sim (y, x)$. Denote by Z the quotient of $D^n \times \partial D^n \cup \partial D^n \times D^n$ under these identifications. This is the same as the quotient of $D^n \times \partial D^n$ under the identifications $(x, y) \sim (y, x)$ for $(x, y) \in \partial D^n \times \partial D^n$.

Define a map $f : D^n \times \mathbb{R}\mathrm{P}^{n-1} \to Z$ by $f(x, L) = (w, z)$ where w and z are defined as in the figure, that is, x is equidistant from $z \in \partial D^n$ and $w \in D^n$ along the line through x parallel to L. Note that if x happens to be the midpoint of the segment zz', then $w = z'$ and there is no way to distinguish between w and z, but since f takes values in the quotient space Z, this is no problem. If $x \in \partial D^n$ then $w = z = x$, independent of L. If $x \in D^n - \partial D^n$ then $w \neq z$, and conversely, given $(w, z) \in D^n \times \partial D^n$ with $w \neq z$ there is a unique (x, L) with $f(x, L) = (w, z)$, namely x is the midpoint of the segment wz and L is the line parallel to this segment. In view of these remarks, we see that Z is the quotient space of $D^n \times \mathbb{R}\mathrm{P}^{n-1}$ under the identifications $(x, L) \sim (x, L')$ if $x \in \partial D^n$. This quotient is precisely $\Sigma^n\mathbb{R}\mathrm{P}^{n-1}$.

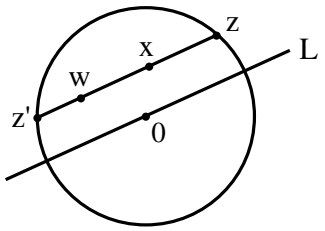

This example illustrates that passing from a CW structure on X to a CW structure on $SP_n(X)$ or $SP(X)$ is not at all straightforward. However, if X is a simplicial complex there is a natural way of putting Δ-complex structures on $SP_n(X)$ and $SP(X)$, as follows. A simplicial complex structure on X gives a CW structure on the product of n copies of X, with cells n-fold products of simplices. Such a product has a canonical barycentric subdivision as a simplicial complex, with vertices the points whose coordinates are barycenters of simplices of X. By induction over skeleta, this just amounts to coning off a simplicial structure on the boundary of each product cell. This simplicial structure on the product of n copies of X is in fact a Δ-complex structure since the vertices of each

of its simplices have a natural ordering given by the dimensions of the cells of which they are barycenters. The action of the symmetric group permuting coordinates respects this Δ-complex structure, taking simplices homeomorphically to simplices, preserving vertex-orderings, so there is an induced Δ-complex structure on the quotient $SP_n(X)$. With the basepoint of X chosen to be a vertex, $SP_n(X)$ is a subcomplex of $SP_{n+1}(X)$ so there is a natural Δ-complex structure on the infinite symmetric product $SP(X)$ as well.

As usual with products, the CW topology on $SP_n(X)$ and $SP(X)$ is in general different from the topology arising from the original definition in terms of product topologies, but one can check that the two topologies have the same compact sets, so the distinction won't matter for our present purposes. For definiteness, we will use the CW topology in what follows, which means restricting X to be a simplicial complex. Since every CW complex is homotopy equivalent to a simplicial complex by Theorem 2.37, and SP_n and SP are homotopy functors, there is no essential loss of generality in restricting from CW complexes to simplicial complexes.

Here is the main result of this section:

Theorem 4.83. *The functor $X \mapsto \pi_i SP(X)$ for $i \geq 1$ coincides with the functor $X \mapsto H_i(X;\mathbb{Z})$ on the category of basepointed connected CW complexes.*

In particular this says that $SP(S^n)$ is a $K(\mathbb{Z}, n)$, and more generally that for a Moore space $M(G,n)$, $SP(M(G,n))$ is a $K(G,n)$.

Now, $SP(X)$ is a commutative, associative H-space with a strict identity element, and this limits its weak homotopy type considerably, at least when X is path-connected:

Corollary 4.84. *A path-connected, commutative, associative H-space with a strict identity element has the weak homotopy type of a product of Eilenberg-MacLane spaces.*

In particular, if X is a connected CW complex then $SP(X)$ is path-connected and has the weak homotopy type of $\prod_n K(H_n(X), n)$. Thus the functor SP essentially reduces to Eilenberg-MacLane spaces.

Proof: Let X be a path-connected, commutative, associative H-space with a strict identity element, and let $G_n = \pi_n(X)$. By Lemma 4.16 there is a map $M(G_n, n) \to X$ inducing an isomorphism on π_n when $n > 1$ and an isomorphism on H_1 when $n = 1$. We can take these maps to be basepoint-preserving, and then they combine to give a map $\bigvee_n M(G_n, n) \to X$. The very nice H-space structure on X allows us to extend this to a homomorphism $f : SP(\bigvee_n M(G_n, n)) \to X$. In general, $SP(\bigvee_\alpha X_\alpha) = \prod_\alpha SP(X_\alpha)$ where this is the 'weak' infinite product, the union of the finite products. This, together with the general fact that the map $\pi_i(Z) \to \pi_i SP(Z) = H_i(Z;\mathbb{Z})$ induced by the inclusion

$Z = SP_1(Z) \hookrightarrow SP(Z)$ is the Hurewicz homomorphism, as we will see at the end of the proof of the theorem, implies that the map f induces an isomorphism on all homotopy groups. Thus we have a weak homotopy equivalence $\prod_n SP(M(G_n, n)) \to X$, and as we noted above, $SP(M(G_n, n))$ is a $K(G_n, n)$. Finally, since each factor $SP(M(G_n, n))$ has only one nontrivial homotopy group, the weak infinite product has the same weak homotopy type as the ordinary infinite product. □

The main step in the proof of the theorem will be to show that for a simplicial pair (X, A) with both X and A connected, there is a long exact sequence $\cdots \to \pi_i SP(A) \to \pi_i SP(X) \to \pi_i SP(X/A) \to \pi_{i-1} SP(A) \to \cdots$. This would follow if the maps $SP(A) \to SP(X) \to SP(X/A)$ formed a fiber bundle or fibration. There is some reason to hope this would be true, because all the fibers of the projection $SP(X) \to SP(X/A)$ are homeomorphic to $SP(A)$. In fact, in terms of the H-space structure on $SP(X)$ as the free abelian monoid generated by X, the fibers are exactly the cosets of the submonoid $SP(A)$. However, the projection $SP(X) \to SP(X/A)$ fails to have the homotopy lifting property, even the special case of lifting paths. For if we take a path in $X - A$ whose endpoint is a point $a \in A$ other than the basepoint, then regarding this as a path in $SP(X/A)$, any lift to $SP(X)$ would have to end at a point of $SP(A)$ which was a multiple of a, so in particular there would be no lift ending at the basepoint of $SP(X)$.

A simpler example to have in mind is the natural projection $M_f \to I$ of the mapping cylinder of a map $f : Z \to Z$. Here all the fibers are copies of Z, but the homotopy lifting property fails for paths if f is not surjective. For example, one can take f to be a proper inclusion $I \hookrightarrow I$.

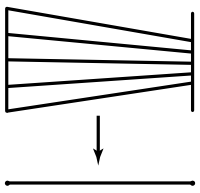

What we will show is that the projection $SP(X) \to SP(X/A)$ has instead the weaker structure of a 'quasifibration' which is still good enough to deduce a long exact sequence of homotopy groups. There are two ways to define a **quasifibration** $p : E \to B$ over a path-connected base space B. The first way is to say that the induced map $p_* : \pi_i(E, p^{-1}(b), x_0) \to \pi_i(B, b)$ is an isomorphism for all $b \in B$, $x_0 \in p^{-1}(b)$, and $i \geq 0$. We have shown in §4.3 that fiber bundles and fibrations have this property for $i > 0$, as a consequence of the homotopy lifting property, and the same reasoning applies for $i = 0$ since we assume B is path-connected.

The alternative definition is to say that the inclusion of each fiber $p^{-1}(b)$ into the homotopy fiber F_b of p over b is a weak homotopy equivalence. Recall that this homotopy fiber is the space of all pairs consisting of a point $x \in E$ and a path γ in B from $p(x)$ to b. The actual fiber $p^{-1}(b)$ is included in F_b as the pairs with γ the constant path at

a point $x \in p^{-1}(b)$. One can see the equivalence of the two definitions by considering the commutative triangle

$$\begin{array}{ccc} \pi_i(E, p^{-1}(b)) & \longrightarrow & \pi_i(E_p, F_b) \\ & \searrow \quad \swarrow & \\ & \pi_i(B, b) & \end{array}$$

where $F_b \to E_p \to B$ is the usual path-fibration construction applied to p. The right-hand map in the diagram is an isomorphism for all $i \geq 0$, and the upper map will be an isomorphism for all $i \geq 0$ iff the inclusion $p^{-1}(b) \hookrightarrow F_b$ is a weak equivalence since $E \simeq E_p$. Hence the two definitions are equivalent. In particular, since all fibers of a fibration over a path-connected base are homotopy equivalent, we see that all the fibers of a quasifibration have the same weak homotopy type.

The more general case of quasifibrations over a base B which is not path-connected is considered in the exercises, but we will not need this generality in what follows.

Returning to the example of the projection $p : M_f \to I$ of the mapping cylinder of a map $f : Z \to Z$, this projection will be a quasifibration if f is a homotopy equivalence. For in this case M_f deformation retracts onto each fiber $p^{-1}(b)$ by Corollary 0.6, hence $\pi_i(M_f, p^{-1}(b)) = 0 = \pi_i(I, b)$ for all i.

A key fact we shall need about quasifibrations is the following lemma, whose proof will be a straightforward application of the results at the beginning of this section.

Lemma 4.85. *A map $p : E \to B$ is a quasifibration if B can be decomposed as the union of open sets V_1 and V_2 such that each of the maps $p^{-1}(V_1) \to V_1$, $p^{-1}(V_2) \to V_2$, and $p^{-1}(V_1 \cap V_2) \to V_1 \cap V_2$ is a quasifibration.*

Proof: To avoid some tedious details we will consider only the case that the fibers of p are path-connected, which will suffice for our present purposes, leaving the general case as an exercise for the reader. This hypothesis on fibers guarantees that all π_0's arising in the proof are trivial. In particular, by an exercise for §4.1 this allows us to terminate long exact sequences of homotopy groups of triples with zeros in the π_0 positions.

Let $U_1 = p^{-1}(V_1)$ and $U_2 = p^{-1}(V_2)$. By applying the five lemma to the long exact sequences of homotopy groups of the triples $(U_k, U_1 \cap U_2, p^{-1}(b))$ and $(V_k, V_1 \cap V_2, b)$ we deduce that the maps $\pi_i(U_k, U_1 \cap U_2) \to \pi_i(V_k, V_1 \cap V_2)$ are isomorphisms for $k = 1, 2$ and all i. Then Proposition 4.81 implies that the maps $\pi_i(E, U_1) \to \pi_i(B, V_1)$ are isomorphisms. The five-lemma and the hypothesis that the maps $\pi_i(U_1, p^{-1}(b)) \to \pi_i(V_1, b)$ are isomorphisms then show that the maps $\pi_i(E, p^{-1}(b)) \to \pi_i(B, b)$ are isomorphisms. □

Proof of 4.83: As we have said, the main step will be to associate a long exact sequence

of homotopy groups to each simplicial pair (X, A) with X and A connected. This will be the long exact sequence of homotopy groups coming from the quasifibration $SP(A) \to SP(X) \to SP(X/A)$, so the major work will be in verifying the quasifibration property. Since SP is a homotopy functor, we are free to replace (X, A) by a homotopy equivalent pair, so let us replace (X, A) by (M, A) where M is the mapping cylinder of the inclusion $A \hookrightarrow X$. This new pair, which we still call (X, A), has some slight technical advantages, as we will see later in the proof.

To begin the proof that the projection $p: SP(X) \to SP(X/A)$ is a quasifibration, let $B_n = SP_n(X/A)$ and $E_n = p^{-1}(B_n)$. Thus E_n consists of those points in $SP(X)$ having at most n coordinates in $X - A$. The following general fact about quasifibrations reduces us to showing that $p: E_n \to B_n$ is a quasifibration:

(*i*) A map $p: E \to B$ is a quasifibration if B is the union of an increasing sequence of subspaces $B_1 \subset B_2 \subset \cdots$, with the direct limit topology, such that each restriction $E_n = p^{-1}(B_n) \to B_n$ is a quasifibration.

The proof of this is the usual compactness argument for direct limits. Namely, to show that $p_*: \pi_i(E, p^{-1}(b)) \to \pi_i(B, b)$ is surjective, represent an element of $\pi_i(B, b)$ by a map $f: (D^i, \partial D^i) \to (B, b)$. The compact image of f lies in some B_n, and since $p: E_n \to B_n$, f is homotopic to the image under p of a map to E_n. Similarly for injectivity, if the composition of a map $f: (D^i, \partial D^i) \to (E, p^{-1}(b))$ with p is nullhomotopic, such a nullhomotopy has compact image, hence lies in some B_n, so since $E_n \to B_n$ is a quasifibration, f is trivial in $\pi_i(E_n, p^{-1}(b))$ hence in $\pi_i(E, p^{-1}(b))$.

The proof that the map $p: p^{-1}(SP_n(X/A)) \to SP_n(X/A)$, which we are abbreviating to $p: E_n \to B_n$, is a quasifibration will be by induction on n, starting with the trivial case $n = 0$ when B_0 is a point. The induction step will consist of showing that p is a quasifibration over a neighborhood of B_{n-1} and over $B_n - B_{n-1}$, then applying the preceding lemma. We first tackle the problem of showing the quasifibration property over a neighborhood of B_{n-1}.

Let $f_t: X \to X$ be a homotopy of the identity map which deformation retracts a neighborhood N of A onto A. Since we have replaced the original X by the mapping cylinder of the inclusion $A \hookrightarrow X$, we can take f_t simply to slide points along the segments $\{a\} \times I$ in the mapping cylinder, with $N = A \times [0, 1/2)$. Let $U \subset E_n$ consist of those points having at least one coordinate in N, or in other words, products with at least one factor in N. Thus U is a neighborhood of E_{n-1} in E_n and $p(U)$ is a neighborhood of B_{n-1} in B_n.

The homotopy f_t induces a homotopy $F_t: E_n \to E_n$ whose restriction to U is a

deformation of U onto E_{n-1}, where by 'deformation' we mean 'deformation retraction in the weak sense.' Since f_t is the identity on A, F_t is the lift of a homotopy $\overline{F}_t : B_n \to B_n$ which restricts to a deformation of $\overline{U} = p(U)$ onto B_{n-1}. We will deduce that the projection $U \to \overline{U}$ is a quasifibration by using one more general fact about quasifibrations:

(*ii*) A map $p : E \to B$ is a quasifibration if there is a deformation F_t of E onto a subspace E', covering a deformation $\overline{F}_t$ of B onto a subspace B', such that the restriction $p : E' \to B'$ is a quasifibration and $F_1 : p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$ is a weak homotopy equivalence for all $b \in B$.

To see this, consider the commutative diagram

$$\begin{array}{ccccc} \big(E, p^{-1}(b)\big) & \xrightarrow{F_1} & \big(E', p^{-1}(\overline{F}_1(b))\big) & \longrightarrow & \big(E, p^{-1}(\overline{F}_1(b))\big) \\ \downarrow & & \downarrow & & \\ (B, b) & \xrightarrow{\overline{F}_1} & (B', \overline{F}_1(b)) & & \end{array}$$

where the upper-right map is inclusion. The composition across the top of the diagram induces an isomorphism on relative homotopy groups via the five-lemma and the hypothesis that $F_1 : p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$ is a weak homotopy equivalence. Since the inclusion $E' \hookrightarrow E$ is a homotopy equivalence, the upper-right map also induces isomorphisms on relative homotopy groups. Hence F_1 induces isomorphisms on relative homotopy groups. The lower map $\overline{F}_1$ is a homotopy equivalence, so from commutativity of the diagram we deduce that the left-hand vertical map induces isomorphisms on homotopy groups since this is true of the right-hand vertical map by assumption.

To apply this to the case at hand we need to verify the condition on $F_1 : p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$. Each point $w \in p^{-1}(b)$ is a commuting product of points in X. Let $\widehat{w}$ be the subproduct whose factors are points in $X - A$, so we have $w = \widehat{w}v$ for v a product of points in A. Since f_1 is the identity on A and F_1 is a homomorphism, we have $F_1(w) = F_1(\widehat{w})v$, which can be written $\widehat{F_1(\widehat{w})}v'v$ with v' also a product of points in A. If we fix $\widehat{w}$ and let v vary over $SP(A)$, we get all points of $p^{-1}(b)$ exactly once, or in other words, we have $p^{-1}(b)$ expressed as the coset $\widehat{w}SP(A)$. The map F_1, $\widehat{w}v \mapsto \widehat{F_1(\widehat{w})}v'v$, takes this coset to the coset $\widehat{F_1(\widehat{w})}SP(A)$ by a map which would be a homeomorphism if the factor v' were not present. Since A is connected, there is a path v'_t from v' to the basepoint, and so by replacing v' with v'_t in the product $\widehat{F_1(\widehat{w})}v'v$ we obtain a homotopy from $F_1 : p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$ to a homeomorphism, so this map is a homotopy equivalence, as desired.

It remains to see that p is a quasifibration over $B_n - B_{n-1}$ and over the intersection of this set with $\overline{U}$. The argument will be the same in both cases.

Identifying $B_n - B_{n-1}$ with $SP_n(X-A)$, the projection $p: E_n - E_{n-1} \to B_n - B_{n-1}$ is the same as the operator $w \mapsto \widehat{w}$ used earlier. The inclusion $SP_n(X-A) \hookrightarrow E_n - E_{n-1}$ provides a section for $p: E_n - E_{n-1} \to B_n - B_{n-1}$, so the induced map $p_*: \pi_i(E_n - E_{n-1}, p^{-1}(b)) \to \pi_i(B_n - B_{n-1}, b)$ is surjective. To see that it is injective, represent an element of its kernel by a map $g: (D^i, \partial D^i) \to (E_n - E_{n-1}, p^{-1}(b))$. A nullhomotopy of pg gives a homotopy of g changing only its coordinates in $X - A$. This homotopy is through maps $(D^i, \partial D^i) \to (E_n - E_{n-1}, p^{-1}(b))$, and ends with a map to $p^{-1}(b)$, so the kernel of p_* is trivial. Thus the projection $E_n - E_{n-1} \to B_n - B_{n-1}$ is a quasifibration, at least if $B_n - B_{n-1}$ is path-connected. But by replacing the original X with the mapping cylinder of the inclusion $A \hookrightarrow X$ for X, we have arranged that $X - A$ is path-connected, hence also $B_n - B_{n-1} = SP_n(X-A)$.

This argument works equally well over any path-connected open subset of $B_n - B_{n-1}$, in particular over $U \cap (B_n - B_{n-1})$, so via the lemma this finishes the proof that $SP(A) \to SP(X) \to SP(X/A)$ is a quasifibration.

This gives us the first two of the three axioms needed for the groups $h_i(X) = \pi_i SP(X)$ to define a reduced homology theory, and there remains only the wedge sum axiom, $h_i(\bigvee_\alpha X_\alpha) \approx \oplus_\alpha h_i(X_\alpha)$, but this is immediate from the evident fact that $SP(\bigvee_\alpha X_\alpha) = \prod_\alpha SP(X_\alpha)$, where this is the 'weak' product, the union of the products of finitely many factors.

The homology theory $h_*(X)$ is defined on the category of connected, basepointed simplicial complexes, with basepoint-preserving maps. The coefficients of this homology theory, the groups $h_i(S^n)$, are the same as for ordinary homology with $\mathbb{Z}$ coefficients since we know this is true for $n = 2$ by the homeomorphism $SP(S^2) \approx \mathbb{C}\mathrm{P}^\infty$, and there are isomorphisms $h_i(X) \approx h_{i+1}(\Sigma X)$ in any reduced homology theory. If the homology theory $h_*(X)$ were defined on the category of all simplicial complexes, without basepoints, then Theorem 2.29 would give natural isomorphisms $h_i(X) \approx H_i(X;\mathbb{Z})$ for all X, and the proof would be complete. However, it is easy to achieve this by defining a new homology theory $h_i'(X) = h_{i+1}(\Sigma X)$, since the suspension of an arbitrary complex is connected and the suspension of an arbitrary map is basepoint-preserving, taking the basepoint to be one of the suspension points. Since $h_i'(X)$ is naturally isomorphic to $h_i(X)$ if X is connected, we are done. □

It is worth noting that the map $\pi_i(Z) \to \pi_i SP(Z) = H_i(Z;\mathbb{Z})$ induced by the inclusion $Z = SP_1(Z) \hookrightarrow SP(Z)$ is the Hurewicz homomorphism. For by naturality this reduces to the case $Z = S^i$, where the map $SP_1(S^i) \hookrightarrow SP(S^i)$ induces on π_i a homomorphism $\mathbb{Z} \to \mathbb{Z}$, which one just has to check is an isomorphism, the Hurewicz homomorphism

being determined only up to sign. The suspension isomorphism gives a further reduction to the case $i = 1$, where the inclusion $SP_1(S^1) \hookrightarrow SP(S^1)$ is a homotopy equivalence, hence induces an isomorphism on π_1.

Exercises

1. Show that Corollary 4.82 remains valid when X and Y are CW complexes and the subspaces U_i and V_i are subcomplexes rather than open sets.

2. Show that a simplicial map $f: K \to L$ is a homotopy equivalence if $f^{-1}(x)$ is contractible for all $x \in L$. [Consider the cover of L by open stars of simplices and the cover of K by the preimages of these open stars.]

3. Show that $SP_n(I) = \Delta^n$.

4. Show that $SP_2(S^1)$ is a Möbius band, and that this is consistent with the description of $SP_2(S^n)$ as a mapping cone given in an example in this section.

5. A map $p: E \to B$ with B not necessarily path-connected is defined to be a quasifibration if the following equivalent conditions are satisfied.

(1) For all $b \in B$ and $x_0 \in p^{-1}(b)$, the map $p_*: \pi_i(E, p^{-1}(b), x_0) \to \pi_i(B, b)$ is an isomorphism for $i > 0$ and $\pi_0(p^{-1}(b), x_0) \to \pi_0(E, x_0) \to \pi_0(B, b)$ is exact.

(2) The inclusion of the fiber $p^{-1}(b)$ into the homotopy fiber F_b of p over b is a weak homotopy equivalence for all $b \in B$.

(3) The restriction of p over each path-component of B is a quasifibration according to the definition in this section.

Show these three conditions are equivalent, and for quasifibrations over non-path-connected base spaces prove Lemma 4.85 and the two facts (i) and (ii) in the proof of Theorem 4.83.

k. Steenrod Squares and Powers

The main objects of study in this section are certain homomorphisms called **Steenrod squares** and **Steenrod powers**:

$$\begin{aligned} Sq^i &: H^n(X;\mathbb{Z}_2) \to H^{n+i}(X;\mathbb{Z}_2) \\ P^i &: H^n(X;\mathbb{Z}_p) \to H^{n+2i(p-1)}(X;\mathbb{Z}_p) \quad \text{for odd primes } p \end{aligned}$$

The terms 'squares' and 'powers' arise from the fact that Sq^i and P^i are related to the maps $\alpha \mapsto \alpha^2$ and $\alpha \mapsto \alpha^p$ sending a cohomology class α to the 2-fold or p-fold cup product with itself. Unlike cup products, however, the operations Sq^i and P^i are stable, i.e., invariant under suspension.

The operations Sq^i generate an algebra $\mathcal{A}_2$, called the Steenrod algebra, such that $H^*(X;\mathbb{Z}_2)$ is a module over $\mathcal{A}_2$ for every space X, and maps between spaces induce homomorphisms of $\mathcal{A}_2$-modules. Similarly, for odd primes p, $H^*(X;\mathbb{Z}_p)$ is a module over a corresponding Steenrod algebra $\mathcal{A}_p$ generated by the P^i's and Bockstein homomorphisms. Like the ring structure given by cup product, these module structures impose strong constraints on spaces and maps between spaces. For example, we will use them to show that there do not exist spaces X having cohomology $H^*(X;Z)$ a polynomial ring $\mathbb{Z}[\alpha]$ unless α has dimension 2 or 4, where there are the familiar examples of $\mathbb{C}\mathrm{P}^\infty$ and $\mathbb{H}\mathrm{P}^\infty$.

This section is divided into two main parts. The first part describes the basic properties of Steenrod squares and powers and gives a number of fairly simple examples and applications. The second part is devoted entirely to the hard work of constructing the squares and powers and showing they satisfy the basic properties listed in the first section, which can be viewed as axioms. More extensive applications will be given in Volume II after spectral sequences have been introduced.

We begin with a few generalities. A **cohomology operation** is a transformation $\Theta = \Theta_X : H^m(X;G) \to H^n(X;H)$ defined for all spaces X, with a fixed choice of m, n, G, and H, satisfying the naturality property that for all maps $f : X \to Y$ there is a commuting diagram

$$\begin{array}{ccc} H^m(Y;G) & \xrightarrow{\Theta_Y} & H^n(Y;H) \\ \downarrow{\scriptstyle f^*} & & \downarrow{\scriptstyle f^*} \\ H^m(X;G) & \xrightarrow{\Theta_X} & H^n(X;H) \end{array}$$

For example, with coefficients in a ring R the transformation $H^m(X;R) \to H^{mp}(X;R)$, $\alpha \mapsto \alpha^p$, is a cohomology operation since $f^*(\alpha^p) = (f^*(\alpha))^p$. This example illustrates that cohomology operations need not be homomorphisms, though this one is when $R = \mathbb{Z}_p$

with p prime. Other examples of cohomology operations we have already encountered are provided by the Bockstein homomorphisms defined in §3.4e. As a more trivial example, a homomorphism $G \to H$ induces change-of-coefficient homomorphisms $H^m(X;G) \to H^m(X;H)$ which can be viewed as cohomology operations.

In spite of their rather general definition, cohomology operations can be described in somewhat more concrete terms:

Proposition 4.86. *For fixed m, n, G, and H there is a bijection between the set of all cohomology operations $\Theta : H^m(X;G) \to H^n(X;H)$ and $H^n(K(G,m);H)$, defined by $\Theta \mapsto \Theta(\iota)$ where $\iota \in H^m(K(G,m);G)$ is a fundamental class.*

Proof: Via CW approximations to spaces, it suffices to restrict attention to CW complexes, so we can identify $H^m(X;G)$ with $[X, K(G,m)]$ as in Theorem 4.17. If $\alpha \in H^m(X;G)$ corresponds to $\varphi : X \to K(G.m)$, i.e., $\varphi^*(\iota) = \alpha$, then $\Theta(\alpha) = \Theta(\varphi^*(\iota)) = \varphi^*(\Theta(\iota))$, so Θ is uniquely determined by $\Theta(\iota)$. Thus $\Theta \mapsto \Theta(\iota)$ is injective. For surjectivity, given an element $\alpha \in H^n(K(G,m);H)$, corresponding to a map $\theta : K(G,m) \to K(H,n)$, then composing with θ defines a transformation $[X, K(G,m)] \to [X, K(H,n)]$, i.e., $\Theta : H^m(X;G) \to H^n(X;H)$, with $\Theta(\iota) = \alpha$. The naturality property for Θ amounts to associativity of the compositions

$$X \xrightarrow{f} Y \xrightarrow{\varphi} K(G,m) \xrightarrow{\theta} K(H,n)$$

and so Θ is a cohomology operation. □

A consequence of the proposition is that cohomology operations cannot decrease dimension, since $K(G,m)$ is $(m-1)$-connected. Moreover, since $H^m(K(G,m);H) \approx Hom(G,H)$, it follows that the only cohomology operations which preserve dimension are given by coefficient homomorphisms.

Most applications of Steenrod squares and powers do not depend on how these operations are actually constructed, but only on their basic properties, specifically, the six or seven properties listed below together with certain relations, called Adém relations, which will be described later. This is similar to the situation for ordinary homology and cohomology, where the axioms generally suffice for most applications. The construction of Steenrod squares and powers and the verification of their basic properties is rather interesting in its own way, but does involve a certain amount of work, so we will delay this until later in the chapter, after we have given some examples and applications elucidating how the operations work in practice.

Here is a list of the main properties which the Steenrod squares $Sq^i : H^n(X;\mathbb{Z}_2) \to H^{n+i}(X;\mathbb{Z}_2)$, $i \geq 0$, will satisfy, beginning with the naturality property:

(1) $Sq^i(f^*(\alpha)) = f^*(Sq^i(\alpha))$ for $f : X \to Y$.
(2) $Sq^i(\alpha + \beta) = Sq^i(\alpha) + Sq^i(\beta)$.
(3) $Sq^i(\alpha \smile \beta) = \sum_j Sq^j(\alpha) \smile Sq^{i-j}(\beta)$ (the Cartan formula).
(4) $Sq^i(\sigma(\alpha)) = \sigma(Sq^i(\alpha))$ where $\sigma : H^n(X;\mathbb{Z}_2) \to H^{n+1}(\Sigma X;\mathbb{Z}_2)$ is the suspension isomorphism of §3.2.
(5) $Sq^i(\alpha) = \alpha^2$ if $i = |\alpha|$, and $Sq^i(\alpha) = 0$ if $i > |\alpha|$.
(6) $Sq^0 = \mathbb{1}$, the identity.
(7) Sq^1 is the $\mathbb{Z}_2$ Bockstein homomorphism β associated with the coefficient sequence $0 \to \mathbb{Z}_2 \to \mathbb{Z}_4 \to \mathbb{Z}_2 \to 0$.

The first part of (5) says that the Steenrod squares extend the squaring operation $\alpha \mapsto \alpha^2$, which has the nice feature of being a homomorphism with $\mathbb{Z}_2$ coefficients. Property (4) says that the Sq^i's are stable operations, invariant under suspension. The actual squaring operation $\alpha \mapsto \alpha^2$ does not have this property since in a suspension ΣX all cup products of positive-dimensional classes are zero, according to an exercise for §3.2.

The fact that Steenrod squares are stable operations extending the cup product square yields the following theorem, which implies that the stable homotopy groups of spheres π_1^s, π_3^s, and π_7^s are nontrivial:

Theorem 4.87. *If $f : S^{2n-1} \to S^n$ has Hopf invariant 1, then $[f] \in \pi_{n-1}^s$ is nonzero, i.e., the iterated suspensions $\Sigma^k f : S^{2n+k-1} \to S^{n+k}$ are all homotopically nontrivial.*

Proof: Associated to a map $f : S^\ell \to S^m$ is the mapping cone $C_f = S^m \cup_f e^{\ell+1}$ with f as the attaching map of the cell $e^{\ell+1}$. Assuming f is basepoint-preserving, we have the relation $C_{\Sigma f} = \Sigma C_f$ where Σ denotes reduced suspension.

If $f : S^{2n-1} \to S^n$ has Hopf invariant 1, then by (5), $Sq^n : H^n(C_f;\mathbb{Z}_2) \to H^{2n}(C_f;\mathbb{Z}_2)$ is nontrivial. By (4), the same is true for $Sq^n : H^{n+k}(\Sigma^k C_f;\mathbb{Z}_2) \to H^{2n+k}(\Sigma^k C_f;\mathbb{Z}_2)$ for all k. If $\Sigma^k f$ were homotopically trivial we would have a retraction $r : \Sigma^k C_f \to S^{n+k}$, pulling a generator $\alpha \in H^{n+k}(S^{n+k};\mathbb{Z}_2)$ back to a generator $r^*(\alpha) \in H^{n+k}(\Sigma^k C_f;\mathbb{Z}_2)$. Then since $Sq^n(\alpha) = 0$ in $H^{2n+k}(S^{n+k};\mathbb{Z}_2)$, property (1) would imply that $Sq^n(r^*(\alpha)) = r^*(Sq^n(\alpha)) = r^*(0) = 0$, a contradiction. □

The Steenrod power operations $P^i : H^n(X;\mathbb{Z}_p) \to H^{n+2i(p-1)}(X;\mathbb{Z}_p)$ for odd primes p will satisfy analogous properties:

(1) $P^i(f^*(\alpha)) = f^*(P^i(\alpha))$ for $f : X \to Y$.
(2) $P^i(\alpha + \beta) = P^i(\alpha) + P^i(\beta)$.

(3) $P^i(\alpha \smile \beta) = \sum_j P^j(\alpha) \smile P^{i-j}(\beta)$ (the Cartan formula).

(4) $P^i(\sigma(\alpha)) = \sigma(P^i(\alpha))$ where $\sigma : H^n(X;\mathbb{Z}_p) \to H^{n+1}(\Sigma X;\mathbb{Z}_p)$ is the suspension isomorphism of §3.2.

(5) $P^i(\alpha) = \alpha^p$ if $2i = |\alpha|$, and $P^i(\alpha) = 0$ if $2i > |\alpha|$.

(6) $P^0 = \mathbb{1}$, the identity.

The germinal property $P^i(\alpha) = \alpha^p$ can only be expected to hold for even-dimensional classes α since for odd-dimensional α the commutativity property of cup product implies that $\alpha^2 = 0$ with $\mathbb{Z}_p$ coefficients if p is odd, and then $\alpha^p = 0$ since $\alpha^2 = 0$. Note that the formula $P^i(\alpha) = \alpha^p$ for $|\alpha| = 2i$ implies that P^i raises dimension by $2i(p-1)$, explaining the appearance of this number.

The Bockstein homomorphism $\beta : H^n(X;\mathbb{Z}_p) \to H^{n+1}(X;\mathbb{Z}_p)$ is not among the P^i's, but this is mainly a matter of notational convenience. As we shall see below, the operation Sq^{2i+1} is the same as the composition $Sq^1Sq^{2i} = \beta Sq^{2i}$, so the Sq^{2i}'s can be regarded as the P^i's for $p = 2$.

One might ask if there are elements of π_*^s detectable by Steenrod powers in the same way that the Hopf maps are detected by Steenrod squares. The answer is yes, but only for the operation P^1. Let us describe a construction for a map $f : S^{2p} \to S^3$ such that in the mapping cone C_f, the operation $P^1 : H^3(C_f;\mathbb{Z}_p) \to H^{2p+1}(C_f;\mathbb{Z}_p)$ is nonzero, hence f in nonzero in π^s_{2p-3}. The construction starts with the fact that a generator of $H^2(K(\mathbb{Z}_p,1);\mathbb{Z}_p)$ has nontrivial p^{th} power, so $P^1 : H^2(K(\mathbb{Z}_p,1);\mathbb{Z}_p) \to H^{2p}(K(\mathbb{Z}_p,1);\mathbb{Z}_p)$ is nontrivial by property (5). This remains true after we suspend to $\Sigma K(\mathbb{Z}_p,1)$, and we showed in §4.4h that $\Sigma K(\mathbb{Z}_p,1)$ has the homotopy type of a wedge of CW complexes X_i, $1 \le i \le p-1$, with $\widetilde{H}_*(X_i;\mathbb{Z})$ consisting only of a $\mathbb{Z}_p$ in each dimension congruent to $2i$ mod $2p-2$. We are interested here in the space $X = X_1$, which has nontrivial $\mathbb{Z}_p$ cohomology in dimensions $2, 3, 2p, 2p+1, \cdots$. Since X is, up to homotopy, a wedge summand of $\Sigma K(\mathbb{Z}_p,1)$, the operation $P^1 : H^3(X;\mathbb{Z}_p) \to H^{2p+1}(X;\mathbb{Z}_p)$ is nontrivial. Since X is simply-connected, the construction in §4.4c shows that we may take X to have $(2p+1)$-skeleton of the form $S^2 \cup e^3 \cup e^{2p} \cup e^{2p+1}$. In fact, using the notion of homology decomposition in §4.4g, we can take this skeleton to be the reduced mapping cone C_g of a map of Moore spaces $g : M(\mathbb{Z}_p, 2p-1) \to M(\mathbb{Z}_p, 2)$. The quotient C_g/S^2 is the reduced mapping cone of the composition of g with the quotient map $M(\mathbb{Z}_p,2) \to M(\mathbb{Z}_p,2)/S^2 = S^3$. The restriction of this map $M(\mathbb{Z}_p,2p-1) \to S^3$ to S^{2p-1} represents an element of order p in $\pi_{2p-1}(S^3)$ since this restriction extends over the $2p$-cell of $M(\mathbb{Z}_p,2p-1)$ which is attached by a map $S^{2p-1} \to S^{2p-1}$ of degree p. Now we use the fact, proved in Volume II using the Serre spectral sequence, that $\pi_i(S^3) = 0$

contains no elements of order p for $i \le 2p-1$. This implies that C_g/S^2 is homotopy equivalent to a CW complex Y obtained from $S^3 \vee S^{2p}$ by attaching a cell e^{2p+1}. The quotient Y/S^{2p} then has the form $S^3 \cup e^{2p+1}$, so it is the reduced mapping cone of a map $f : S^{2p} \to S^3$. By construction there is a map $C_g \to C_f$ inducing an isomorphism on $\mathbb{Z}_p$ cohomology in dimensions 3 and $2p+1$, so the operation P^1 is nontrivial in $H^*(C_f;\mathbb{Z}_p)$ since this was true for C_g which was the $(2p+1)$-skeleton of X.

It is a perhaps disappointing fact that no other squares or powers besides Sq^1, Sq^2, Sq^4, Sq^8, and P^1 detect elements of homotopy groups of spheres. (We include Sq^1 in this list since for any p the mod p Bockstein detects a map $S^n \to S^n$ of degree p.) We will show this for certain Sq^i's and P^i's later in this section. The general case is the 'mod p Hopf invariant one' theorem, which will be proved in [VBKT].

The Cartan formulas can be put into a more concise form by defining *total* Steenrod square and power operations by $Sq = Sq^0 + Sq^1 + \cdots$ and $P = P^0 + P^1 + \cdots$. These act on $H^*(X;\mathbb{Z}_p)$ since by property (5), only a finite number of Sq^i's or P^i's can be nonzero on a given cohomology class. The Cartan formulas then say that $Sq(\alpha \smile \beta) = Sq(\alpha) \smile Sq(\beta)$ and $P(\alpha \smile \beta) = P(\alpha) \smile P(\beta)$, so Sq and P are ring homomorphisms.

We can use Sq and P to compute how the operations Sq^i and P^i act in projective spaces and lens spaces using the general formulas

$$Sq^i(\alpha^n) = \tbinom{n}{i}\alpha^{n+i} \quad \text{for} \quad \alpha \in H^1(X;\mathbb{Z}_2)$$
$$P^i(\alpha^n) = \tbinom{n}{i}\alpha^{n+i(p-1)} \quad \text{for} \quad \alpha \in H^2(X;\mathbb{Z}_p)$$

To derive the first formula, properties (5) and (6) give $Sq(\alpha) = \alpha + \alpha^2 = \alpha(1+\alpha)$ so $Sq(\alpha^n) = Sq(\alpha)^n = \alpha^n(1+\alpha)^n = \sum_i \binom{n}{i}\alpha^{n+i}$ hence $Sq^i(\alpha^n) = \binom{n}{i}\alpha^{n+i}$. The second formula is obtained in similar fashion: $P(\alpha) = \alpha + \alpha^p = \alpha(1+\alpha^{p-1})$ so $P(\alpha^n) = \alpha^n(1+\alpha^{p-1})^n = \sum_i \binom{n}{i}\alpha^{n+i(p-1)}$.

In Lemma 3.41 we described how binomial coefficients can be computed modulo a prime p:

$$\tbinom{m}{n} \equiv \textstyle\prod_i \binom{m_i}{n_i} \mod p, \qquad \text{where } m = \sum_i m_i p^i \text{ and } n = \sum_i n_i p^i \text{ are the } p\text{-adic expansions of } m \text{ and } n.$$

Taking $p = 2$ for example, and considering the extreme cases of a dyadic expansion consisting of a single 1 or all 1's, we obtain:

$$Sq(\alpha^{2^k}) = \alpha^{2^k} + \alpha^{2^{k+1}}$$
$$Sq(\alpha^{2^k-1}) = \alpha^{2^k-1} + \alpha^{2^k} + \alpha^{2^k+1} + \cdots + \alpha^{2^{k+1}-2}$$

for $\alpha \in H^1(X;\mathbb{Z}_2)$. More generally, the coefficients of $Sq(\alpha^n)$ can be read off from the $(n+1)^{st}$ row of the mod 2 Pascal triangle, shown in the figure below, where dots denote zeros.

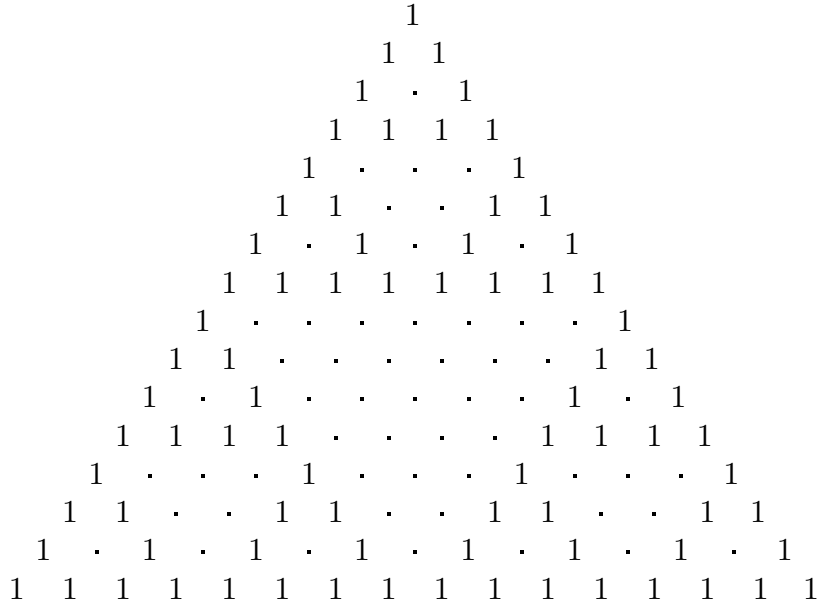

Example: Stable splittings. These formulas tell us how to compute Steenrod squares for $\mathbb{R}\mathrm{P}^\infty$, hence also for any suspension of $\mathbb{R}\mathrm{P}^\infty$. The formulas for $Sq(\alpha^{2^k})$ and $Sq(\alpha^{2^k-1})$ written down above show that all the powers of the generator $\alpha \in H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$ are tied together by Steenrod squares since the first formula connects α inductively to all the powers α^{2^k} and the second formula connects these powers to all the other powers. This shows that no suspension $\Sigma^k\mathbb{R}\mathrm{P}^\infty$ has the homotopy type of a wedge sum $X \vee Y$ with both X and Y having nontrivial cohomology. In the case of $\mathbb{R}\mathrm{P}^\infty$ itself we could have deduced this from the ring structure of $H^*(\mathbb{R}\mathrm{P}^\infty;Z_2) \approx \mathbb{Z}_2[\alpha]$, but cup products become trivial in a suspension.

The same reasoning shows that $\mathbb{C}\mathrm{P}^\infty$ and $\mathbb{H}\mathrm{P}^\infty$ have no nontrivial stable splittings. The $\mathbb{Z}_2$ cohomology in these cases is again $\mathbb{Z}_2[\alpha]$, though with α no longer 1-dimensional. However, we still have $Sq(\alpha) = \alpha + \alpha^2$ since these spaces have no nontrivial cohomology in the dimensions between α and α^2, so we have $Sq^{2i}(\alpha^n) = \binom{n}{i}\alpha^{n+i}$ in the complex case and $Sq^{4i}(\alpha^n) = \binom{n}{i}\alpha^{n+i}$ in the quaternionic case. Then the arguments from the real case carry over using the operations Sq^{2i} and Sq^{4i} in place of Sq^i.

Suppose we consider the same question for $K(\mathbb{Z}_3,1)$ instead of $\mathbb{R}\mathrm{P}^\infty$. Taking cohomology with $\mathbb{Z}_3$ coefficients, the Bockstein β is nonzero on odd-dimensional classes in $H^*(K(\mathbb{Z}_3,1);\mathbb{Z}_3)$, thus tying them to the even-dimensional classes, so we only need to see which even-dimensional classes are connected by P^i's. The even-dimensional part of $H^*(K(\mathbb{Z}_3,1);\mathbb{Z}_3)$ is a polynomial algebra $\mathbb{Z}_3[\alpha]$ with $|\alpha| = 2$, so we have $P^i(\alpha^n) =$

$\binom{n}{i}\alpha^{n+i(p-1)} = \binom{n}{i}\alpha^{n+4i}$ by our earlier formula. Since P^i raises dimension by $4i$, there is no chance that all the even-dimensional cohomology will be connected by the P^i's. In fact, we showed in §4.4h that $\Sigma K(\mathbb{Z}_3,1) \simeq X_1 \vee X_2$ where X_1 has the cohomology of $\Sigma K(\mathbb{Z}_3,1)$ in dimensions congruent to 2 and 3 mod 4, while X_2 has the remaining cohomology. Thus the best one could hope would be that all the odd powers of α are connected by P^i's and likewise all the even powers are connected, since this would imply that neither X_1 nor X_2 splits nontrivially. This is indeed the case, as one sees by an examination of the coefficients in the formula $P^i(\alpha^n) = \binom{n}{i}\alpha^{n+4i}$. The Pascal triangle mod 3 begins:

$$
\begin{array}{ccccccccccc}
 & & & & & 1 & & & & & \\
 & & & & 1 & & 1 & & & & \\
 & & & 1 & & 2 & & 1 & & & \\
 & & 1 & & \cdot & & \cdot & & 1 & & \\
 & 1 & & 1 & & \cdot & & 1 & & 1 & \\
1 & & 2 & & 1 & & 1 & & 2 & & 1 \\
1 & & \cdot & & \cdot & & 2 & & \cdot & & \cdot \\
1 & & 1 & & \cdot & & 2 & & 2 & & \cdot \\
1 & & 2 & & 1 & & 2 & & 1 & & 2 \\
1 & & \cdot & & \cdot & & \cdot & & \cdot & & \cdot
\end{array}
\qquad
\begin{array}{c}
 \\
 \\
 \\
 \\
 \\
 \\
 \\
 \\
 \\
 \\
1
\end{array}
$$

Here $P(\alpha^n)$ is determined by the $(n+1)^{st}$ row. For example the sixth row says that $P(\alpha^5) = \alpha^5 + 2\alpha^7 + \alpha^9 + \alpha^{11} + 2\alpha^{13} + \alpha^{15}$. A few moments' thought shows that the rows which compute $P(\alpha^n)$ for $n = k3^m - 1$ have all nonzero entries, and these rows together with the rows right after them suffice to connect the powers of α in the desired way, so that X_1 and X_2 have no stable splittings. One can also see that although $\Sigma^2 X_1$ and X_2 have isomorphic homology groups, they are not homotopy equivalent, even stably, since the operations P^i act differently in the two spaces, e.g., P^2 is trivial on suspensions of α but not on suspensions of α^2.

The situation for $K(\mathbb{Z}_p,1)$ for larger primes p is entirely similar, with $\Sigma K(\mathbb{Z}_p,1)$ splitting as a wedge of $p-1$ spaces. The same arguments work more generally for $K(\mathbb{Z}_{p^i},1)$, though for $i > 1$ the usual Bockstein β is identically zero so one has to use instead a Bockstein involving $\mathbb{Z}_{p^i}$ coefficients. We leave the details of these arguments as exercises.

Example: Maps of $\mathbb{H}\mathrm{P}^\infty$. We can use the operations P^i together with a little number theory to demonstrate an interesting distinction between $\mathbb{H}\mathrm{P}^\infty$ and $\mathbb{C}\mathrm{P}^\infty$, namely, we will show that if a map $f\colon \mathbb{H}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ has $f^*(\gamma) = d\gamma$ for γ a generator of $H^4(\mathbb{H}\mathrm{P}^\infty;\mathbb{Z})$, then the integer d, which we call the *degree* of f, must be a square. By constrast, since $\mathbb{C}\mathrm{P}^\infty$ is a $K(\mathbb{Z},2)$, maps $\mathbb{C}\mathrm{P}^\infty \to \mathbb{C}\mathrm{P}^\infty$ exist carrying a generator $\alpha \in H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ onto any given multiple of itself. Explicitly, the map $z \mapsto z^d$, $z \in \mathbb{C}$, induces a map f of $\mathbb{C}\mathrm{P}^\infty$

with $f^*(\alpha) = d\alpha$, but commutativity of $\mathbb{C}$ is needed for this construction so it doesn't extend to the quaternionic case.

We shall deduce the action of Steenrod powers on $H^*(\mathbb{H}\mathrm{P}^\infty;\mathbb{Z}_p)$ from their action on $H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}_p)$. We have the earlier formula $P^i(\alpha^n) = \binom{n}{i}\alpha^{n+i(p-1)}$ for α a generator of $H^2(X;\mathbb{Z}_p)$. There is a natural quotient map $\mathbb{C}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ arising from the definition of both spaces as quotients of S^∞. This map takes the 4-cell of $\mathbb{C}\mathrm{P}^\infty$ homeomorphically onto the 4-cell of $\mathbb{H}\mathrm{P}^\infty$, so the induced map on cohomology sends a generator $\gamma \in H^4(\mathbb{H}\mathrm{P}^\infty;\mathbb{Z}_p)$ to α^2, hence γ^n to α^{2n}. Thus the formula $P^i(\alpha^{2n}) = \binom{2n}{i}\alpha^{2n+i(p-1)}$ implies that $P^i(\gamma^n) = \binom{2n}{i}\gamma^{n+i(p-1)/2}$. For example, $P^1(\gamma) = 2\gamma^{(p+1)/2}$.

Now let $f:\mathbb{H}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ be any map. Applying the formula $P^1(\gamma) = 2\gamma^{(p+1)/2}$ in two ways, we get:

$$\begin{aligned} & P^1 f^*(\gamma) = f^* P^1(\gamma) = f^*(2\gamma^{(p+1)/2}) = 2d^{(p+1)/2}\gamma^{(p+1)/2} \\ \text{and} \quad & P^1 f^*(\gamma) = P^1(d\gamma) = 2d\gamma^{(p+1)/2} \end{aligned}$$

Hence the degree d satisfies $d^{(p+1)/2} \equiv d \mod p$ for all odd primes p. Thus either d is 0 mod p or $d^{(p-1)/2} \equiv 1 \mod p$. In both cases d is a square mod p since the congruence $d^{(p-1)/2} \equiv 1 \mod p$ is equivalent to d being a nonzero square mod p, the multiplicative group of nonzero elements of the field $\mathbb{Z}_p$ being cyclic of order $p-1$.

The argument is completed by appealing to the number theory fact that an integer which is a square mod p for all sufficiently large primes p must be a square. This can be deduced from quadratic reciprocity and Dirichlet's theorem on primes in arithmetic progressions as follows. Suppose on the contrary that the result is false for the integer d. Consider primes p not dividing d. Since the product of two squares in $\mathbb{Z}_p$ is again a square, we may assume that d is a product of distinct primes $q_1, \cdots, q_n$, where one of these 'primes' is allowed to be -1 if d is negative. In terms of the Legendre symbol $\left(\frac{d}{p}\right)$ which is defined to be $+1$ if d is a square mod p and -1 otherwise, we have

$$\left(\frac{d}{p}\right) = \left(\frac{q_1}{p}\right)\cdots\left(\frac{q_n}{p}\right)$$

The left side is $+1$ for all large p by hypothesis, so it will suffice to see that p can be chosen to give each term on the right an arbitrary preassigned value. The values of $\left(\frac{-1}{p}\right)$ and $\left(\frac{2}{p}\right)$ depend only on p mod 8, and the four combinations of values are realized by the four residues $1, 3, 5, 7 \mod 8$. Having specified the value of p mod 8, the quadratic reciprocity law then says that for odd primes q, specifying $\left(\frac{q}{p}\right)$ is equivalent to specifying $\left(\frac{p}{q}\right)$. Thus we need only choose p in the appropriate residue classes mod 8 and mod q_i for each odd q_i. By the Chinese remainder theorem, this means specifying p modulo 8 times

a product of odd primes. Dirichlet's theorem guarantees that in fact infinitely many primes p exist satisfying this congruence condition.

It is known that the integers which are realizable as degrees of maps $\mathbb{H}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ are exactly the odd squares and zero. The construction of maps of odd square degree will be given in Volume II using localization techniques, following [Sullivan]. Ruling out nonzero even squares can be done using K-theory; see [Feder-Gitler], which also treats maps $\mathbb{H}\mathrm{P}^n \to \mathbb{H}\mathrm{P}^n$.

The preceding calculations can also be used to show that every map $\mathbb{H}\mathrm{P}^n \to \mathbb{H}\mathrm{P}^n$ must have a fixed point if $n > 1$. For, taking $p = 3$, the element $P^1(\gamma)$ lies in $H^8(\mathbb{H}\mathrm{P}^n; \mathbb{Z}_3)$ which is nonzero if $n > 1$, so, when the earlier argument is specialized to the case $p = 3$, the congruence $d^{(p+1)/2} \equiv d \mod p$ becomes $d^2 = d$ in $\mathbb{Z}_3$, which is satisfied only by 0 and 1 in $\mathbb{Z}_3$. Hence $f^* : \widetilde{H}^*(\mathbb{H}\mathrm{P}^n; \mathbb{Z}_3) \to \widetilde{H}^*(\mathbb{H}\mathrm{P}^n; \mathbb{Z}_3)$ is either zero or the identity. In both cases the Lefschetz number $\lambda(f)$, which is the sum of the traces of the maps $f^* : H^{4i}(\mathbb{H}\mathrm{P}^n; \mathbb{Z}_3) \to H^{4i}(\mathbb{H}\mathrm{P}^n; \mathbb{Z}_3)$, is nonzero, so the Lefschetz fixed point theorem gives the result.

Example: Vector fields on spheres. Let us now apply Steenrod squares to determine the maximum number of orthogonal tangent vector fields on a sphere in all cases except when the dimension of the sphere is congruent to $-1 \mod 16$. The first step is to rephrase the question in terms of Stiefel manifolds. Recall from §3.4d and §4.3 the space $V_{n,k}$ of orthonormal k-frames in $\mathbb{R}^n$. Projection of a k-frame onto its first vector gives a map $p : V_{n,k} \to S^{n-1}$, and a section for this projection, i.e., a map $f : S^{n-1} \to V_{n,k}$ such that $pf = \mathbb{1}$, is exactly a set of $k-1$ orthonormal tangent vector fields $v_1, \cdots, v_{k-1}$ on S^{n-1} since f assigns to each $x \in S^{n-1}$ an orthonormal k-frame $(x, v_1(x), \cdots, v_{k-1}(x))$.

We described a cell structure on $V_{n,k}$ at the end of §3.4d, and we claim that the $(n-1)$-skeleton of this cell structure is $\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1}$ if $2k-1 \le n$. The cells of $V_{n,k}$ were products $e^{i_1} \times \cdots \times e^{i_m}$ with $n > i_1 > \cdots > i_m \ge n-k$, so the products with a single factor account for all of the $(2n-2k)$-skeleton, hence all of the $(n-1)$-skeleton if $n-1 \le 2n-2k$, i.e., if $2k-1 \le n$. The cells which are products with a single factor are the homeomorphic images of cells of $\mathbb{R}\mathrm{P}^{n-1}$ under a map $\mathbb{R}\mathrm{P}^{n-1} \to SO(n) \to SO(n)/SO(n-k) = V_{n,k}$. This map collapses $\mathbb{R}\mathrm{P}^{n-k-1}$ to a point, so we get the desired conclusion that $\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1}$ is the $(n-1)$-skeleton if $2k-1 \le n$.

Now suppose we have $f : S^{n-1} \to V_{n,k}$ such that $pf = \mathbb{1}$, as before. In particular, f^* is surjective on $H^{n-1}(-; \mathbb{Z}_2)$. If we deform f to a cellular map, with image in the $(n-1)$-skeleton, then by the preceding paragraph this will give a map $g : S^{n-1} \to \mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1}$ if $2k-1 \le n$, and this map will still induce a surjection, hence an

isomorphism, on $H^{n-1}(-;\mathbb{Z}_2)$. If we choose n and k so that $\binom{n-k}{k-1} \equiv 1 \mod 2$, then $Sq^{k-1}: H^{n-k}(\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1};\mathbb{Z}_2) \to H^{n-1}(\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1};\mathbb{Z}_2)$ will be nonzero, but this contradicts the existence of the map g since the operation $Sq^{k-1}: H^{n-k}(S^{n-1};\mathbb{Z}_2) \to H^{n-1}(S^{n-1};\mathbb{Z}_2)$ is zero.

In order to guarantee that $\binom{n-k}{k-1} \equiv 1 \mod 2$, write $n = 2^r(2s+1)$ and choose $k = 2^r + 1$. Assume for the moment that $s \geq 1$. Then $\binom{n-k}{k-1} = \binom{2^{r+1}s-1}{2^r}$, and in view of the rule for computing binomial coefficients in $\mathbb{Z}_2$, this is nonzero since the dyadic expansion of $2^{r+1}s-1$ ends with a string of 1's including a 1 in the single digit where the expansion of 2^r is nonzero. Note that the earlier condition $2k-1 \leq n$ is satisfied since it becomes $2^{r+1}+1 \leq 2^{r+1}s + 2^r$ and we assume $s \geq 1$.

Summarizing, we have shown that for $n = 2^r(2s+1)$, the sphere S^{n-1} cannot have 2^r orthonormal tangent vector fields, if $s \geq 1$. This is also trivially true for $s = 0$ since if S^{n-1} had n orthonormal tangent vector fields, this would say $V_{n,n+1}$ was nonempty, which is certainly false.

It is easy to see that this result is best possible when $r \leq 3$ by explicitly constructing $2^r - 1$ orthonormal vector fields on S^{n-1} when $n = 2^r m$. When $r = 1$, view S^{n-1} as the unit sphere in $\mathbb{C}^m$, and then $x \mapsto ix$ defines a tangent vector field since the unit complex numbers 1 and i are orthogonal and multiplication by a unit complex number is an isometry of $\mathbb{C}$, so x and ix are orthogonal in each coordinate of $\mathbb{C}^m$, hence are orthogonal. When $r = 2$ the same construction works with $\mathbb{H}$ in place of $\mathbb{C}$, using the three maps $x \mapsto ix$, $x \mapsto jx$, and $x \mapsto kx$ to define orthonormal tangent vector fields on the unit sphere in $\mathbb{H}^m$. When $r = 3$ we can follow the same procedure with the octonions, constructing 7 orthonormal tangent vector fields to the unit sphere in $\mathbb{O}^m$ via an orthonormal basis $1, i, j, k, \cdots$ for $\mathbb{O}$.

The upper bound of $2^r - 1$ for the number of orthogonal vector fields to S^{n-1} is not best possible in the remaining case that $n \equiv 0 \mod 16$. The optimal upper bound is obtained instead using K-theory; see [VBKT] or [Husemoller]. The construction of the requisite number of vector fields is again algebraic, this time using Clifford algebras.

Adém Relations and the Steenrod Algebra

When Steenrod squares or powers are composed, the compositions satisfy certain

relations, unfortunately rather complicated, known as **Adém relations**:

$$Sq^a Sq^b = \sum_j \tbinom{b-j-1}{a-2j} Sq^{a+b-j} Sq^j \qquad \text{if } a < 2b$$

$$P^a P^b = \sum_j (-1)^{a+j} \tbinom{(p-1)(b-j)-1}{a-pj} P^{a+b-j} P^j \qquad \text{if } a < pb$$

$$P^a \beta P^b = \sum_j (-1)^{a+j} \tbinom{(p-1)(b-j)}{a-pj} \beta P^{a+b-j} P^j - \sum_j (-1)^{a+j} \tbinom{(p-1)(b-j)-1}{a-pj-1} P^{a+b-j} \beta P^j \qquad \text{if } a \le pb$$

By convention, the binomial coefficient $\binom{m}{n}$ is taken to be zero if m or n is negative or if $m < n$. Also $\binom{m}{0} = 1$ for $m \ge 0$.

For example, taking $a = 1$ in the Adém relation for the Steenrod squares we have $Sq^1 Sq^b = (b-1)Sq^{b+1}$, so $Sq^1 Sq^{2i} = Sq^{2i+1}$ and $Sq^1 Sq^{2i+1} = 0$. The relations $Sq^1 Sq^{2i} = Sq^{2i+1}$ and $Sq^1 = \beta$ explain the earlier comment that Sq^{2i} is the analog of P^i for $p = 2$.

The **Steenrod algebra** $\mathcal{A}_2$ is defined to be the algebra over $\mathbb{Z}_2$ which is the quotient of the algebra of polynomials in the noncommuting variables $Sq^1, Sq^2, \cdots$ by the two-sided ideal generated by the Adém relations, i.e., by the polynomials given by the differences between the left and right sides of the Adém relations. In similar fashion, $\mathcal{A}_p$ for odd p is defined to be the algebra over $\mathbb{Z}_p$ formed by polynomials in the noncommuting variables $\beta, P^1, P^2, \cdots$ modulo the Adém relations and the relation $\beta^2 = 0$. Thus for every space X, $H^*(X; \mathbb{Z}_p)$ is a module over $\mathcal{A}_p$. The Steenrod algebra is a graded algebra, the elements of degree k being those which map $H^n(X; \mathbb{Z}_p)$ to $H^{n+k}(X; \mathbb{Z}_p)$ for all n.

The next proposition implies that $\mathcal{A}_2$ is generated as an algebra by the elements Sq^{2^k}, and $\mathcal{A}_p$ for p odd is generated by β and the elements P^{p^k}.

Proposition 4.88. *If i is not a power of 2, there is a relation $Sq^i = \sum_{0<j<i} a_j Sq^{i-j} Sq^j$ with coefficients $a_j \in \mathbb{Z}_2$. Similarly, if i is not a power of p there is a relation $P^i = \sum_{0<j<i} a_j P^{i-j} P^j$ with $a_j \in \mathbb{Z}_p$.*

One says that Sq^i or P^i is **decomposable** if such a relation holds, expressing it as a linear combination of compositions of Sq^j's, or P^j's and β's, with $j < i$. In other words, Sq^i or P^i is decomposable if it is in the image of the multiplication map $\widetilde{\mathcal{A}}_p \otimes \widetilde{\mathcal{A}}_p \to \mathcal{A}_p$ where $\widetilde{\mathcal{A}}_p \subset \mathcal{A}_p$ consists of the elements without degree 0 terms. The operation Sq^{2^k} is indecomposable since for α a generator of $H^1(\mathbb{R}\mathrm{P}^\infty; \mathbb{Z}_2)$ we saw that $Sq^{2^k}(\alpha^{2^k}) = \alpha^{2^{k+1}}$ but $Sq^i(\alpha^{2^k}) = 0$ for $0 < i < 2^k$. Similarly P^{p^k} is indecomposable since if $\alpha \in H^2(\mathbb{C}\mathrm{P}^\infty; \mathbb{Z}_p)$ is a generator then $P^{p^k}(\alpha^{p^k}) = \alpha^{p^{k+1}}$ but $P^i(\alpha^{p^k}) = 0$ for $0 < i < p^k$ and also $\beta(\alpha^{p^k}) = 0$.

Proof of 4.88: The argument is the same for $p = 2$ and p odd, so we describe the latter case. The idea is to write i as the sum of integers $a > 0$ and $b > 0$ with $a < pb$ in such a way that the coefficient of the $j = 0$ term in the Adém relation for P^aP^b is nonzero. Then one can solve this relation for $P^{a+b} = P^i$.

Let $i = i_0 + i_1p + \cdots + i_kp^k$ be the p-adic representation of i, with $i_k \neq 0$. Let $b = p^k$ and $a = i - p^k$, so $b > 0$ and $a > 0$ if i is not a power of p. The claim is that $\binom{(p-1)b-1}{a}$ is nonzero in $\mathbb{Z}_p$. The p-adic expansion of $(p-1)b - 1 = (p^{k+1} - 1) - p^k$ is $(p-1)+(p-1)p+\cdots+(p-2)p^k$, and for a the p-adic expansion is $i_0+i_1p+\cdots+(i_k-1)p^k$. Hence $\binom{(p-1)b-1}{a} \equiv \binom{p-1}{i_0} \cdots \binom{p-2}{i_k-1}$ and in each factor of the latter product the numerator is nonzero in $\mathbb{Z}_p$ so the product is nonzero in $\mathbb{Z}_p$. When $p = 2$ the last factor is omitted, and the product is still nonzero in $\mathbb{Z}_2$. □

This proposition implies that if an element of π^s_* is detected by a Sq^i or P^i then i must be a power of 2 for Sq^i and a power of p for P^i. For if Sq^i is decomposable, then the map $Sq^i : H^n(C_f : \mathbb{Z}_2) \to H^{n+i}(C_f; \mathbb{Z}_2)$ must be trivial since it is a sum of compositions which pass through trivial cohomology groups, and similarly for P^i.

Another consequence of decomposability is the following:

Theorem 4.89. *Suppose $H^*(X; \mathbb{Z}_p)$ is the polynomial algebra $\mathbb{Z}_p[\alpha]$ on a generator α of dimension n, possibly truncated by the relation $\alpha^m = 0$ for $m > p$. Then if $p = 2$, n must be a power of 2, and if p is an odd prime, n must be of the form $p^k\ell$ where ℓ is an even divisor of $2(p-1)$.*

As we mentioned in §3.2, there is a stronger theorem that n must be 1, 2, 4, or 8 when $p = 2$, and n must be an even divisor of $2(p-1)$ when p is an odd prime. We also gave examples showing the necessity of the hypothesis $m > p$ in the case of a truncated polynomial algebra.

Proof: In the case $p = 2$, $Sq^n(\alpha) = \alpha^2 \neq 0$. If n is not a power of 2 then Sq^n decomposes into compositions $Sq^{n-j}Sq^j$ with $0 < j < n$. Such compositions must be zero since they pass through the group $H^{n+j}(X; \mathbb{Z}_2)$ which is zero for $0 < j < n$.

For odd p, the fact that α^2 is nonzero implies that n is even, say $n = 2k$. Then $P^k(\alpha) = \alpha^p \neq 0$. Since P^k can be expressed in terms of P^{p^i}'s, some P^{p^i} must be nonzero in $H^*(X; \mathbb{Z}_p)$. This implies that $2p^i(p-1)$, the amount by which P^{p^i} raises dimension, must be a multiple of n since $H^*(X; \mathbb{Z}_p)$ is concentrated in dimensions which are multiples of n. Since n divides $2p^i(p-1)$, it must be a power of p times a divisor of $2(p-1)$, and this divisor must be even since n is even and p is odd. □

Corollary 4.90. *If $H^*(X;\mathbb{Z})$ is a polynomial algebra $\mathbb{Z}[\alpha]$, possibly truncated by $\alpha^m = 0$ with $m > 3$, then $|\alpha| = 2$ or 4.*

Proof: Passing to $\mathbb{Z}_2$ coefficients, the theorem implies that $|\alpha|$ is a power of 2, and taking $\mathbb{Z}_3$ coefficients we see that $|\alpha|$ is also a power of 3 times a divisor of $2(3-1) = 4$. □

In particular, the octonionic projective plane $\mathbb{O}\mathrm{P}^2$, constructed in §4.3 by attaching a 16-cell to S^8 via the Hopf map $S^{15} \to S^8$, does not generalize to an octonionic projective n-space $\mathbb{O}\mathrm{P}^n$ with $n \geq 3$.

We have seen how the Adém relations put strong limits on detecting elements of π_*^s by Steenrod squares and powers. However, the Adém relations can also be used in a positive way to detect elements of π_*^s:

Proposition 4.91. *If $\eta \in \pi_1^s$ is represented by the Hopf map $S^3 \to S^2$, then η^2 is nonzero in π_2^s. Similarly, the other two Hopf maps represent elements $\nu \in \pi_3^s$ and $\sigma \in \pi_7^s$ whose squares are nontrivial in π_6^s and π_{14}^s.*

Proof: Let $\eta : S^{n+1} \to S^n$ be a suspension of the Hopf map, with $C_\eta = S^n \cup_\eta e^{n+2}$. Letting C and S denote unreduced cone and suspension, we have then the following diagram where Φ is the characteristic map of the cell e^{n+2}.

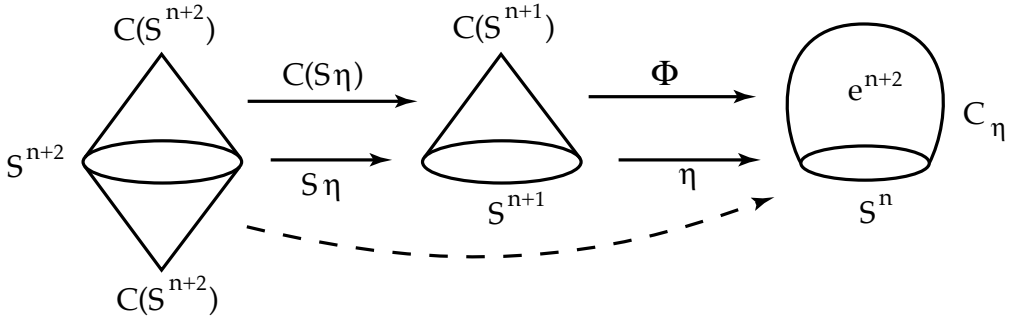

If we assume that the composition $S\eta \circ \eta$ is nullhomotopic, then there is an extension of $S\eta \circ \eta : S^{n+2} \to S^n$ over the lower cone $C(S^{n+2})$. The diagram then gives a map $S^{n+3} = S(S^{n+2}) \to C_\eta$ which we use to attach a cell e^{n+4} to C_η, forming the space Y. Since C_η is the $(n+2)$-skeleton of Y, the map $Sq^2 : H^n(Y) \to H^{n+2}(Y)$ is an isomorphism. The quotient Y/S^n is $C_{S^2\eta}$ and the projection $Y \to Y/S^n$ induces an isomorphism on H^{n+2} and H^{n+4}, from which it follows that $Sq^2 : H^{n+2}(Y) \to H^{n+4}(Y)$ is an isomorphism. Thus the composition $Sq^2Sq^2 : H^n(Y) \to H^{n+4}(Y)$ is an isomorphism. But this is impossible in view of the Adém relation $Sq^2Sq^2 = Sq^3Sq^1$, since $Sq^1 : H^n(Y) \to H^{n+1}(Y) = 0$ is trivial.

The same argument shows that ν^2 and σ^2 are nontrivial using the relations $Sq^4Sq^4 = Sq^7Sq^1 + Sq^6Sq^2$ and $Sq^8Sq^8 = Sq^{15}Sq^1 + Sq^{14}Sq^2 + Sq^{12}Sq^4$. □

This line of reasoning does not work for odd primes since for the element $\alpha \in \pi^s_{2p-3}$ detected by P^1, the Adém relation for P^1P^1 is $P^1P^1 = 2P^2$, which is not helpful. And in fact $\alpha^2 = 0$ by the skew-commutativity of the product in π^s_*.

When dealing with $\mathcal{A}_2$ it is often convenient to abbreviate notation by writing a monomial $Sq^{i_1}Sq^{i_2}\cdots$ as Sq^I where I is the finite sequence of non-negative integers $i_1, i_2, \cdots$ (or equivalently, an infinite sequence with only finitely many nonzero terms). Call Sq^I **admissible** if no Adém relation can be applied to it, i.e., if $i_j \geq 2i_{j+1}$ for all j. The Adém relations imply that every monomial Sq^I can be written as a sum of admissible monomials. For if Sq^I is not admissible, it contains a pair Sq^aSq^b to which the Adém relation can be applied, yielding a sum of terms Sq^J for which $J > I$ with respect to the lexicographic ordering on finite sequences of integers. These Sq^J's have the same degree $i_1 + \cdots + i_k$ as Sq^I, and since the number of monomials Sq^I of a fixed degree is finite, successive applications of the Adém relations eventually reduce any Sq^I to a sum of admissible monomials.

For odd p, elements of $\mathcal{A}_p$ are linear combinations of monomials $\beta^{\varepsilon_1}P^{i_1}\beta^{\varepsilon_2}P^{i_2}\cdots$ with each $\varepsilon_j = 0$ or 1. Such a monomial is **admissible** if $i_j \geq \varepsilon_{j+1} + pi_{j+1}$ for all j, which again means that no Adém relation can be applied to it. As with $\mathcal{A}_2$, the Adém relations suffice to reduce every monomial to a linear combination of admissible monomials, by the same argument as before but now using the lexicographic ordering on tuples $(\varepsilon_1 + pi_1, \varepsilon_2 + pi_2, \cdots)$.

Define the **excess** of the admissible monomial Sq^I to be $\sum_j (i_j - 2i_{j+1})$, the amount by which Sq^I exceeds being admissible. Similarly, the excess of $\beta^{\varepsilon_1}P^{i_1}\beta^{\varepsilon_2}P^{i_2}\cdots$ is defined to be $\sum_j 2(i_j - pi_{j+1}) - \varepsilon_{j+1}$.

As we explained at the beginning of this chapter, cohomology operations correspond to elements in the cohomology of Eilenberg-MacLane spaces. Here is a somewhat difficult but important theorem which will be proved in Volume II.

Theorem. *For each prime p, $H^*(K(\mathbb{Z}_p, n); \mathbb{Z}_p)$ is the free commutative algebra on the generators $\Theta(\iota_n)$ where $\iota_n \in H^n(K(\mathbb{Z}_p, n); \mathbb{Z}_p)$ is a generator and Θ ranges over all admissible monomials of excess less than n.*

Here 'free commutative algebra' means 'polynomial algebra' when $p = 2$ and 'polynomial algebra on even-dimensional generators tensor exterior algebra on odd-dimensional generators' when p is odd. We will say something about the rationale behind the 'excess less than n' condition in a moment.

For example, in the cases $n = 1, 2$ we have:

$$\begin{aligned}
K(\mathbb{Z}_2, 1): &\quad \mathbb{Z}_2[\iota] \\
K(\mathbb{Z}_p, 1): &\quad \Lambda_{\mathbb{Z}_p}[\iota] \otimes \mathbb{Z}_p[\beta\iota] \\
K(\mathbb{Z}_2, 2): &\quad \mathbb{Z}_2[\iota, Sq^1\iota, Sq^2Sq^1\iota, Sq^4Sq^2Sq^1\iota, \cdots] \\
K(\mathbb{Z}_p, 2): &\quad \mathbb{Z}_p[\iota, \beta P^1\beta\iota, \beta P^pP^1\beta\iota, \beta P^{p^2}P^pP^1\beta\iota, \cdots] \\
&\qquad\qquad \otimes\ \Lambda_{\mathbb{Z}_p}[\beta\iota, P^1\beta\iota, P^pP^1\beta\iota, P^{p^2}P^pP^1\beta\iota, \cdots]
\end{aligned}$$

The theorem says in particular that the admissible monomials in $\mathcal{A}_p$ are linearly independent, hence form a basis. For if some linear combination of admissible monomials were zero, then it would be zero when applied to the class ι_n, but if we choose n larger than the excess of each monomial in the linear combination, this would contradict the freeness of the algebra $H^*(K(\mathbb{Z}_p, n); \mathbb{Z}_p)$. (For a proof of linear independence using only techniques from the present chapter, see [Steenrod-Epstein].) Even thought the multiplicative structure of the Steenrod algebra is rather complicated, the Adém relations provide a way of performing calculations algorithmically by sytematically reducing all products to sums of admissible monomials.

Another consequence of the theorem is that all cohomology operations with $\mathbb{Z}_p$ coefficients are polynomials in the Sq^i's when $p = 2$ and polynomials in the P^i's and β when p is odd, in view of Proposition 4.86. We can also conclude that $\mathcal{A}_p$ consists precisely of all the $\mathbb{Z}_p$ cohomology operations which are stable, i.e., commute with suspension. For consider the map $\Sigma K(\mathbb{Z}_p, n) \to K(\mathbb{Z}_p, n+1)$ which pulls ι_{n+1} back to the suspension of ι_n. This map induces an isomorphism on homotopy groups π_i for $i \le 2n$ and a surjection for $i = 2n+1$, by Theorem 4.30, hence the same is true for homology and cohomology. Letting n go to infinity, the limit $\varprojlim \widetilde{H}^*(K(\mathbb{Z}_p, n); \mathbb{Z}_p)$ then exists in a strong sense. On the one hand, this limit is exactly the stable operations by Proposition 4.86 and the definition of a stable operation. On the other hand, the preceding theorem implies that this limit is exactly $\mathcal{A}_p$ since it says that all elements of $H^*(K(\mathbb{Z}_p, n); \mathbb{Z}_p)$ below dimension $2n$ are uniquely expressible as sums of admissible monomials applied to ι_n.

Now let us explain why the condition 'excess less than n' in the theorem is natural and optimal, at least in the case $p = 2$. For a monomial $Sq^I = Sq^{i_1}Sq^{i_2}\cdots$ the definition of the excess $e(I)$ can be rewritten as an equation $i_1 = e(I) + i_2 + i_3 + \cdots$. Thus if $e(I) > n$, we have $i_1 > |Sq^{i_2}Sq^{i_3}\cdots(\iota_n)|$ hence $Sq^I(\iota_n) = 0$. And if $e(I) = n$ then $Sq^I(\iota_n) = (Sq^{i_2}Sq^{i_3}\cdots(\iota_n))^2$ and either $Sq^{i_2}Sq^{i_3}\cdots$ has excess less than n or it has excess equal to n and we can repeat the process to write $Sq^{i_2}Sq^{i_3}\cdots(\iota_n) = (Sq^{i_3}\cdots(\iota_n))^2$, etc., eventually writing $Sq^I(\iota_n)$ as a power of an element $Sq^J(\iota_n)$ with $e(J) < n$, so that

$Sq^I(\iota_n)$ is already in the algebra generated by the elements $Sq^J(\iota_n)$ with $e(J) < n$. The situation for odd p is nearly the same. For an admissible monomial $P^I = \beta^{\varepsilon_1} P^{i_1} \beta^{\varepsilon_2} P^{i_2} \cdots$ of excess $e(I) > n$ we must have $P^I(\iota_n) = 0$, and if $e(I) = n$ then either $P^I(\iota_n)$ is a power $(P^J(\iota_n))^k$ with $e(J) < n$, or, if P^I begins with β, then $P^I(\iota_n) = \beta\big((P^J(\iota_n))^k\big)$ with $e(J) < n$. However, in the latter case there doesn't seem to be an *a priori* reason why the algebra generated by the classes $P^J(\iota_n)$ with $e(J) < n$ should be closed under the Bockstein operation.

There is another set of relations among Steenrod squares equivalent to the Adém relations and somewhat easier to remember:

$$\sum_j \tbinom{k}{j} Sq^{2n-k-j-1} Sq^{n+j} = 0$$

When $k = 0$ this is simply the relation $Sq^{2n-1}Sq^n = 0$, and the cases $k > 0$ are obtained from this via Pascal's triangle. For example, from the relation $Sq^7Sq^4 = 0$ we derive the further relations:

$$\begin{aligned}
Sq^7Sq^4 &= 0\\
Sq^6Sq^4 + Sq^7Sq^3 &= 0\\
Sq^5Sq^4 + Sq^7Sq^2 &= 0\\
Sq^4Sq^4 + Sq^5Sq^3 + Sq^6Sq^2 + Sq^7Sq^1 &= 0\\
Sq^3Sq^4 + Sq^7Sq^0 &= 0\\
Sq^2Sq^4 + Sq^3Sq^3 + Sq^6Sq^0 &= 0\\
Sq^1Sq^4 + Sq^3Sq^2 + Sq^5Sq^0 &= 0\\
Sq^0Sq^4 + Sq^1Sq^3 + Sq^2Sq^2 + Sq^3Sq^1 + Sq^4Sq^0 &= 0
\end{aligned}$$

These relations are not in simplest possible form. For example, $Sq^5Sq^3 = 0$ in the fourth row and $Sq^3Sq^2 = 0$ in the seventh row, instances of $Sq^{2n-1}Sq^n = 0$. For Steenrod powers there are similar relations $\sum_j \binom{k}{j} P^{pn-k-j-1}P^{n+j} = 0$ derived from the basic relation $P^{pn-1}P^n = 0$. We leave it to the interested reader to show that these relations follow from the Adém relations.

Constructing the Squares and Powers

Now we turn to the construction of the Steenrod squares and powers, and the proof of their basic properties including the Adém relations. As will be seen, this all hinges on the commutativity property of cup product and the fact that cohomology is maps into Eilenberg-MacLane spaces. The case $p = 2$ is in some ways simpler than the case p odd, so

in the first part of the development we will specialize p to 2 whenever there is a significant difference between the two cases.

For a space X with basepoint x_0, let $X^{\wedge p}$ denote the smash product $X\wedge\cdots\wedge X$ of p copies of X. There is a map $T: X^{\wedge p}\to X^{\wedge p}$, $T(x_1,\cdots,x_p)=(x_2,\cdots,x_p,x_1)$, permuting the factors cyclically. When $p=2$ this is just the transposition $(x_1,x_2)\mapsto(x_2,x_1)$. The map T generates an action of $\mathbb{Z}_p$ on $X^{\wedge p}$. There is also the standard action of $\mathbb{Z}_p$ on S^∞ viewed as the union of the unit spheres S^{2n-1} in $\mathbb{C}^n$, a generator of $\mathbb{Z}_p$ rotating each $\mathbb{C}$ factor through an angle $2\pi/p$, with quotient space an infinite-dimensional lens space L^∞, or $\mathbb{R}\mathrm{P}^\infty$ when $p=2$. On the product $S^\infty\times X^{\wedge p}$ there is then the diagonal action $g(s,x)=(g(s),g(x))$ for $g\in\mathbb{Z}_p$. Let $\Gamma(X)$ denote the orbit space $(S^\infty\times X^{\wedge p})/\mathbb{Z}_p$ of this diagonal action. This is the same as the Borel construction $S^\infty\times_{\mathbb{Z}_p}X^{\wedge p}$ described in §3.4h. The projection $S^\infty\times X^{\wedge p}\to S^\infty$ induces a projection $\pi:\Gamma(X)\to L^\infty$ with $\pi^{-1}(z)=X^{\wedge p}$ for all $z\in L^\infty$ since the action of $\mathbb{Z}_p$ on S^∞ is free. This projection $\Gamma(X)\to L^\infty$ is in fact a fiber bundle, though we shall not need this fact and so we leave the proof as an exercise. The $\mathbb{Z}_p$ action on $X^{\wedge p}$ fixes the basepoint $x_0\in X^{\wedge p}$, hence the inclusion $S^\infty\times\{x_0\}\hookrightarrow S^\infty\times X^{\wedge p}$ induces an inclusion $L^\infty\hookrightarrow\Gamma(X)$. The composition $L^\infty\hookrightarrow\Gamma(X)\to L^\infty$ is the identity, so in fiber bundle terminology this subspace $L^\infty\subset\Gamma(X)$ is a section of the bundle. Let $\Lambda(X)$ denote the quotient $\Gamma(X)/L^\infty$ obtained by collapsing the section L^∞ to a point. Note that the fibers $X^{\wedge p}$ in $\Gamma(X)$ are still embedded in the quotient $\Lambda(X)$ since each fiber meets the section L^∞ in a single point.

If we replace S^∞ by S^1 in these definition, we get subspaces $\Gamma^1(X)\subset\Gamma(X)$ and $\Lambda^1(X)\subset\Lambda(X)$. All these spaces have natural CW structures if X is a CW complex having x_0 as a 0-cell. Namely, L^∞ is given its standard CW structure with one cell in each dimension. This is lifted to a CW structure on S^∞ with p cells in each dimension, and then T freely permutes the product cells of $S^\infty\times X^{\wedge p}$ so there is induced a quotient CW structure on $\Gamma(X)$. The section $L^\infty\subset\Gamma(X)$ is a subcomplex, so the quotient $\Lambda(X)$ inherits a CW structure from $\Gamma(X)$. For example, if the n-skeleton of X is S^n with its usual CW structure, then the pn-skeleton of $\Lambda(X)$ is S^{pn} with its usual CW structure.

We remark also that Γ, Γ^1, Λ, and Λ^1 are functors: A map $f:(X,x_0)\to(Y,y_0)$ induces maps $\Gamma(f):\Gamma(X)\to\Gamma(Y)$, etc., in the evident way.

For brevity we shall denote $H^*(-;\mathbb{Z}_p)$ simply as $H^*(-)$. For $n>0$ let K_n denote a CW complex $K(\mathbb{Z}_p,n)$ with $(n-1)$-skeleton a point and n-skeleton S^n. Let $\iota\in H^n(K_n)$ be the fundamental class. It will be notationally convenient to regard an element $\alpha\in H^n(X)$ also as a map $\alpha:X\to K_n$ such that $\alpha^*(\iota)=\alpha$. Here we are assuming X is a CW complex.

From §3.2 we have a reduced external cup product $H^*(X)^{\otimes p} \to \widetilde{H}^*(X^{\wedge p})$ where $\widetilde{H}^*(X)^{\otimes p}$ denotes the p-fold tensor product of $\widetilde{H}^*(X)$ with itself. In particular, for each $\alpha \in H^n(X)$, $n > 0$, we have its p-fold external cup product $\alpha^{\otimes p} \in \widetilde{H}^*(X^{\wedge p})$. Our first task will be construct an element $\lambda(\alpha) \in H^{pn}(\Lambda(X))$ restricting to $\alpha^{\otimes p}$ in each fiber $X^{\wedge p} \subset \Lambda(X)$. The construction is based on on the following simple fact:

Lemma 4.92. *$T^*(\alpha^{\otimes p}) = \alpha^{\otimes p}$, where $T : X^{\wedge p} \to X^{\wedge p}$ cyclically permutes the factors.*

Proof: By definition, $X^{\wedge p} = X^p/Y$ where Y is the subspace of the p-fold product X^p consisting of points with at least one coordinate equal to the basepoint x_0. The isomorphism $\widetilde{H}^*(X^{\wedge p}) \approx \widetilde{H}^*(X)^{\otimes p}$ is the isomorphism $H^*(X^p, Y) \approx H^*(X, x_0)^{\otimes p}$ given by relative external cup product. In these terms, $\alpha^{\otimes p}$ is $\pi_1^*(\alpha) \smile \cdots \smile \pi_p^*(\alpha) \in H^*(X^p, Y)$ where $\pi_i : X^p \to X$ is projection onto the i^{th} factor, and $\alpha \in H^n(X, x_0)$. Then

$$T^*(\pi_1^*(\alpha) \smile \cdots \smile \pi_p^*(\alpha)) = \pi_2^*(\alpha) \smile \cdots \smile \pi_p^*(\alpha) \smile \pi_1^*(\alpha) = (-1)^{(p-1)n} \pi_1^*(\alpha) \smile \cdots \smile \pi_p^*(\alpha).$$

If p is odd the coefficient $(-1)^{(p-1)n}$ is $+1$, and if $p = 2$ we have $\mathbb{Z}_2$ coefficients so signs are irrelevant. Hence $T^*(\alpha^{\otimes p}) = \alpha^{\otimes p}$ in both cases. □

Taking $X = K_n$ and $\alpha = \iota$, the equation $T^*(\iota^{\otimes p}) = \iota^{\otimes p}$ means that there is a homotopy $\iota^{\otimes p} T \simeq \iota^{\otimes p} : K_n^{\wedge p} \to K_{pn}$. This homotopy defines a map $\Gamma^1(K_n) \to K_{pn}$ since $\Gamma^1(X) = I \times X^{\wedge p}/(0, x) \sim (1, T(x))$. The homotopy may be taken to be basepoint-preserving since K_{pn} is simply-connected if $n > 0$. Hence the map $\Gamma^1(K_n) \to K_{pn}$ has a quotient map $\lambda_1 : \Lambda^1(K_n) \to K_{pn}$. Since K_n is obtained from S^n by attaching cells of dimension greater than n, $\Lambda(K_n)$ is obtained from $\Lambda^1(K_n)$ by attaching cells of dimension greater than $pn + 1$. There are then no obstructions to extending λ_1 to a map $\lambda : \Lambda(K_n) \to K_{pn}$ since $\pi_i(K_{pn}) = 0$ for $i > pn$.

By construction, the restriction of λ to each fiber $K_n^{\wedge p}$ is homotopic to $\iota^{\otimes p}$. Note that this property determines λ uniquely, up to homotopy, since the restriction map $H^{pn}(\Lambda(K_n)) \to H^{pn}(K_n^{\wedge p})$ is injective, the pn-skeleton of $\Lambda(K_n)$ being contained in $K_n^{\wedge p}$. We shall have occasion to use this argument again in the proof, so we refer to it as the 'uniqueness argument.'

For an arbitrary class $\alpha \in H^n(X)$ let $\lambda(\alpha)$ be the composition $\Lambda(X) \xrightarrow{\Lambda(\alpha)} \Lambda(K_n) \xrightarrow{\lambda} K_{pn}$. This restricts to $\alpha^{\otimes p}$ in each fiber $X^{\wedge p}$ since $\Lambda(\alpha)$ restricts to $\alpha^{\wedge p}$ in each fiber.

Now we are ready to define some cohomology operations. There is an inclusion $L^\infty \times X \hookrightarrow \Gamma(X)$ as the quotient of the 'diagonal' embedding $S^\infty \times X \hookrightarrow S^\infty \times X^{\wedge p}$, $(s, x) \mapsto (s, x, \cdots, x)$. Composing with the quotient map $\Gamma(X) \to \Lambda(X)$ we get a map $\nabla : L^\infty \times X \to$

$\Lambda(X)$ inducing $\nabla^*: H^*(\Lambda(X)) \to H^*(L^\infty \times X) \approx H^*(L^\infty) \otimes H^*(X)$. For each $\alpha \in H^n(X)$ the element $\nabla^*(\lambda(\alpha)) \in H^{pn}(L^\infty \times X)$ may be written in the form

$$\nabla^*(\lambda(\alpha)) = \sum_i \omega_{(p-1)n-i} \otimes \theta_i(\alpha)$$

where ω_j is a generator of $H^j(L^\infty)$ and $\theta_i(\alpha) \in H^{n+i}(X)$. Thus θ_i increases dimension by i. When $p = 2$ there is no ambiguity about ω_j. For odd p we choose ω_1 to be the class dual to the 1-cell of L^∞ in its standard cell structure, then we take ω_2 to be the Bockstein $\beta\omega_1$ and we set $\omega_{2j} = \omega_2^j$ and $\omega_{2j+1} = \omega_1\omega_2^j$.

It is clear that θ_i is a cohomology operation since $\theta_i(\alpha) = \alpha^*(\theta_i(\iota))$. For $p = 2$ we set $Sq^i(\alpha) = \theta_i(\alpha)$. For odd p we shall show that $\theta_i = 0$ unless $i = 2k(p-1)$ or $2k(p-1)+1$. The operation P^k will be defined to be a constant times $\theta_{2k(p-1)}$, and $\theta_{2k(p-1)+1}$ will be a constant times βP^k, for β the mod p Bockstein.

Theorem 4.93. *The operations Sq^i satisfy the properties (1)-(7).*

Proof: We have already observed that the θ_i's are cohomology operations, so property (1) holds. Because cohomology operations cannot decrease dimension, $\theta_i = 0$ for $i < 0$. Also, $\theta_i = 0$ for $i > (p-1)n$ since the factor $\omega_{(p-1)n-i}$ vanishes in this case. The basic property that $\lambda(\alpha)$ restricts to $\alpha^{\otimes p}$ in each fiber implies that $\theta_{(p-1)n}(\alpha) = \alpha^p$ since $\omega_0 = 1$. This gives property (5).

To show the additivity property (2) we will show $\lambda(\alpha+\beta) = \lambda(\alpha) + \lambda(\beta)$ using the following diagram:

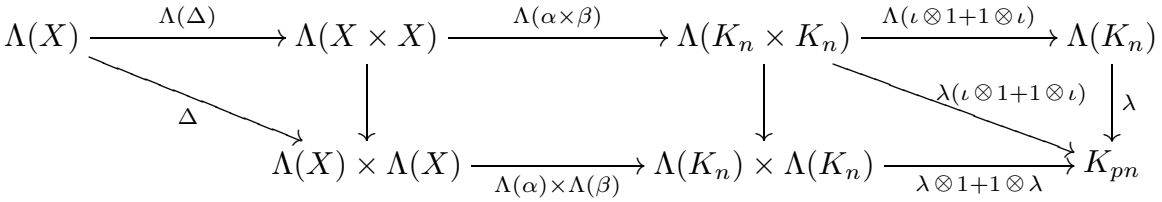

Here Δ is a generic symbol for diagonal maps $x \mapsto (x, x)$. These relate tensor to cup products via $\Delta^*(\varphi \otimes \psi) = \varphi \smile \psi$. The two unlabeled vertical maps are induced by $(s, x_1, y_1, \cdots, x_p, y_p) \mapsto (s, x_1, \cdots, x_p, s, y_1, \cdots, y_p)$. It is obvious that the central square commutes, and also the left and upper right triangles. To show that the remaining triangle commutes it suffices to show this when we restrict to the fiber $(K_n \times K_n)^{\wedge p}$ by the 'uniqueness argument' used earlier, since $(K_n \times K_n)^{\wedge p}$ contains the pn-skeleton of $\Lambda(K_n \times K_n)$. On the fiber the two routes to K_{pn} are $\iota^{\otimes p} \otimes 1 + 1 \otimes \iota^{\otimes p}$ and $(\iota \otimes 1 + 1 \otimes \iota)^{\otimes p}$. These are equal since raising to the p^{th} power is an additive homomorphism of the field $\mathbb{Z}_p$.

The composition $\Lambda(X) \to \Lambda(K_n)$ is $\Lambda(\alpha+\beta)$, so the composition $\Lambda(X) \to K_{pn}$ across the top of the diagram is $\lambda(\alpha+\beta)$. Going from $\Lambda(X)$ to K_{pn} across the bottom of the diagram gives $\lambda(\alpha)+\lambda(\beta)$ since the second and third steps of this composition give $\lambda(\alpha)\otimes 1 + 1\otimes\lambda(\beta)$. Thus $\lambda(\alpha+\beta) = \lambda(\alpha)+\lambda(\beta)$, which implies that each θ_i is an additive homomorphism.

Next we turn to the Cartan formula. For any prime p we will show that $\lambda(\alpha\smile\beta) = (-1)^{p(p-1)mn/2}\lambda(\alpha)\lambda(\beta)$ for $m=|\alpha|$ and $n=|\beta|$. This implies (3) when $p=2$ since if we let $\omega=\omega_1$, hence $\omega_j=\omega^j$, then

$$\begin{aligned}\sum_i Sq^i(\alpha\smile\beta)\otimes\omega^{n+m-i} &= \nabla^*\big(\lambda(\alpha\smile\beta)\big) = \nabla^*\big(\lambda(\alpha)\smile\lambda(\beta)\big)\\ &= \nabla^*\big(\lambda(\alpha)\big)\smile\nabla^*\big(\lambda(\beta)\big)\\ &= \sum_j Sq^j(\alpha)\otimes\omega^{n-j}\smile\sum_k Sq^k(\beta)\otimes\omega^{m-k}\\ &= \sum_i\Big(\sum_{j+k=i} Sq^j(\alpha)\smile Sq^k(\beta)\Big)\otimes\omega^{n+m-i}\end{aligned}$$

To obtain the formula $\lambda(\alpha\smile\beta)=\pm\lambda(\alpha)\smile\lambda(\beta)$ we use a diagram rather like the previous one:

$$\begin{array}{ccccccc}
\Lambda(X) & \xrightarrow{\Lambda(\Delta)} & \Lambda(X\wedge X) & \xrightarrow{\Lambda(\alpha\wedge\beta)} & \Lambda(K_m\wedge K_n) & \xrightarrow{\Lambda(\iota_m\otimes\iota_n)} & \Lambda(K_{m+n}) \\
& \searrow^{\Delta} & \downarrow & & \downarrow & \searrow^{\lambda(\iota_m\otimes\iota_n)} & \downarrow{\scriptstyle\lambda} \\
& & \Lambda(X)\wedge\Lambda(X) & \xrightarrow[\Lambda(\alpha)\wedge\Lambda(\beta)]{} & \Lambda(K_m)\wedge\Lambda(K_n) & \xrightarrow[\lambda(\iota_m)\otimes\lambda(\iota_n)]{} & K_{pm+pn}
\end{array}$$

The composition $\Lambda(X)\to K_{pm+pn}$ going across the top of the diagram is $\lambda(\alpha\smile\beta)$ since the map $\Lambda(X)\to\Lambda(K_{m+n})$ is $\Lambda(\alpha\smile\beta)$. The composition $\Lambda(X)\wedge\Lambda(X)\to K_{pm+pn}$ is $\lambda(\alpha)\otimes\lambda(\beta)$ so the composition $\Lambda(X)\to K_{pm+pn}$ across the bottom of the diagram is $\lambda(\alpha)\smile\lambda(\beta)$. Except for the lower triangle, the commutativity of the diagram is obvious from the definitions. It remains to see that the lower triangle commutes up to the sign $(-1)^{p(p-1)mn/2}$. Since $(K_m\wedge K_n)^{\wedge p}$ is the $(pm+pn)$-skeleton of $\Lambda(K_m\wedge K_n)$, restriction to this fiber is injective on H^{pm+pn}. On this fiber the two routes around the triangle give $(\iota_m\otimes\iota_n)^{\otimes p}$ and $\iota_m^{\otimes p}\otimes\iota_n^{\otimes p}$. These differ by a permutation which is the product of $(p-1)+(p-2)+\cdots+1 = p(p-1)/2$ transpositions of adjacent factors. Since ι_m and ι_n have dimensions m and n, this permutation introduces a sign $(-1)^{p(p-1)mn/2}$ by the commutativity property of cup product. This finishes the verification of the Cartan formula when $p=2$.

Before proceeding further we need to make an explicit calculation to show that Sq^0 is the identity on $H^1(S^1)$. Viewing S^1 as the one-point compactification of $\mathbb{R}$, with the point at infinity as the basepoint, the 2-sphere $S^1 \wedge S^1$ becomes the one-point compactification of $\mathbb{R}^2$. The map $T : S^1 \wedge S^1 \to S^1 \wedge S^1$ then corresponds to reflecting $\mathbb{R}^2$ across the line $x = y$, so after a rotation of coordinates this becomes reflection of S^2 across the equator. Hence $\Gamma^1(S^1)$ is obtained from the shell $I \times S^2$ by identifying its inner and outer boundary spheres via a reflection across the equator. The diagonal $\mathbb{R}\mathrm{P}^1 \times S^1 \subset \Gamma^1(S^1)$ is a torus, obtained from the equatorial annulus $I \times S^1 \subset I \times S^2$ by identifying the two ends via the identity map since the equator is fixed by the reflection. This $\mathbb{R}\mathrm{P}^1 \times S^1$ represents the same element of $H_2(\Gamma^1(S^1); \mathbb{Z}_2)$ as the fiber sphere $S^1 \wedge S^1$ since the upper half of the shell is a 3-cell whose mod 2 boundary in $\Gamma^1(S^1)$ is the union of these two surfaces.

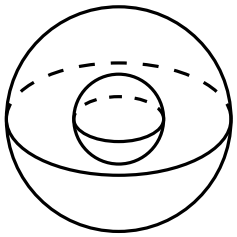

Let $\alpha \in H^1(S^1)$ be a generator, and consider the element $\nabla^*(\lambda(\alpha))$ in $H^2(\mathbb{R}\mathrm{P}^\infty \times S^1) \approx Hom(H_2(\mathbb{R}\mathrm{P}^\infty \times S^1; \mathbb{Z}_2), \mathbb{Z}_2)$. A basis for $H_2(\mathbb{R}\mathrm{P}^\infty \times S^1; \mathbb{Z}_2)$ is represented by $\mathbb{R}\mathrm{P}^2 \times \{x_0\}$ and $\mathbb{R}\mathrm{P}^1 \times S^1$. A cocycle representing $\nabla^*(\lambda(\alpha))$ takes the value 0 on $\mathbb{R}\mathrm{P}^2 \times \{x_0\}$ since $\mathbb{R}\mathrm{P}^\infty \times \{x_0\}$ collapses to a point in $\Lambda(S^1)$ and $\lambda(\alpha)$ lies in $H^2(\Lambda(S^1))$. On $\mathbb{R}\mathrm{P}^1 \times S^1$, $\nabla^*(\lambda(\alpha))$ takes the value 1 since when $\lambda(\alpha)$ is pulled back to $\Gamma(S^1)$ it takes the same value on the homologous cycles $\mathbb{R}\mathrm{P}^1 \times S^1$ and $S^1 \wedge S^1$, namely 1 by the defining property of $\lambda(\alpha)$ since $\alpha \otimes \alpha \in H^2(S^1 \wedge S^1)$ is a generator. Thus $\nabla^*(\lambda(\alpha)) = \omega_1 \otimes \alpha$ and hence $Sq^0(\alpha) = \alpha$ by the definition of Sq^0.

We use this calculation to prove that Sq^i commutes with the suspension σ, where σ is defined by $\sigma(\alpha) = \varepsilon \otimes \alpha \in H^*(S^1 \wedge X)$ for ε a generator of $H^1(S^1)$ and $\alpha \in H^*(X)$. We have just seen that $Sq^0(\varepsilon) = \varepsilon$. By (5), $Sq^1(\varepsilon) = \varepsilon^2 = 0$ and $Sq^i(\varepsilon) = 0$ for $i > 1$. The Cartan formula then gives $Sq^i(\sigma(\alpha)) = Sq^i(\varepsilon \otimes \alpha) = \sum_j Sq^j(\varepsilon) \otimes Sq^{i-j}(\alpha) = \varepsilon \otimes Sq^i(\alpha) = \sigma(Sq^i(\alpha))$.

From this it follows that Sq^0 is the identity on $H^n(S^n)$ for all $n > 0$. Since S^n is the n-skeleton of K_n, this implies that Sq^0 is the identity on the fundamental class ι_n, hence Sq^0 is the identity on all positive dimensional classes.

Property (7) is proved similarly: Sq^1 coincides with the Bockstein β on the generator $\omega \in H^1(\mathbb{R}\mathrm{P}^2)$ since both equal ω^2. Hence $Sq^1 = \beta$ on the iterated suspensions of ω, and the n-fold suspension of $\mathbb{R}\mathrm{P}^2$ is the $(n+2)$-skeleton of K_{n+1}. □

Theorem 4.94. *The Adém relations hold for Steenrod squares.*

Proof: The idea is to imitate the construction of $\Lambda(X)$ using $\mathbb{Z}_p \times \mathbb{Z}_p$ in place of $\mathbb{Z}_p$. The Adém relations will then arise from the symmetry of $\mathbb{Z}_p \times \mathbb{Z}_p$ interchanging the two

factors.

The group $\mathbb{Z}_p \times \mathbb{Z}_p$ acts on $S^\infty \times S^\infty$ via the product action $(g,h)(s,t) = (g(s),h(t))$, with quotient $L^\infty \times L^\infty$. There is also an action of $\mathbb{Z}_p \times \mathbb{Z}_p$ on $X^{\wedge p^2}$, obtained by writing points of $X^{\wedge p^2}$ as p^2-tuples (x_{ij}) with subscripts i and j varying from 1 to p, and then letting the first $\mathbb{Z}_p$ act on the first subscript and the second $\mathbb{Z}_p$ act on the second. Factoring out the diagonal action of $\mathbb{Z}_p \times \mathbb{Z}_p$ on $S^\infty \times S^\infty \times X^{\wedge p^2}$ gives a quotient space $\Gamma_2(X)$. This projects to $L^\infty \times L^\infty$ with a section, and collapsing the section gives $\Lambda_2(X)$. The fibers of the projection $\Lambda_2(X) \to L^\infty \times L^\infty$ are $X^{\wedge p^2}$ since the action of $\mathbb{Z}_p \times \mathbb{Z}_p$ on $S^\infty \times S^\infty$ is free. We could also obtain $\Lambda_2(X)$ from the product $S^\infty \times S^\infty \times X^{p^2}$ by first collapsing the subspace of points having at least one X coordinate equal to the basepoint x_0, then factoring out the $\mathbb{Z}_p \times \mathbb{Z}_p$ action.

It is interesting to compare $\Lambda_2(X)$ with $\Lambda(\Lambda(X))$. The latter space is the quotient of $S^\infty \times (S^\infty \times X^p)^p$ in which one first identifies all points having at least one X coordinate equal to x_0, then one factors out by an action of the wreath product $\mathbb{Z}_p \wr \mathbb{Z}_p$, the group of order p^{p+1} defined by a split exact sequence $0 \to \mathbb{Z}_p^p \to \mathbb{Z}_p \wr \mathbb{Z}_p \to \mathbb{Z}_p \to 0$ with conjugation by the quotient group $\mathbb{Z}_p$ given by cyclic permutations of the p $\mathbb{Z}_p$ factors of $\mathbb{Z}_p^p$. In the coordinates $(s, t_1, x_{11}, \cdots, x_{1p}, \cdots, t_p, x_{p1}, \cdots, x_{pp})$ the i^{th} factor $\mathbb{Z}_p$ of $\mathbb{Z}_p^p$ acts in the block $(t_i, x_{i1}, \cdots, x_{ip})$, and the quotient $\mathbb{Z}_p$ acts by cyclic permutation of the index i and by rotation in the s coordinate. There is a natural map $\Lambda_2(X) \to \Lambda(\Lambda(X))$ induced by $(s,t,x_{11},\cdots,x_{1p},\cdots,x_{p1},\cdots,x_{pp}) \mapsto (s,t,x_{11},\cdots,x_{1p},\cdots,t,x_{p1},\cdots,x_{pp})$. In $\Lambda_2(X)$ one is factoring out by the action of $\mathbb{Z}_p \times \mathbb{Z}_p$. This is the subgroup of $\mathbb{Z}_p \wr \mathbb{Z}_p$ obtained by restricting the action of the quotient $\mathbb{Z}_p$ on $\mathbb{Z}_p^p$ to the diagonal subgroup $\mathbb{Z}_p \subset \mathbb{Z}_p^p$, where this action becomes trivial so that one has the direct product $\mathbb{Z}_p \times \mathbb{Z}_p$.

Since it suffices to prove that the Adém relations hold on the class $\iota \in H^n(K_n)$, we take $X = K_n$. There is a map $\lambda_2 : \Lambda_2(K_n) \to K_{p^2 n}$, unique up to homotopy, with the property that its restriction to each fiber is $\iota^{\otimes p} r$. This is constructed by the same method used to construct λ. One starts with a map representing $\iota^{\otimes p^2}$ in a fiber, then extends over the part of $\Lambda_2(K_n)$ projecting to the 1-skeleton of $L^\infty \times L^\infty$ using the commutativity property of cup product, then one extends inductively over higher-dimensional cells of $\Lambda_2(K_n)$ using the fact that $K_{p^2 n}$ is an Eilenberg-MacLane space.

Now we consider the diagram:

$$\begin{array}{ccccc} L^\infty \times L^\infty \times K_n & \xrightarrow{\nabla_2} & \Lambda_2(K_n) & \xrightarrow{\lambda_2} & K_{p^2 n} \\ \downarrow{\scriptstyle 1\!\!1 \times \nabla} & & \downarrow & \nearrow{\scriptstyle \lambda(\lambda)} & \\ L^\infty \times \Lambda(K_n) & \xrightarrow{\nabla} & \Lambda(\Lambda(K_n)) & & \end{array} \qquad \begin{array}{ccc} s,t,x & \longrightarrow & s,t,x_{11},x_{12},x_{21},x_{22} \\ \downarrow & & \downarrow \\ s,t,x_1,x_2 & \longrightarrow & s,t_1,x_{11},x_{12},t_2,x_{21},x_{22} \end{array}$$

Here the unlabelled map is the one defined above. The other three maps in the square are

also induced by various sorts of diagonal maps, indicated by the coordinates in the right part of the figure for the case $p = 2$. It is clear that the square commutes. Commutativity of the triangle, up to homotopy, follows from the uniqueness property of λ_2.

The element $\nabla_2^*\lambda_2^*(\iota)$ may be written in the form $\sum_{r,s} \omega_r \otimes \omega_s \otimes \varphi_{rs}$. Switching the two $\mathbb{Z}_p$ factors of $\mathbb{Z}_p \times \mathbb{Z}_p$ gives the relation $\varphi_{rs} = (-1)^{rs+p(p-1)n/2}\varphi_{sr}$ since we get a sign $(-1)^{rs}$ from permuting ω_r and ω_s, and interchanging the subscripts on coordinates x_{ij} involves $p(p-1)/2$ transpositions, one for each pair (i,j) with $1 \le i < j \le p$.

On the other hand, if we pull ι back across the bottom of the diagram, we obtain the element $\nabla_2^*\lambda_2^*(\iota)$ in the form

$$\sum_i \omega_{(p-1)pn-i} \otimes \theta_i\big(\sum_j \omega_{(p-1)n-j} \otimes \theta_j(\iota)\big) = \sum_{i,j} \omega_{(p-1)pn-i} \otimes \theta_i\big(\omega_{(p-1)n-j} \otimes \theta_j(\iota)\big)$$

Now we specialize to $p = 2$. The Cartan formula converts this last expression into

$$\sum_{i,j,k} \omega^{2n-i} \otimes Sq^k(\omega^{n-j}) \otimes Sq^{i-k}Sq^j(\iota) = \sum_{i,j,k} \tbinom{n-j}{k}\omega^{2n-i} \otimes \omega^{n-j+k} \otimes Sq^{i-k}Sq^j(\iota)$$

To write this expression more symmetrically with respect to the two ω terms, let $n-j+k = 2n-\ell$. Then we get

$$\sum_{i,j,\ell} \tbinom{n-j}{n+j-\ell}\omega^{2n-i} \otimes \omega^{2n-\ell} \otimes Sq^{i+\ell-n-j}Sq^j(\iota)$$

In view of the symmetry property of φ_{rs}, switching i and ℓ in this formula leaves it unchanged. Hence we get the relation

$$(*) \qquad \sum_j \tbinom{n-j}{n+j-\ell} Sq^{i+\ell-n-j}Sq^j(\iota) = \sum_j \tbinom{n-j}{n+j-i} Sq^{i+\ell-n-j}Sq^j(\iota)$$

This holds for all n, i, and ℓ, and the idea is to choose these numbers so that the left side of this equation has only one nonzero term. Given integers r and s, let $n = 2^r - 1 + s$ and $\ell = n + s$, so that $\binom{n-j}{n+j-\ell} = \binom{2^r-1-(j-s)}{j-s}$. If r is sufficiently large, this will be 0 unless $j = s$. This is because the dyadic expansion of $2^r - 1$ consists entirely of 1's, so the expansion of $2^r - 1 - (j-s)$ will have 0's in the positions where the expansion of $j-s$ has 1's, hence these positions contribute factors of $\binom{0}{1} = 0$ to $\binom{2^r-1-(j-s)}{j-s}$. Thus with n and ℓ chosen as above, the relation $(*)$ becomes

$$Sq^iSq^s(\iota) = \sum_j \tbinom{2^r-1+s-j}{2^r-1+s+j-i} Sq^{i+s-j}Sq^j(\iota) = \sum_j \tbinom{2^r+s-j-1}{i-2j} Sq^{i+s-j}Sq^j(\iota)$$

where the latter equality comes from the general relation $\binom{x}{y} = \binom{x}{x-y}$.

The final step is to show that $\binom{2^r+s-j-1}{i-2j} = \binom{s-j-1}{i-2j}$ if $i < 2s$. Both of these binomial coefficients are zero if $i < 2j$, and if $i \geq 2j$ then we have $2j \leq i < 2s$, so $j < s$, hence $s-j-1 \geq 0$. The term 2^r then makes no difference in $\binom{2^r+s-j-1}{i-2j}$ if r is large since this 2^r contributes only a single 1 to the dyadic expansion of $2^r+s-j-1$, far to the left of all the nonzero entries in the dyadic expansions of $s-j-1$ and $i-2j$.

This gives the Adém relations for the classes ι of dimension $n = 2^r - 1 + s$ with r large. This implies the relations hold for all classes of these dimensions. Since we can suspend repeatedly to make any class have dimension of this form, the Adém relations must hold for all cohomology classes. □

Steenrod Powers

Now we assume the prime p is odd.

Lemma 4.95. *$\theta_i = 0$ unless $i = 2k(p-1)$ or $2k(p-1)+1$.*

Proof: The group of automorphisms of $\mathbb{Z}_p$ is cyclic of order $p-1$, and we can represent a generator by an automorphism which is multiplication by some number r. Define a map $\varphi : S^\infty \times X^{\wedge p} \to S^\infty \times X^{\wedge p}$ by $\varphi(s,(x_i)) = (s^r,(x_{ri}))$ where subscripts are taken mod p and s^r means raise each coordinate of s, regarded as a unit vector in $\mathbb{C}^\infty$, to the r^{th} power and renormalize the resulting vector to have unit length. Then if γ is a generator of the $\mathbb{Z}_p$ action on $S^\infty \times X^{\wedge p}$, we have $\varphi(\gamma(s,(x_i))) = \varphi(e^{2\pi/p}s,(x_{i+1})) = (e^{2r\pi/p}s^r,(x_{ri+r})) = \gamma^r(\varphi(s,(x_i)))$. This means that φ takes orbits to orbits, hence induces maps $\varphi : \Gamma(X) \to \Gamma(X)$ and $\varphi : \Lambda(X) \to \Lambda(X)$. Restricting to the first coordinate, there is also an induced map $\varphi : L^\infty \to L^\infty$. Taking $X = K_n$, these maps fit into a diagram

$$\begin{array}{ccccc} L^\infty \times K_n & \xrightarrow{\nabla} & \Lambda(K_n) & \xrightarrow{\lambda} & K_{pn} \\ {\scriptstyle \varphi\times 1\!\!1}\downarrow & & {\scriptstyle \varphi}\downarrow & \nearrow_{\lambda} & \\ L^\infty \times K_n & \xrightarrow{\nabla} & \Lambda(K_n) & & \end{array}$$

The square obviously commutes, and the triangle commutes up to homotopy since it suffices to verify this in the fiber $X^{\wedge p}$, and here the map φ is an even permutation of the coordinates since it has even order $p-1$, so φ preserves $\iota^{\otimes p}$.

Commutativity of the diagram means that $\sum_i \omega_{(p-1)n-i} \otimes \theta_i(\iota)$ is invariant under $\varphi^* \otimes 1\!\!1$. The map φ induces multiplication by r in $\pi_1(L^\infty)$, hence also in $H_1(L^\infty)$ and $H^1(L^\infty;\mathbb{Z}_p)$, sending ω_1 to $r\omega_1$. Since ω_2 was chosen to be the Bockstein of ω_1 it is also multiplied by r. We chose r to have multiplicative order $p-1$ mod p, so $\varphi^*(\omega_\ell) = \omega_\ell$

only when the total number of ω_1 and ω_2 factors in ω_ℓ is a multiple of $p-1$, i.e., when ℓ has the form $2k(p-1)$ or $2k(p-1)-1$. Thus in the formula $\sum_i \omega_{(p-1)n-i}\otimes\theta_i(\iota)$ we must have $\theta_i(\iota)=0$ unless i is congruent to 0 or 1 mod $2(p-1)$. □

Since $\theta_0 : H^n(X)\to H^n(X)$ is a cohomology operation which preserves dimension, it must be defined by a coefficient homomorphism $\mathbb{Z}_p\to\mathbb{Z}_p$, multiplication by some $a_n\in\mathbb{Z}_p$. We claim that these a_n's satisfy

$$a_{m+n} = (-1)^{p(p-1)mn/2}a_m a_n \qquad \text{and} \qquad a_n = (-1)^{p(p-1)n(n-1)/4}a_1$$

To see this, recall the formula $\lambda(\alpha\smile\beta) = (-1)^{p(p-1)mn/2}\lambda(\alpha)\lambda(\beta)$ for $\alpha\in H^m$ and $\beta\in H^n$. From the definition of the θ_i's this implies $\theta_0(\alpha\smile\beta) = (-1)^{p(p-1)mn/2}\theta_0(\alpha)\theta_0(\beta)$, which gives the first part of the claim. The second part follows from this by induction on n.

Lemma 4.96. *$a_1 = \pm m!$ where $p = 2m+1$.*

Proof: It suffices to compute $\theta_0(\alpha)$ where α is any nonzero one-dimensional class, so the simplest thing is to choose α to be a generator of $H^1(S^1)$, say a generator coming from a generator of $H^1(S^1;\mathbb{Z})$; this determines α up to a sign. Since $H^i(S^1)=0$ for $i>1$, we have $\theta_i(\alpha)=0$ for $i>0$, so the defining formula for $\theta_0(\alpha)$ has the form $\nabla^*(\lambda(\alpha)) = \omega_{p-1}\otimes\theta_0(\alpha) = a_1\omega_{p-1}\otimes\alpha$ in $H^p(L^\infty\times S^1)$. To compute a_1 there is no harm in replacing L^∞ by a finite-dimensional lens space here, say L^p, the p-skeleton of L^∞. Thus we may restrict the bundle $\Lambda(S^1)\to L^\infty$ to a bundle $\Lambda^p(S^1)\to L^p$ with the same fibers $(S^1)^{\wedge p} = S^p$. We regard S^1 as the one-point compactification of $\mathbb{R}$ with basepoint the added point at infinity, and then $(S^1)^{\wedge p}$ becomes the one-point compactification $\widehat{\mathbb{R}^p}$ of $\mathbb{R}^p$ with $\mathbb{Z}_p$ acting by permuting the coordinates of $\mathbb{R}^p$ cyclically, preserving the origin and the point at infinity. This action defines the bundle $\Gamma^p(S^1)\to L^p$ with fibers $S^p = \widehat{\mathbb{R}^p}$, containing a zero-section and a section at infinity, and $\Lambda^p(S^1)$ is obtained by collapsing the section at infinity. We can also describe $\Lambda^p(S^1)$ as the one-point compactification of the complement of the section at infinity in $\Gamma^p(S^1)$, since the base space L^p is compact. This complement of the section at infinity is a bundle $E\to L^p$ with fibers $\mathbb{R}^p$, defined like $\Gamma^p(S^1)$ but using $\mathbb{R}^p$ in place of $\widehat{\mathbb{R}^p}$. In general, the one-point compactification of a fiber bundle E over a compact base space with fibers $\mathbb{R}^n$ is called the **Thom space** $T(E)$ of the bundle, and a class in $H^n(T(E))$ which restricts to a generator of H^n of the one-point compactification $\widehat{\mathbb{R}^n}$ of each fiber $\mathbb{R}^n$ is called a **Thom class**. In our situation, $\lambda(\alpha)$ is such a Thom class.

Our first task is to describe subbundles $E_0, E_1, \cdots, E_m$ of E, where E_0 has fiber $\mathbb{R}$ and the other E_j's have fiber $\mathbb{R}^2$, so $p = 2m+1$. The bundle E comes from the linear

transformation $T:\mathbb{R}^p \to \mathbb{R}^p$ permuting the coordinates cyclically. We claim there is a decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ with V_0 one-dimensional and the other V_j's two-dimensional, such that $T(V_j) = V_j$ for all j, with $T|V_0$ the identity and $T|V_j$ a rotation by the angle $2\pi j/p$ for $j > 0$. Thus T defines an action of $\mathbb{Z}_p$ on V_j and we can define E_j just as E was defined, as the quotient $(S^p \times V_j)/\mathbb{Z}_p$ with respect to the diagonal action.

An easy way to get the decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ is to regard $\mathbb{R}^p$ as a module over the principal ideal domain $\mathbb{R}[t]$ by setting $tv = T(v)$ for $v \in \mathbb{R}^p$. Then $\mathbb{R}^p$ is isomorphic as a module to the module $\mathbb{R}[t]/(t^p - 1)$ since T permutes coordinates cyclically; this amounts to identifying the standard basis vectors $v_1, \cdots, v_p$ in $\mathbb{R}^p$ with $1, t, \cdots, t^{p-1}$. The polynomial $t^p - 1$ factors over $\mathbb{C}$ into the linear factors $t - e^{2\pi ij/p}$ for $j = 0, \cdots, p-1$. Combining complex conjugate factors, this gives a factorization over $\mathbb{R}$, $t^p - 1 = (t-1)\prod_{1\le j\le m}(t^2 - 2(\cos\varphi_j)t + 1)$, where $\varphi_j = 2\pi j/p$. These are distinct monic irreducible factors, so the module $\mathbb{R}[t]/(t^p - 1)$ splits as $\mathbb{R}[t]/(t-1)\bigoplus_{1\le j\le m}\mathbb{R}[t]/(t^2 - 2(\cos\varphi_j)t + 1)$ by the basic structure theory of modules over a principal ideal domain. This translates into a decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ with $T(V_j) \subset V_j$. Here V_0 corresponds to $\mathbb{R}[t]/(t-1) \approx \mathbb{R}$ with t acting as the identity, and V_j for $j > 0$ corresponds to $\mathbb{R}[t]/(t^2 - 2(\cos\varphi_j)t + 1)$. The latter module is isomorphic to $\mathbb{R}^2$ with t acting as rotation by the angle φ_j since the characteristic polynomial of this rotation is readily computed to be $t^2 - 2(\cos\varphi_j)t + 1$, hence this rotation satisfies $t^2 - 2(\cos\varphi_j)t + 1 = 0$ so there is a module homomorphism $\mathbb{R}[t]/(t^2 - 2(\cos\varphi_j)t + 1) \to \mathbb{R}^2$ which is obviously an isomorphism.

From the decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ and the action of T on each factor we can see that the only vectors fixed by T are those in the line V_0. The vectors $(x, \cdots, x)$ are fixed by T, so V_0 must be this diagonal line.

Next we compute Thom classes for the bundles E_j. This is easy for E_0 which is the product $L^p \times \mathbb{R}$, so the projection $E_0 \to \mathbb{R}$ one-point compactifies to a map $T(E_0) \to S^1$ and we can pull back the chosen generator $\alpha \in H^1(S^1)$ to a Thom class for E_0. The other E_j's have 2-dimensional fibers, which we now view as $\mathbb{C}$ rather than $\mathbb{R}^2$. Just as E_j is the quotient of $S^p \times \mathbb{C}$ via the identifications $(v, z) \sim (e^{2\pi i/p}v, e^{2\pi ij/p}z)$, we can define a bundle $\overline{E}_j \to \mathbb{C}\mathrm{P}^m$ with fiber $\mathbb{C}$ by the identifications $(v, z) \sim (\lambda v, \lambda^j z)$ for $\lambda \in S^1 \subset \mathbb{C}$. We then have the left half of the following commutative diagram

$$\begin{array}{ccccc} E_j & \xrightarrow{\widetilde{q}} & \overline{E}_j & \xrightarrow{\widetilde{f}} & \overline{E}_1 \\ \downarrow & & \downarrow & & \downarrow \\ L^p & \xrightarrow{q} & \mathbb{C}\mathrm{P}^m & \xrightarrow{f} & \mathbb{C}\mathrm{P}^m \end{array}$$

where the quotient map $\widetilde{q}$ restricts to a homeomorphism on each fiber. The maps $\widetilde{f}$ and f are induced by the map $S^p \times \mathbb{C} \to S^p \times \mathbb{C}$ sending (v, z) to (v^j, z) where v^j means

raise each coordinate of v to the j^{th} power and then rescale to get a vector of unit length. The map $\widetilde{f}$ is well-defined since equivalent pairs $(v,z)\sim(\lambda v,\lambda^j z)$ in $\overline{E}_j$ are carried to pairs (v^j,z) and $(\lambda^j v^j,\lambda^j z)$ which are equivalent in $\overline{E}_1$.

Since both $\widetilde{q}$ and $\widetilde{f}$ restrict to homeomorphisms in each fiber, they extend to maps of Thom spaces which pull a Thom class for $\overline{E}_1$ back to Thom classes for $\overline{E}_j$ and E_j. To construct a Thom class for $\overline{E}_1$, observe that the Thom space $T(\overline{E}_1)$ is homeomorphic to $\mathbb{C}\mathrm{P}^{m+1}$, namely, view the sphere $S^p=S^{2m+1}$ as the unit sphere in $\mathbb{C}^{m+1}$, and then the inclusion $S^p\times\mathbb{C}\hookrightarrow\mathbb{C}^{m+1}\times\mathbb{C}=\mathbb{C}^{m+2}$ induces a map $g\colon\overline{E}_1\to\mathbb{C}\mathrm{P}^{m+1}$ since the equivalence relation defining $\overline{E}_1$ is $(v,z)\sim(\lambda v,\lambda z)$ for $\lambda\in S^1$. It is evident that g is a homeomorphism onto the complement of the point $[0,\cdots,0,1]$ in $\mathbb{C}\mathrm{P}^{m+1}$, so sending the point at infinity in $T(\overline{E}_1)$ to $[0,\cdots,0,1]$ gives an extension of g to a homeomorphism $T(\overline{E}_1)\approx\mathbb{C}\mathrm{P}^{m+1}$. Under this homeomorphism the one-point compactifications of the fibers of $\overline{E}_1$ correspond to the 2-spheres S^2_v consisting of $[0,\cdots,0,1]$ and the points $[v,z]\in\mathbb{C}\mathrm{P}^{m+1}$ with fixed $v\in S^p$ and varying $z\in\mathbb{C}$. Each S^2_v is a $\mathbb{C}\mathrm{P}^1$ in $\mathbb{C}\mathrm{P}^{m+1}$ equivalent to the standard $\mathbb{C}\mathrm{P}^1$ under a homeomorphism of $\mathbb{C}\mathrm{P}^{m+1}$ coming from a linear isomorphism of $\mathbb{C}^{m+2}$, so a generator γ of $H^2(\mathbb{C}\mathrm{P}^{m+1})$ is a Thom class, restricting to a generator of $H^2(S^2_v)$ for each v. We choose γ to be the $\mathbb{Z}_p$-reduction of a generator of $H^2(\mathbb{C}\mathrm{P}^{m+1};\mathbb{Z})$, so γ is determined up to a sign. (This construction of Thom classes for the bundles E_j works with $\mathbb{Z}$ coefficients as well as $\mathbb{Z}_p$ coefficients.)

A slightly different view of Thom classes will be useful. For the bundle $E\to L^p$, for example, we have isomorphisms

$$\begin{aligned}
\widetilde{H}^*(T(E)) &\approx H^*(T(E),\infty) \qquad \text{where } \infty \text{ is the compactification point}\\
&\approx H^*(T(E),T(E)-L^p) \qquad \text{where } L^p \text{ is embedded in } T(E) \text{ as the zero-section, so that } T(E)-L^p \text{ deformation retracts to } \infty\\
&\approx H^*(E,E-L^p) \qquad \text{by excision.}
\end{aligned}$$

Thus we can view a Thom class as lying in $H^*(E,E-L^p)$, and similarly for the bundles E_j.

We have projections $\pi_j\colon E\to E_j$ coming from the projections $V_0\oplus V_1\oplus\cdots\oplus V_m\to V_j$ in fibers. If $\tau_j\in H^*(E_j,E_j-L^p)$ denotes the Thom class constructed above, then we have the pullback $\pi_j^*(\tau_j)\in H^*(E,E-\pi_j^{-1}(L^p))$, and the cup product $\prod_j\pi_j^*(\tau_j)\in H^*(E,E-L^p)$ is a Thom class for E as one sees by applying Lemma 3.10 repeatedly in each fiber. Under the isomorphism $H^*(E,E-L^p)\approx\widetilde{H}^*(T(E))$, the class $\prod_j\pi_j^*(\tau_j)$ corresponds to $\pm\lambda(\alpha)$ since both classes restrict to $\pm\alpha^{\otimes p}$ in each fiber $S^p\subset T(E)$ and $\lambda(\alpha)$ is uniquely determined by its restriction to fibers.

Now we can finish the proof. The class $\nabla^*(\lambda(\alpha))$ is obtained by restricting $\lambda(\alpha) \in H^p(T(E))$ to the diagonal $T(E_0)$, then pulling back to $L^p \times S^1$ via the quotient map $L^p \times S^1 \to T(E_0)$ which collapses the section at infinity to a point. Restricting $\prod_j \pi_j^*(\tau_j)$ to $H^p(E_0, E_0 - L^p) \approx H^p(T(E_0))$ gives $\tau_0 \smile e_1 \smile \cdots \smile e_m$ where $e_j \in H^2(E_0)$ is the image of τ_j under $H^2(E_j, E_j - L^p) \to H^2(E_j) \approx H^2(L^p) \approx H^2(E_0)$, these last two isomorphisms coming from including L^p in E_j and E_0 as the zero-section, to which they deformation retract. To compute e_j, we use the diagram

$$\begin{array}{ccccc}
H^2(E_j, E_j - L^p) & \xleftarrow{\widetilde{q}^*} & H^2(\overline{E}_j, \overline{E}_j - \mathbb{C}\mathrm{P}^m) & \xleftarrow{\widetilde{f}^*} & H^2(\overline{E}_1, \overline{E}_1 - \mathbb{C}\mathrm{P}^m) \\
\downarrow & & \downarrow & & \downarrow \\
H^2(E_j) & \xleftarrow{\widetilde{q}^*} & H^2(\overline{E}_j) & \xleftarrow{\widetilde{f}^*} & H^2(\overline{E}_1, \overline{E}_1) \\
\downarrow\approx & & \downarrow\approx & & \downarrow\approx \\
H^2(L^p) & \xleftarrow{q^*} & H^2(\mathbb{C}\mathrm{P}^m) & \xleftarrow{f^*} & H^2(\mathbb{C}\mathrm{P}^m)
\end{array}$$

The Thom class for $\overline{E}_1$ lies in the upper-right group. Following this class across the top of the diagram and then down to the lower-left corner gives the element e_j. To compute e_j we take the alternate route through the lower-right corner of the diagram. The image of the Thom class for $\overline{E}_1$ in the lower-right $H^2(\mathbb{C}\mathrm{P}^m)$ is the generator γ since $T(\overline{E}_1) = \mathbb{C}\mathrm{P}^{m+1}$. The map f^* is multiplication by j since f has degree j on $\mathbb{C}\mathrm{P}^1 \subset \mathbb{C}\mathrm{P}^m$. And $q^*(\gamma) = \pm\omega_2$ since q restricts to a homeomorphism on the 2-cell of L^p in the CW structure defined in Example 7 of §2.2. Thus $e_j = \pm j\omega_2$, and so $\tau_0 \smile e_1 \smile \cdots \smile e_m = \pm m!\tau_0 \smile \omega_2^m = \pm m!\tau_0 \smile \omega_{p-1}$. Since τ_0 was the pullback of α via the projection $T(E_0) \to S^1$, when we pull τ_0 back to $L^p \times S^1$ via ∇ we get $1 \otimes \alpha$, so $\tau_0 \smile e_1 \smile \cdots \smile e_m$ pulls back to $\pm m!\omega_{p-1} \otimes \alpha$. Hence $a_1 = \pm m!$. □

The lemma implies in particular that a_n is not zero in $\mathbb{Z}_p$, so it makes sense to define:

$$P^i(\alpha) = (-1)^i a_n^{-1} \theta_{2i(p-1)}(\alpha) \qquad \text{for } \alpha \in H^n(X)$$

The factor a_n^{-1} guarantees that P^0 is the identity. The factor $(-1)^i$ is inserted in order to make $P^i(\alpha) = \alpha^p$ if $|\alpha| = 2i$, as we show next. We know that $\theta_{2i(p-1)}(\alpha) = \alpha^p$, so what must be shown is that $(-1)^i a_{2i}^{-1} = 1$, or equivalently, $a_{2i} = (-1)^i$.

To do this we need a small number theory fact: $\big((p-1)/2\big)!^2 \equiv (-1)^{(p+1)/2} \mod p$. For consider the multiplicative group $\mathbb{Z}_p^*$ of nonzero elements of $\mathbb{Z}_p$. Half of the elements of this group are $1, 2, \cdots, (p-1)/2$ and the other half are the negatives of these, so $(p-1)! \equiv \big((p-1)/2\big)!^2 (-1)^{(p-1)/2}$. On the other hand, $(p-1)!$ is the product of all the elements of $\mathbb{Z}_p^*$, and this group is cyclic of even order, so the product of all its elements is the

unique element of order 2, which is -1, since all the other nontrivial elements cancel their inverses in this product. Thus $(p-1)! \equiv -1$ and hence $\big((p-1)/2\big)!^2 \equiv (-1)^{(p+1)/2} \mod p$.

Therefore

$$\begin{aligned} a_{2i} &= (-1)^{p(p-1)2i(2i-1)/4}\big((p-1)/2\big)!^{2i} \\ &= (-1)^{p[(p-1)/2]i(2i-1)}(-1)^{i(p+1)/2} \\ &= (-1)^{i(p-1)/2}(-1)^{i(p+1)/2} \qquad \text{since } p \text{ and } 2i-1 \text{ are odd} \\ &= (-1)^{ip} = (-1)^i \qquad \text{since } p \text{ is odd.} \end{aligned}$$

Theorem 4.97. *The operations P^i satisfy the properties (1)-(6) and the Adém relations.*

Proof: The first two properties, naturality and additivity, along with the fact that $P^i(\alpha) = 0$ if $2i > |\alpha|$, are inherited from the θ_i's. Property (6) and the other half of (5) have just been shown above. For the Cartan formula we have, for $\alpha \in H^m$ and $\beta \in H^n$, $\lambda(\alpha \smile \beta) = (-1)^{p(p-1)mn/2}\lambda(\alpha)\lambda(\beta)$ and hence

$$\sum_i \omega_{(p-1)(m+n)-i} \otimes \theta_i(\alpha \smile \beta) = \\ (-1)^{p(p-1)mn/2}\Big(\sum_j \omega_{(p-1)m-j} \otimes \theta_j(\alpha)\Big)\Big(\sum_k \omega_{(p-1)n-k} \otimes \theta_k(\beta)\Big)$$

Recall that $\omega_{2r} = \omega_2^r$ and $\omega_{2r+1} = \omega_1\omega_2^r$, with $\omega_1^2 = 0$. Therefore terms with i even on the left side of the equation can only come from terms with j and k even on the right side. This leads to the second equality in the following sequence:

$$\begin{aligned} P^i(\alpha \smile \beta) &= (-1)^i a_{m+n}^{-1}\theta_{2i(p-1)}(\alpha \smile \beta) \\ &= (-1)^i a_{m+n}^{-1}(-1)^{p(p-1)mn/2}\sum_j \theta_{2(i-j)(p-1)}(\alpha)\theta_{2j(p-1)}(\beta) \\ &= \sum_j (-1)^{i-j} a_m^{-1}\theta_{2(i-j)(p-1)}(\alpha)(-1)^j a_n^{-1}\theta_{2j(p-1)}(\beta) \\ &= \sum_j P^{i-j}(\alpha)P^j(\beta) \end{aligned}$$

Property (4), the invariance of P^i under suspension, follows from the Cartan formula just as in the case $p = 2$, using the fact that P^0 is the only P^i which can be nonzero on one-dimensional classes, by (5).

It remains to prove the Adém relations for Steenrod powers. We will need a Bockstein calculation:

Lemma 4.98. $\beta\theta_{2k} = -\theta_{2k+1}$.

Proof: Let us first reduce the problem to showing that $\beta\nabla^*(\lambda(\iota)) = 0$. If we compute $\beta\nabla^*(\lambda(\iota))$ using the product formula for β we get

$$\beta\Big(\sum_i \omega_{(p-1)n-i}\otimes\theta_i(\iota)\Big) = \sum_i\Big(\beta\omega_{(p-1)n-i}\otimes\theta_i(\iota) + (-1)^i\omega_{(p-1)n-i}\otimes\beta\theta_i(\iota)\Big)$$

Since $\beta\omega_{2j-1} = \omega_{2j}$ and $\beta\omega_{2j} = 0$, the terms with $i = 2k$ and $i = 2k+1$ give, respectively,

$$\sum_k \omega_{(p-1)n-2k}\otimes\beta\theta_{2k}(\iota) \quad\text{and}\quad \sum_k \omega_{(p-1)n-2k}\otimes\theta_{2k+1}(\iota) - \sum_k \omega_{(p-1)n-2k-1}\otimes\beta\theta_{2k+1}(\iota)$$

The coefficient of $\omega_{(p-1)n-2k}$ is $\beta\theta_{2k}(\iota) + \theta_{2k+1}(\iota)$, and this must vanish if the whole sum is zero since we are in the tensor product $H^*(L^\infty)\otimes H^*(K_n)$. Thus $\beta\theta_{2k}(\iota) = -\theta_{2k+1}(\iota)$ and hence $\beta\theta_{2k}(\alpha) = -\theta_{2k+1}(\alpha)$ for all α. Note that $\beta\theta_{2k} = -\theta_{2k+1}$ implies $\beta\theta_{2k+1} = 0$ since $\beta^2 = 0$.

In order to show that $\beta\nabla^*(\lambda(\iota)) = 0$ we first compute $\beta\lambda(\iota)$. We may assume K_n has a single n-cell and a single $(n+1)$-cell, attached by a map of degree p. Let φ and ψ be the cellular cochains assigning the value 1 to the n-cell and the $(n+1)$-cell, respectively, so $\delta\varphi = p\psi$. In $K_n^{\wedge p}$ we then have

$$(*)\qquad \delta(\varphi^{\otimes p}) = \sum_i (-1)^{in}\varphi^{\otimes i}\otimes\delta\varphi\otimes\varphi^{\otimes(p-i-1)} = p\sum_i (-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$$

where the tensor notation means cellular cross product, e.g., $\varphi^{\otimes p}$ is the cellular cochain dual to the np-cell $e^n\times\cdots\times e^n$ of $K_n^{\wedge p}$. The formula $(*)$ holds also in $\Lambda(K_n)$ since the latter space has only one $(np+1)$-cell not in $K_n^{\wedge p}$, with cellular boundary zero. Namely, this cell is the product of the 1-cell of L^∞ and the np-cell of $K_n^{\wedge p}$ with one end of this product attached to the np-cell by the identity map and the other end by the cyclic permutation T, which has degree $+1$ since p is odd, so these two terms in the boundary of this cell cancel, and there are no other terms since the rest of the attachment of this cell is at the basepoint.

The formula $(*)$ says that $\sum_i (-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$ represents $\beta\lambda(\iota)$. Via the quotient map $\Gamma(K_n)\to\Lambda(K_n)$, the class $\lambda(\iota)$ pulls back to a class $\gamma(\iota)$ with $\beta\gamma(\iota)$ also represented by $\sum_i (-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$. To see what happens when we pull $\beta\gamma(\iota)$ back to $\beta\nabla^*(\lambda(\iota))$ via the inclusion $L^\infty\times K_n\hookrightarrow\Gamma(K_n)$, consider the following commutative diagram:

$$\begin{array}{ccccc} H^*(\Gamma(K_n)) & \xrightarrow{\pi^*} & H^*(S^\infty\times K_n^p) & \xrightarrow{\tau} & H^*(\Gamma(K_n)) \\ \downarrow & & \downarrow & & \downarrow \\ H^*(L^\infty\times K_n) & \xrightarrow{\pi^*} & H^*(S^\infty\times K_n) & \xrightarrow{\tau} & H^*(L^\infty\times K_n) \end{array}$$

In the left-hand square the maps π^* are induced by the covering space projections $\pi : S^\infty \times K_n^{\wedge p} \to \Gamma(K_n)$ and $\pi : S^\infty \times K_n \to L^\infty \times K_n$ arising from the free $\mathbb{Z}_p$ actions. The vertical maps are induced by the diagonal inclusion $S^\infty \times K \hookrightarrow S^\infty \times K_n^{\wedge p}$. The maps τ are the transfer homomorphisms defined in §3.4h. Recall the definition: If $\pi : \widetilde{X} \to X$ is an m-sheeted covering space, a chain map $C_*(X) \to C_*(\widetilde{X})$ is defined by sending a singular simplex $\sigma : \Delta^k \to X$ to the sum of its m lifts to $\widetilde{X}$, and τ is the induced map on cohomology. The key property of τ is that $\tau\pi^* : H^*(X) \to H^*(X)$ is multiplication by m, for any choice of coefficient group, since when we project the m lifts of $\sigma : \Delta^k \to X$ back to X we get $m\sigma$. When X is a CW complex and $\widetilde{X}$ is given the lifted CW structure, then τ can also be defined in cellular cohomology by the same procedure.

The class $\beta\gamma(\iota)$ is in the image of $\tau : H^*(S^\infty \times K_n^{\wedge p}) \to H^*(\Gamma(K_n))$ since it is represented by $\sum_i (-1)^{in} \varphi^{\otimes i} \otimes \psi \otimes \varphi^{\otimes(p-i-1)} = \sum_i T^i(\psi \otimes \varphi^{\otimes(p-1)}) = \tau(1 \otimes \psi \otimes \varphi^{\otimes(p-1)})$ where $T : K_n^{\wedge p} \to K_n^{\wedge p}$ permutes the factors cyclically.

Since $\beta\gamma(\iota)$ is in the image of the upper τ in the diagram, the image of $\beta\gamma(\iota)$ in $H^*(L^\infty \times K_n)$, which is $\nabla^*(\beta\lambda(\iota))$, is in the image of the lower τ since the right-hand square commutes. The map π^* in the lower row is obviously onto since $\widetilde{H}^*(S^\infty) = 0$, so $\nabla^*(\beta\lambda(\iota))$ is in the image of the composition $\tau\pi^*$ across the bottom of the diagram. But this composition is multiplication by p, which is zero for $\mathbb{Z}_p$ coefficients, so $\beta\nabla^*(\lambda(\iota)) = \nabla^*(\beta\lambda(\iota)) = 0$. □

The derivation of the Adém relations now follows the pattern for the case $p = 2$. Setting $p = 2m+1$ in our earlier formula for $\nabla_2^*\lambda_2^*(\iota)$ gives $\sum_{i,j} \omega_{2mpn-i} \otimes \theta_i(\omega_{2mn-j} \otimes \theta_j(\iota))$. The only nonzero θ_i's are $\theta_{2i(p-1)} = (-1)^i a_n P^i$ and $\theta_{2i(p-1)+1} = -\beta\theta_{2i(p-1)}$ so we have

$$\begin{aligned}
\sum_{i,j} \omega_{2mpn-i} \otimes \theta_i(\omega_{2mn-j} \otimes \theta_j(\iota)) = \\
&\sum_{i,j} (-1)^{i+j} a_{2mn} a_n \omega_{2m(pn-2i)} \otimes P^i(\omega_{2m(n-2j)} \otimes P^j(\iota)) \\
-\ &\sum_{i,j} (-1)^{i+j} a_{2mn} a_n \omega_{2m(pn-2i)} \otimes P^i(\omega_{2m(n-2j)-1} \otimes \beta P^j(\iota)) \\
-\ &\sum_{i,j} (-1)^{i+j} a_{2mn} a_n \omega_{2m(pn-2i)-1} \otimes \beta P^i(\omega_{2m(n-2j)} \otimes P^j(\iota)) \\
+\ &\sum_{i,j} (-1)^{i+j} a_{2mn} a_n \omega_{2m(pn-2i)-1} \otimes \beta P^i(\omega_{2m(n-2j)-1} \otimes \beta P^j(\iota))
\end{aligned}$$

Factoring out the nonzero constant $a_{2mn}a_n$ and applying the Cartan formula to expand

the P^i terms gives

$$\begin{aligned}
&\sum_{i,j,k}(-1)^{i+j}\tbinom{m(n-2j)}{k}\omega_{2m(pn-2i)}\otimes\omega_{2m(n-2j+2k)}\otimes P^{i-k}P^j(\iota)\\
-\ &\sum_{i,j,k}(-1)^{i+j}\tbinom{m(n-2j)-1}{k}\omega_{2m(pn-2i)}\otimes\omega_{2m(n-2j+2k)-1}\otimes P^{i-k}\beta P^j(\iota)\\
-\ &\sum_{i,j,k}(-1)^{i+j}\tbinom{m(n-2j)}{k}\omega_{2m(pn-2i)-1}\otimes\omega_{2m(n-2j+2k)}\otimes \beta P^{i-k}P^j(\iota)\\
+\ &\sum_{i,j,k}(-1)^{i+j}\tbinom{m(n-2j)-1}{k}\omega_{2m(pn-2i)-1}\otimes\omega_{2m(n-2j+2k)}\otimes P^{i-k}\beta P^j(\iota)\\
-\ &\sum_{i,j,k}(-1)^{i+j}\tbinom{m(n-2j)-1}{k}\omega_{2m(pn-2i)-1}\otimes\omega_{2m(n-2j+2k)-1}\otimes \beta P^{i-k}\beta P^j(\iota)
\end{aligned}$$

using the fact that $P^k(\omega_{2r}) = \binom{r}{k}\omega_{2r+2k(p-1)}$ and $P^k(\omega_{2r+1}) = \binom{r}{k}\omega_{2r+2k(p-1)+1}$, as we showed earlier in the section.

Letting $\ell = mn + j - k$, so that $n - 2j + 2k = pn - 2\ell$, the first of the five preceding summations becomes

$$\sum_{i,j,\ell}(-1)^{i+j}\tbinom{m(n-2j)}{mn+j-\ell}\omega_{2m(pn-2i)}\otimes\omega_{2m(pn-2\ell)}\otimes P^{i+\ell-mn-j}P^j(\iota)$$

and similarly for the other four summations. Now we bring in the symmetry property $\varphi_{rs} = (-1)^{rs+mnp}\varphi_{sr}$. Of the five summations, only the first has both ω terms with even subscripts, so this term by itself must be symmetric with respect to switching i and ℓ, up to a sign which will be $+$ if we choose n to be even, as we will do. This gives the relation:

$$(1)\qquad \sum_j(-1)^{i+j}\tbinom{m(n-2j)}{mn+j-\ell}P^{i+\ell-mn-j}P^j(\iota) = \sum_j(-1)^{\ell+j}\tbinom{m(n-2j)}{mn+j-i}P^{i+\ell-mn-j}P^j(\iota)$$

Similarly, the second, third, and fourth summations involve ω's with subscripts of opposite parity, yielding the relation:

$$\begin{aligned}
(2)\qquad &\sum_j(-1)^{i+j}\tbinom{m(n-2j)-1}{mn+j-\ell}P^{i+\ell-mn-j}\beta P^j(\iota) =\\
&\sum_j(-1)^{\ell+j}\tbinom{m(n-2j)}{mn+j-i}\beta P^{i+\ell-mn-j}P^j(\iota) - \sum_j(-1)^{\ell+j}\tbinom{m(n-2j)-1}{mn+j-i}P^{i+\ell-mn-j}\beta P^j(\iota)
\end{aligned}$$

The relations (1) and (2) will yield the two Adém relations, so we won't need to consider the relation arising from the fifth summation.

To get the first Adém relation from (1) we choose n and ℓ so that the left side of (1) has only one term, namely we take $n = 2(1 + p + \cdots + p^{r-1}) + 2s$ and $\ell = mn + s$ for given integers r and s. Then

$$\tbinom{m(n-2j)}{mn+j-\ell} = \tbinom{p^r-1-(p-1)(j-s)}{j-s}$$

and if r is large, this binomial coefficient is 1 if $j = s$ and 0 otherwise since if the rightmost nonzero digit in the p-adic expansion of the 'denominator' $j - s$ is x, then the corresponding digit of the 'numerator' $(p-1)[(1 + p + \cdots + p^{r-1}) - (j - s)]$ is obtained by reducing $(p-1)(1-x) \mod p$, giving $x - 1$, and $\binom{x-1}{x} = 0$. Then (1) becomes

$$\begin{aligned}
(-1)^{i+s}P^iP^s(\iota) &= \sum_j (-1)^{\ell+j}\tbinom{m(n-2j)}{mn+j-i}P^{i+s-j}P^j(\iota) \\
\text{or}\qquad P^iP^s(\iota) &= \sum_j (-1)^{i+j}\tbinom{m(n-2j)}{mn+j-i}P^{i+s-j}P^j(\iota) \qquad \text{since} \quad \ell \equiv s \mod 2 \\
&= \sum_j (-1)^{i+j}\tbinom{m(n-2j)}{i-pj}P^{i+s-j}P^j(\iota) \qquad \text{since} \quad \tbinom{x}{y} = \tbinom{x}{x-y} \\
&= \sum_j (-1)^{i+j}\tbinom{p^r-1-(p-1)(s-j)}{i-pj}P^{i+s-j}P^j(\iota)
\end{aligned}$$

If r is large and $i < ps$, the term p^r in the binomial coefficient can be omitted since we may assume $i \geq pj$, hence $j < s$, so $-1 + (p-1)(s-j) \geq 0$ and the p^r has no effect on the binomial coefficient if r is large. This shows the first Adém relation holds for the class ι, and the general case follows as in the case $p = 2$.

To get the second Adém relation we choose $n = 2p^r + 2s$ and $\ell = mn + s$. Reasoning as before, the left side of (2) then reduces to $(-1)^{i+s}P^i\beta P^s(\iota)$ and (2) becomes

$$\begin{aligned}
P^i\beta P^s(\iota) = &\sum_j (-1)^{i+j}\tbinom{(p-1)(p^r+s-j)}{i-pj}\beta P^{i+s-j}P^j(\iota) \\
&- \sum_j (-1)^{i+j}\tbinom{(p-1)(p^r+s-j)-1}{i-pj-1}P^{i+s-j}\beta P^j(\iota)
\end{aligned}$$

This time the term p^r can be omitted if r is large and $i \leq ps$. □

Exercises

1. Determine all cohomology operations $H^1(X;\mathbb{Z}) \to H^n(X;\mathbb{Z})$, $H^2(X;\mathbb{Z}) \to H^n(X;\mathbb{Z})$, and $H^1(X;\mathbb{Z}_p) \to H^n(X;\mathbb{Z}_p)$ for p prime.

2. Use cohomology operations to show that the spaces $(S^1 \times \mathbb{C}\mathrm{P}^\infty)/(S^1 \times \{x_0\})$ and $S^3 \times \mathbb{C}\mathrm{P}^\infty$ are not homotopy equivalent.

3. Since there is a fiber bundle $S^2 \to \mathbb{C}\mathrm{P}^3 \to \mathbb{H}\mathrm{P}^2$, one might ask whether there is an analogous bundle $S^4 \to \mathbb{H}\mathrm{P}^3 \to \mathbb{O}\mathrm{P}^2$. Use Steenrod powers for the prime 3 to show that such a bundle cannot exist. [The Gysin sequence can be used to determine the map on cohomology induced by the bundle projection $\mathbb{H}\mathrm{P}^3 \to \mathbb{O}\mathrm{P}^2$.]

4. Show there is no fiber bundle $S^7 \to S^{23} \to \mathbb{O}\mathrm{P}^2$. [Compute the cohomology ring of the mapping cone of the projection $S^{23} \to \mathbb{O}\mathrm{P}^2$ via Poincaré duality or the Thom isomorphism.]

5. Show that the subalgebra of $\mathcal{A}_2$ generated by Sq^i for $i \leq 2$ has dimension 8 as a vector space over $\mathbb{Z}_2$, with multiplicative structure encoded in the following diagram, where diagonal lines indicate left-multiplication by Sq^1 and horizontal lines indicate left-multiplication by Sq^2.

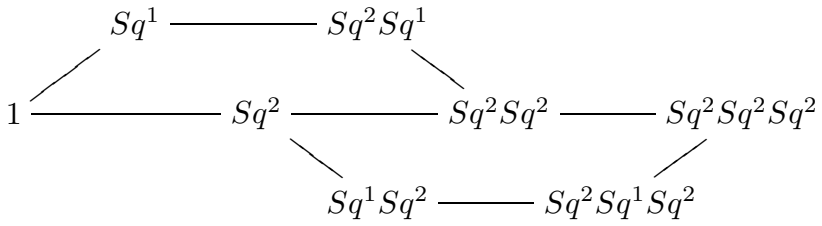

Chapter 4
Homotopy Theory

Homotopy theory begins with the homotopy groups $\pi_n(X)$, which are the natural higher-dimensional analogs of the fundamental group. These higher homotopy groups have certain formal similarities with homology groups. For example, $\pi_n(X)$ turns out to be always abelian for $n \geq 2$, and there are relative homotopy groups which fit into a long exact sequence just like the long exact sequence of homology groups. However, the higher homotopy groups are much harder to compute than either homology groups or the fundamental group, due to the fact that neither the excision property for homology nor van Kampen's theorem for π_1 hold for higher homotopy groups.

In spite of these computational difficulties, homotopy groups are of great theoretical significance. One reason for this is Whitehead's theorem that a map between CW complexes which induces isomorphisms on all homotopy groups is a homotopy equivalence. The stronger statement that two CW complexes with isomorphic homotopy groups are homotopy equivalent is usually false, however. One of the rare cases when a CW complex does have its homotopy type uniquely determined by its homotopy groups is when it has just a single nontrivial homotopy group. Such spaces, known as Eilenberg-MacLane spaces, turn out to play a fundamental role in algebraic topology for a variety of reasons. Perhaps the most important is their close connection with cohomology: Cohomology classes in a CW complex correspond bijectively with homotopy classes of maps from the complex into an Eilenberg-MacLane space. Thus cohomology has a strictly homotopy-theoretic interpretation, and there is an analogous but more subtle homotopy-theoretic interpretation of homology, explained in §4.4f.

A more elementary and direct connection between homotopy and homology is the Hurewicz theorem, asserting that the first nonzero homotopy group $\pi_n(X)$ of a simply-connected space X is isomorphic to the first nonzero homology group $\widetilde{H}_n(X)$. This result, along with its relative version, is one of the cornerstones of algebraic topology.

Though the excision property does not always hold for homotopy groups, in some important special cases there is a range of dimensions in which it does hold. This leads to the idea of 'stable homotopy groups,' the beginning of 'stable homotopy theory.' Perhaps the major unsolved problem in algebraic topology is the computation of the stable homotopy groups of spheres. Near the end of §4.2 we give some tables of known calculations which give a clear indication of the complexity of the problem.

Included in §4.3 is a brief introduction to fiber bundles, which generalize covering spaces and play a somewhat analogous role for higher homotopy groups. It would easily

be possible to devote a whole book to the subject of fiber bundles, particularly the special case of vector bundles, but here we use fiber bundles only to provide a few basic examples and to motivate their more flexible homotopy-theoretic generalization, fibrations, which are the main subject of §4.3. Among other things, fibrations allow one to describe, in theory at least, how the homotopy type of an arbitrary CW complex is built up from its homotopy groups by an inductive procedure of forming 'twisted products' of Eilenberg-MacLane spaces. This is the notion of a Postnikov tower. In favorable cases, including all simply-connected CW complexes, the data which determines a homotopy type can also be described, as a sequence of cohomology classes. If these are all zero, one has a product of Eilenberg-MacLane spaces, and otherwise not. Unfortunately these are cohomology classes in rather complicated spaces in general, so this is not a practical way of classifying homotopy types, but it is useful for various more theoretical purposes.

1. Homotopy Groups

Let I^n be the n-dimensional unit cube, the product of n copies of the interval $[0,1]$. The boundary ∂I^n of I^n is the subspace consisting of points with at least one coordinate equal to 0 or 1. For a space X with basepoint $x_0 \in X$, define $\pi_n(X, x_0)$ to be the set of homotopy classes of maps $f : (I^n, \partial I^n) \to (X, x_0)$, where homotopies f_t are also required to satisfy $f_t(\partial I^n) = x_0$. The definition extends to the case $n = 0$ by taking I^0 to be a point and ∂I^0 to be empty, so $\pi_0(X, x_0)$ is just the set of path-components of X.

When $n \geq 2$, a sum operation in $\pi_n(X, x_0)$, generalizing the composition operation in π_1, is defined by:

$$(f+g)(s_1, s_2, \cdots, s_n) = \begin{cases} f(2s_1, s_2, \cdots, s_n), & s_1 \in [0, 1/2] \\ g(2s_1 - 1, s_2, \cdots, s_n), & s_1 \in [1/2, 1] \end{cases}$$

This is clearly well-defined on homotopy classes. Since only the first coordinate is involved in the sum operation, the same arguments as for π_1 show that $\pi_n(X, x_0)$ is a group, with the constant map $I^n \to x_0$ as the zero element and inverses given by $-f(s_1, s_2, \cdots, s_n) = f(1 - s_1, s_2, \cdots, s_n)$. The additive notation for the group operation is used because $\pi_n(X, x_0)$ is abelian for $n \geq 2$. Namely, $f + g \simeq g + f$ via the homotopy indicated in the figures below which show how f and g slide around each other in the coordinates s_1 and s_2, the other coordinates being fixed.

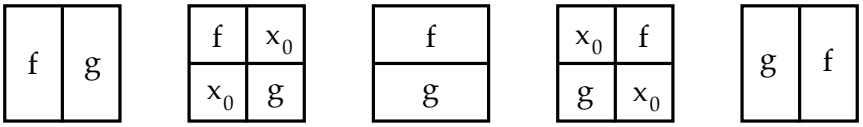

Maps $(I^n, \partial I^n) \to (X, x_0)$ are the same as maps of the quotient $I^n/\partial I^n = S^n$ to X taking the basepoint $s_0 = \partial I^n/\partial I^n$ to x_0. This means that we can also view $\pi_n(X, x_0)$ as homotopy classes of maps $(S^n, s_0) \to (X, x_0)$, where homotopies are through maps $(S^n, s_0) \to (X, x_0)$. In this interpretation of $\pi_n(X, x_0)$, the sum $f + g$ is the composition $S^n \xrightarrow{c} S^n \vee S^n \xrightarrow{f \vee g} X$ where c collapses the equator $S^{n-1} \subset S^n$ to a point and we choose the basepoint $s_0 \in S^{n-1}$.

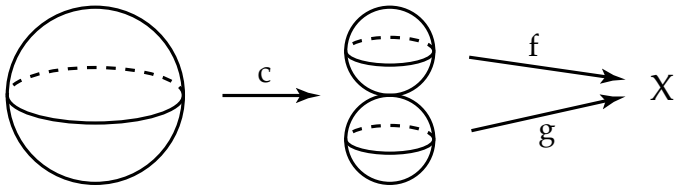

We will show shortly that if X is path-connected, different choices of the basepoint x_0 always produce isomorphic groups $\pi_n(X, x_0)$, just as for π_1. In this case one often writes $\pi_n(X)$ for $\pi_n(X, x_0)$.

Perhaps the simplest noncontractible spaces are spheres, so let us look at some of the known calculations of the groups $\pi_i(S^n)$ shown in the table below, extracted from [Toda62].

$\pi_i(S^n)$

		$i \longrightarrow$											
		1	2	3	4	5	6	7	8	9	10	11	12
n	1	$\mathbb{Z}$	0	0	0	0	0	0	0	0	0	0	0
$\downarrow$	2	0	$\mathbb{Z}$	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{12}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_3$	$\mathbb{Z}_{15}$	$\mathbb{Z}_2$	$\mathbb{Z}_2 \times \mathbb{Z}_2$
	3	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{12}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_3$	$\mathbb{Z}_{15}$	$\mathbb{Z}_2$	$\mathbb{Z}_2 \times \mathbb{Z}_2$
	4	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z} \times \mathbb{Z}_{12}$	$\mathbb{Z}_2 \times \mathbb{Z}_2$	$\mathbb{Z}_2 \times \mathbb{Z}_2$	$\mathbb{Z}_{24} \times \mathbb{Z}_3$	$\mathbb{Z}_{15}$	$\mathbb{Z}_2$
	5	0	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{30}$
	6	0	0	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	0	$\mathbb{Z}$	$\mathbb{Z}_2$
	7	0	0	0	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	0	0
	8	0	0	0	0	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	0

This is an intriguing mixture of pattern and chaos. The dominant feature is the zeros below the diagonal, and we shall show in Corollary 4.8 that $\pi_i(S^n) = 0$ for $i < n$. From the first row one would guess also that $\pi_i(S^1) = 0$ for all $i > 1$. This is indeed true and will follow from Proposition 4.1.

The coincidences in the second and third rows can hardly be overlooked, and in fact are the case $n = 1$ of isomorphisms $\pi_i(S^{2n}) \approx \pi_{i-1}(S^{2n-1}) \times \pi_i(S^{4n-1})$ which hold for $n = 1, 2, 4$ and all i. The case $n = 2$ says that each entry in the fourth row is the product

of the entry diagonally above it to the left and the entry three units below it. Actually, these isomorphisms hold for all n if one factors out 2-torsion — elements of order a power of 2.

The next regular feature in the table is the sequence of $\mathbb{Z}$'s down the diagonal. This is an illustration of the Hurewicz theorem, which asserts that for a simply-connected space X, the first nonzero homotopy group $\pi_n(X)$ is isomorphic to the first nonzero homology group $\widetilde{H}_n(X)$. The proof shows moreover that the isomorphism $\pi_n(S^n) \approx \mathbb{Z}$ associates to a map $f : S^n \to S^n$ its degree, so $\pi_n(S^n)$ is generated by the identity map of S^n.

What else is notable about the table? One may observe that all the groups above the diagonal are finite except for $\pi_3(S^2)$, $\pi_7(S^4)$, and $\pi_{11}(S^6)$. In §4.4c we use cup products in cohomology to show that $\pi_{4k-1}(S^{2k})$ contains a $\mathbb{Z}$ direct summand for all $k \geq 1$. It is a theorem of Serre which will be proved in Volume II that $\pi_i(S^n)$ is finite for $i > n$ except for $\pi_{4k-1}(S^{2k})$, which is the direct sum of $\mathbb{Z}$ with a finite group. So all the complexity of the homotopy groups of spheres resides in finite abelian groups. The problem thus reduces to computing the p-torsion in $\pi_i(S^n)$ for each prime p.

An especially interesting feature of the table which might not be quite obvious from the relatively small sample shown is that along each diagonal, the groups $\pi_{n+k}(S^n)$ with k fixed and varying n eventually become independent of n for large enough n. This 'stability' property will be proved in §4.2, where we give more extensive tables of these stable homotopy groups of spheres.

Returning to general theory, let us consider the question of how $\pi_n(X, x_0)$ depends on the choice of the basepoint x_0. Given a path $\gamma : I \to X$ from $x_0 = \gamma(0)$ to another basepoint $x_1 = \gamma(1)$, we may associate to each map $f : (I^n, \partial I^n) \to (X, x_1)$ a new map $\gamma f : (I^n, \partial I^n) \to (X, x_0)$ obtained by shrinking the domain of f to a smaller concentric cube in I^n, and then inserting the path γ on each radial segment in the shell between this smaller cube and ∂I^n.

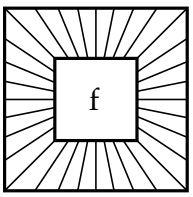

Note that when $n = 1$, γf according to this definition is the composition of three paths: γ, f, and the inverse of γ. So to be consistent, a different notation from 'γf' should really be used when $n = 1$. However, since we are mainly interested in the cases $n > 1$, we shall leave it to the reader to make the necessary notational changes when $n = 1$.

It is clear that a homotopy of γ or f through maps fixing ∂I or ∂I^n, respectively, yields a homotopy of γf through maps $(I^n, \partial I^n) \to (X, x_0)$. Further basic properties are:

(1) $\gamma(f + g) \simeq \gamma f + \gamma g$

(2) $(\gamma\eta)f \simeq \gamma(\eta f)$

(3) $1f \simeq f$ where 1 denotes the constant path.

The homotopies in (2) and (3) are obvious. For (1), we may first deform f and g to be constant on the right and left halves of I^n, respectively, then we excise a progressively wider symmetric middle slab of $\gamma f + \gamma g$ until it becomes $\gamma(f+g)$:

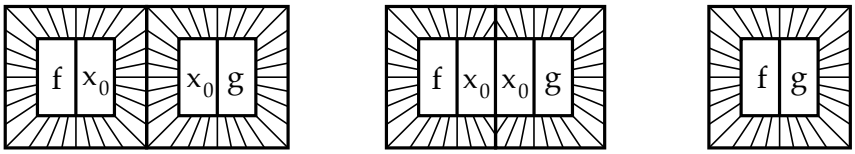

An explicit formula for this homotopy is

$$h_t(s_1, \cdots, s_n) = \gamma(f+0)\big((1-t)s_1, s_2, \cdots, s_n\big) + \gamma(0+g)\big((1-t)s_1 + t, s_2, \cdots, s_n\big)$$

for $0 \le t \le \frac{1}{2}$. The second plus sign makes sense since $\gamma(f+0)\big((1-t), s_2, \cdots, s_n\big) = \gamma(0+g)\big(t, s_2, \cdots, s_n\big)$ for $0 \le t \le \frac{1}{2}$. Since $h_0 = \gamma(f+0)+\gamma(0+g)$ and $h_{1/2} = \gamma(f+g)$, we have $\gamma(f+g) \simeq \gamma(f+0) + \gamma(0+g) \simeq \gamma f + \gamma g$.

If we define a change-of-basepoint transformation $\beta_\gamma : \pi_n(X, x_0) \to \pi_n(X, x_1)$ by $\beta_\gamma([f]) = [\gamma f]$, then (1) shows that β_γ is a homomorphism, while (2) and (3) imply that β_γ is an isomorphism with inverse $\beta_{\overline{\gamma}}$ where $\overline{\gamma}$ is the inverse path of γ, $\overline{\gamma}(s) = \gamma(1-s)$. Thus if X is path-connected, different choices of basepoint x_0 yield isomorphic groups $\pi_n(X, x_0)$, which may then be written simply as $\pi_n(X)$.

Now let us restrict attention to loops γ at the basepoint x_0. Since $\beta_{\gamma\eta} = \beta_\gamma\beta_\eta$, the association $[\gamma] \mapsto \beta_\gamma$ defines a homomorphism from $\pi_1(X, x_0)$ to $Aut(\pi_n(X, x_0))$, the group of automorphisms of $\pi_n(X, x_0)$. This is called the **action of π_1 on π_n**, each element of π_1 acting as an automorphism $[f] \mapsto [\gamma f]$ of π_n. When $n = 1$ this is the action of π_1 on itself by inner automorphisms. When $n > 1$, the action makes the abelian group $\pi_n(X, x_0)$ into a module over the group ring $\mathbb{Z}[\pi_1(X, x_0)]$. Elements of $\mathbb{Z}[\pi_1]$ are finite sums $\sum_i n_i\gamma_i$ with $n_i \in \mathbb{Z}$ and $\gamma_i \in \pi_1$, multiplication being defined by distributivity and the multiplication in π_1. The module structure on π_n is given by $\big(\sum_i n_i\gamma_i\big)\alpha = \sum_i n_i(\gamma_i\alpha)$ for $\alpha \in \pi_n$. For brevity one sometimes says π_n is a π_1-module rather than a $\mathbb{Z}[\pi_1]$-module.

In the literature, a space with trivial π_1 action on π_n is called 'n-simple,' and 'simple' means 'n-simple for all n.' It would be nice to have more descriptive terms for these properties. In this book we shall call a space with trivial action of π_1 on all π_n's **abelian** since the defining property generalizes the condition that π_1 be abelian. This is consistent with a long-established usage of the term 'nilpotent' to refer to spaces with nilpotent π_1 and nilpotent action of π_1 on all higher homotopy groups; see [Hilton-Mislin-Roitberg].

We next observe that π_n is a functor. Namely, a map $\varphi:(X,x_0)\to(Y,y_0)$ induces $\varphi_*:\pi_n(X,x_0)\to\pi_n(Y,y_0)$ defined by $\varphi_*([f])=[\varphi f]$. It is immediate from the definitions that φ_* is well-defined and a homomorphism for $n\geq 1$. The functor properties $(\varphi\psi)_*=\varphi_*\psi_*$ and $\mathbb{1}_*=\mathbb{1}$ are also evident, as is the fact that if $\varphi_t:(X,x_0)\to(Y,y_0)$ is a homotopy then $\varphi_{0*}=\varphi_{1*}$.

In particular, a homotopy equivalence $(X,x_0)\simeq(Y,y_0)$ in the basepointed sense induces isomorphisms on all homotopy groups π_n. This is true even if basepoints are not required to be stationary during homotopies, by an exercise at the end of the section.

Homotopy groups behave very nicely with respect to covering spaces:

Proposition 4.1. *If $p:(\widetilde{X},\widetilde{x}_0)\to(X,x_0)$ is a covering space, then $p_*:\pi_n(\widetilde{X},\widetilde{x}_0)\to\pi_n(X,x_0)$ is an isomorphism for all $n\geq 2$.*

Proof: If $n\geq 2$, the lifting criterion of Proposition 1.19 implies that a map $(S^n,s_0)\to(X,x_0)$ lifts to $(\widetilde{X},\widetilde{x}_0)$ since $\pi_1(S^n)=0$. So p_* is surjective on π_n. Injectivity is immediate from the covering homotopy property, just as in Proposition 1.17 which treated the case $n=1$. □

In particular, $\pi_n(X,x_0)=0$ for $n\geq 2$ whenever X has a contractible universal cover. This applies for example to S^1, so we obtain the first row of the table of homotopy groups of spheres shown earlier. More generally, the n-dimensional torus T^n, the product of n circles, has universal cover $\mathbb{R}^n$, so $\pi_i(T^n)=0$ for $i>0$. This is in marked contrast to the homology groups $H_i(T^n)$ which are nonzero for all $i\leq n$. Spaces with $\pi_n=0$ for $n\geq 2$ are sometimes called **aspherical**.

The behavior of homotopy groups with respect to products is very simple:

Proposition 4.2. *For a product $\prod_\alpha X_\alpha$ of an arbitrary collection of path-connected spaces X_α there are isomorphisms $\pi_n(\prod_\alpha X_\alpha)\approx\prod_\alpha\pi_n(X_\alpha)$ for all n.*

Proof: A map $f:Y\to\prod_\alpha X_\alpha$ is the same thing as a collection of maps $f_\alpha:Y\to X_\alpha$. Taking Y to be S^n and $S^n\times I$ gives the result. □

The Long Exact Sequence

To define relative groups $\pi_n(X,A,x_0)$ for a pair (X,A) with a basepoint $x_0\in A$, regard I^{n-1} as the face of I^n with the last coordinate $s_n=0$ and let J^{n-1} be the closure of ∂I^n-I^{n-1}, the union of the remaining faces of I^n. Then $\pi_n(X,A,x_0)$ for $n\geq 1$ is defined to be the set of homotopy classes of maps $(I^n,\partial I^n,J^{n-1})\to(X,A,x_0)$, with homotopies through maps of the same form. There does not seem to be a completely

satisfactory way of defining $\pi_0(X, A, x_0)$, so we shall leave this undefined. Note that $\pi_n(X, x_0, x_0) = \pi_n(X, x_0)$, so absolute homotopy groups are a special case of relative homotopy groups, just as with homology groups.

A sum operation is defined in $\pi_n(X, A, x_0)$ by the same formulas as for $\pi_n(X, x_0)$, except that the coordinate s_n now plays a special role and is no longer available for the sum operation. Thus $\pi_n(X, A, x_0)$ is a group for $n \geq 2$, and this group is abelian for $n \geq 3$. For $n = 1$ we have $I^1 = [0, 1]$, $I^0 = \{0\}$, and $J^0 = \{1\}$, so $\pi_1(X, A, x_0)$ is the set of homotopy classes of paths in X from a varying point in A to the fixed basepoint $x_0 \in A$. In general this is not a group in any natural way.

Just as $\pi_n(X, x_0)$ can be regarded as homotopy classes of maps $(S^n, s_0) \to (X, x_0)$, so $\pi_n(X, A, x_0)$ can be viewed as homotopy classes of maps $(D^n, S^{n-1}, s_0) \to (X, A, x_0)$, since collapsing J^{n-1} to a point converts $(I^n, \partial I^n, J^{n-1})$ into (D^n, S^{n-1}, s_0). From this viewpoint, addition is done via the map $c\colon D^n \to D^n \vee D^n$ collapsing $D^{n-1} \subset D^n$ to a point.

A map $\varphi\colon (X, A, x_0) \to (Y, B, y_0)$ induces maps $\varphi_*\colon \pi_n(X, A, x_0) \to \pi_n(Y, B, y_0)$ which are homomorphisms for $n \geq 2$ and have properties analogous to those in the absolute case: $(\varphi\psi)_* = \varphi_*\psi_*$, $\mathbb{1}_* = \mathbb{1}$, and $\varphi_* = \psi_*$ if $\varphi \simeq \psi$.

Probably the most useful feature of the relative groups $\pi_n(X, A, x_0)$ is that they fit into a long exact sequence

$$\cdots \longrightarrow \pi_n(A, x_0) \xrightarrow{i} \pi_n(X, x_0) \xrightarrow{j} \pi_n(X, A, x_0) \xrightarrow{\partial} \pi_{n-1}(A, x_0) \longrightarrow \cdots \longrightarrow \pi_0(X, x_0)$$

Here i is induced by inclusion, j comes from regarding a map $(I^n, \partial I^n) \to (X, x_0)$ as a map $(I^n, \partial I^n, J^{n-1}) \to (X, A, x_0)$, and ∂ from restricting maps $(I^n, \partial I^n, J^{n-1}) \to (X, A, x_0)$ to I^{n-1}. The maps i, j, and ∂ are clearly homomorphisms when group structures are defined in domain and range.

Theorem 4.3. *This sequence is exact.*

Near the end of the sequence, where group structures are not defined, exactness still makes sense: The image of one map is the kernel of the next, i.e., those elements mapping to the homotopy class of the constant map.

Our proof of this theorem will use a general observation: A map $f\colon (I^n, \partial I^n, J^{n-1}) \to (X, A, x_0)$ represents zero in $\pi_n(X, A, x_0)$ iff it is homotopic rel ∂I^n to a map with image contained in A. For if we have such a homotopy to a map g, then $[f] = [g]$ in $\pi_n(X, A, x_0)$, and $[g] = 0$ via the homotopy obtained by composing g with a deformation retraction of I^n onto J^{n-1}. Conversely, if $[f] = 0$ via a homotopy $F\colon I^{n+1} \to X$, then by restricting F to a family of n-disks in I^{n+1} starting with I^n and ending with J^n, all these disks

having the same boundary, then we get a homotopy from f to a map into A, stationary on ∂I^n.

Proof: With just about the same effort we can derive the long exact sequence of a triple (X, A, B, x_0) with $x_0 \in B \subset A \subset X$:

$$\cdots \longrightarrow \pi_n(A, B, x_0) \xrightarrow{i} \pi_n(X, B, x_0) \xrightarrow{j} \pi_n(X, A, x_0) \xrightarrow{\partial} \pi_{n-1}(A, B, x_0) \longrightarrow \cdots \longrightarrow \pi_1(X, A, x_0)$$

When $B = x_0$ this reduces to the exact sequence for the pair (X, A, x_0), though the latter sequence continues on two more steps to $\pi_0(X, x_0)$. Exactness at these last two steps we leave as a simple exercise.

Exactness at $\pi_n(X, B, x_0)$: The composition ji is zero since a map $(I^n, \partial I^n, J^{n-1}) \to (A, B, x_0)$ represents zero in $\pi_n(X, A, x_0)$ by the remarks preceding the proof. Conversely, if $f : (I^n, \partial I^n, J^{n-1}) \to (X, B, x_0)$ represents zero in $\pi_n(X, A, x_0)$, then it is homotopic rel ∂I^n to a map with image in A, hence the class $[f] \in \pi_n(X, B, x_0)$ is in the image of i.

Exactness at $\pi_n(X, A, x_0)$: The composition ∂j is zero since the restriction of a map $(I^n, \partial I^n, J^{n-1}) \to (X, B, x_0)$ to I^{n-1} has image lying in B, and hence represents zero in $\pi_{n-1}(A, B, x_0)$. Conversely, suppose the restriction of $f : (I^n, \partial I^n, J^{n-1}) \to (X, A, x_0)$ to I^{n-1} represents zero in $\pi_{n-1}(A, B, x_0)$, so there is a homotopy F from $f\,|\,I^{n-1}$ to a map with image in B, rel ∂I^{n-1}. We can tack F onto f to get a new map $f_1 : (I^n, \partial I^n, J^{n-1}) \to (X, B, x_0)$ which, as a map $(I^n, \partial I^n, J^{n-1}) \to (X, A, x_0)$, is homotopic to f by the homotopy which tacks on increasingly longer initial segments of F. So $[f] = [f_1] \in Im(j)$.

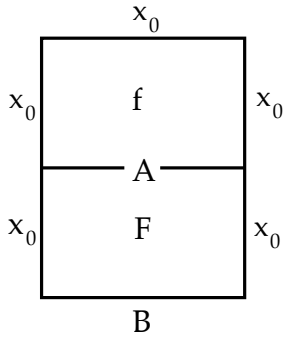

Exactness at $\pi_n(A, B, x_0)$: The composition $i\partial$ is zero since the restriction of a map $f : (I^{n+1}, \partial I^{n+1}, J^n) \to (X, B, x_0)$ to I^n is homotopic rel ∂I^n to a constant map via f itself. The converse is trivial if B is a point, since a nullhomotopy $f_t : (I^n, \partial I^n) \to (X, x_0)$ of $f_0 : (I^n, \partial I^n) \to (A, x_0)$ gives a map $F : (I^{n+1}, \partial I^{n+1}, J^n) \to (X, A, x_0)$ with $\partial([F]) = [f_0]$. Thus the proof is finished in this case. For a general B, let F be a nullhomotopy of $f : (I^n, \partial I^n, J^{n-1}) \to (A, B, x_0)$ through maps $(I^n, \partial I^n, J^{n-1}) \to (X, B, x_0)$, and let g be the restriction of F to $I^{n-1} \times I$, as in the first of the two pictures. Reparametrizing the n^{th} and $(n+1)^{st}$ coordinates as shown in the second picture, we see that f with g tacked on is in the image of ∂. But as we noted in the preceding paragraph, tacking g onto f gives the same element of $\pi_n(A, B, x_0)$. □

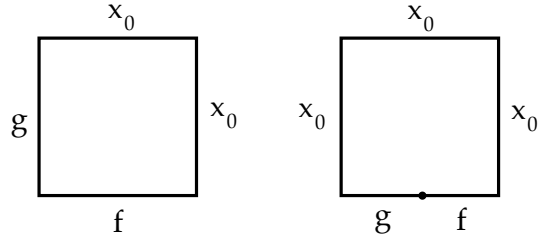

Example. Let CX be the cone on a path-connected space X, so CX is the quotient space of $X \times I$ obtained by collapsing $X \times \{0\}$ to a point. We can view X as the subspace $X \times \{1\} \subset CX$. Since CX is contractible, the long exact sequence of homotopy groups gives isomorphisms $\pi_n(CX, X, x_0) \approx \pi_{n-1}(X, x_0)$. Taking $n = 2$, we can thus realize any group G, abelian or not, as a relative π_2 by choosing X to have $\pi_1(X) \approx G$.

The long exact sequence of homotopy groups is clearly natural: A map $(X, A, x_0) \to (Y, B, y_0)$ induces a map between the long exact sequences for these two pairs of spaces, with commuting squares.

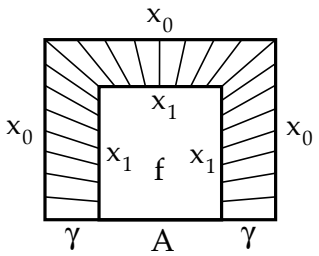

For relative homotopy groups there are basepoint-change isomorphisms β_γ analogous to those in the abolute case. One starts with a path γ in $A \subset X$ from x_0 to x_1, and this induces $\beta_\gamma : \pi_n(X, A, x_1) \to \pi_n(X, A, x_0)$, $\beta_\gamma[f] = [\gamma f]$, where γf is defined as in the picture. The basic properties of these relative β_γ's parallel closely those developed earlier in the absolute case. However, even though absolute homotopy groups $\pi_n(X, x_0)$ are the special case of relative groups $\pi_n(X, A, x_0)$ in which $A = x_0$, the relative β_γ's do not specialize to the absolute β_γ's. So separate proofs must be given in the two cases. But the proofs are quite similar, so we leave this to the exercises.

Restricting γ to loops at the basepoint, the association $\gamma \mapsto \beta_\gamma$ defines an action of $\pi_1(A, x_0)$ on $\pi_n(X, A, x_0)$, as in the absolute case. In fact, it is clear from the definitions that $\pi_1(A, x_0)$ acts on the whole long exact sequence of homotopy groups for (X, A, x_0), the action commuting with the various maps in the sequence.

A space X with basepoint x_0 is called **n-connected** if $\pi_i(X, x_0) = 0$ for $i \le n$. Thus 0-connected means path-connected and 1-connected means simply-connected. Since n-connected implies 0-connected, the choice of the basepoint x_0 is not significant. The condition of being n-connected can be expressed without mention of a basepoint since it is an easy exercise to check that the following three conditions are equivalent.

(1) Every map $S^i \to X$ is homotopic to a constant map.
(2) Every map $S^i \to X$ extends to a map $D^{i+1} \to X$.
(3) $\pi_i(X, x_0) = 0$ for all $x_0 \in X$.

Thus X is n-connected if any one of these three conditions holds for all $i \le n$. Similarly, in the relative case it is not hard to see that the following four conditions are equivalent, for $i > 0$:

(1) Every map $(D^i, \partial D^i) \to (X, A)$ is homotopic, staying fixed on ∂D^i, to a map $D^i \to A$.
(2) Every map $(D^i, \partial D^i) \to (X, A)$ is homotopic, through maps $(D^i, \partial D^i) \to (X, A)$, to a map $D^i \to A$.
(3) Every map $(D^i, \partial D^i) \to (X, A)$ is homotopic, through maps $(D^i, \partial D^i) \to (X, A)$, to a constant map $D^i \to A$.
(4) $\pi_i(X, A, x_0) = 0$ for all $x_0 \in A$.

When $i = 0$ we did not define the relative π_0, and (1)-(3) are trivially equivalent since D^0 is a point and ∂D^0 is empty. The pair (X, A) is called **n-connected** if (1)-(4) hold for all $0 < i \le n$ and (1)-(3) hold for $i = 0$. Note that X is n-connected iff (X, x_0) is n-connected for some x_0, hence for all x_0. It is sometimes convenient to say a space or pair is ∞-connected if it is n-connected for all n.

Whitehead's Theorem

Since CW complexes are built using attaching maps whose domains are spheres, it is perhaps not too surprising that homotopy groups of CW complexes carry a lot of information. Whitehead's theorem makes this explicit:

Theorem 4.4. *If a map $f : X \to Y$ between connected CW complexes induces isomorphisms $f_* : \pi_n(X) \to \pi_n(Y)$ for all n, then f is a homotopy equivalence. In case f is the inclusion of a subcomplex $X \hookrightarrow Y$, the conclusion is stronger: X is a deformation retract of Y.*

The proof will follow rather easily from a more technical result which is useful in a number of other contexts as well:

Lemma 4.5. *Let (X, A) be a CW pair and (Y, B) a pair with Y path-connected and B nonempty. Suppose that for each $n > 0$ such that $X - A$ has cells of dimension n, $\pi_n(Y, B, y_0) = 0$ for all $y_0 \in B$. Then every map $f : (X, A) \to (Y, B)$ is homotopic* rel A *to a map $X \to B$.*

Proof: Assume inductively that f has already been homotoped to take the skeleton X^{k-1} to B. If Φ is the characteristic map of a cell e^k of $X - A$, the composition $f\Phi : (D^k, \partial D^k) \to (Y, B)$ can be homotoped into B rel ∂D^k since $\pi_k(Y, B) = 0$ by hypothesis. This homotopy induces a homotopy of f on the quotient space $X^{k-1} \cup e^k$ of $X^{k-1} \amalg D^k$, a homotopy rel X^{k-1}. Doing this for all k-cells of $X - A$ simultaneously, and taking the constant homotopy on A, we obtain a homotopy of $f \,|\, X^k \cup A$ to a map into B. By the homotopy extension property, this homotopy extends to a homotopy defined on

all of X, completing the induction step. The induction starts with the case $k = 0$, where we use the assumption that Y is path-connected and B is nonempty to deform f to take 0-cells into B.

If the cells of $X - A$ are of bounded dimension, finitely many applications of the induction step finish the proof. To handle the case that $X - A$ has cells of arbitrarily large dimension, we perform the homotopy of the induction step during the t-interval $[1 - 1/2^k, 1 - 1/2^{k+1}]$. Any finite skeleton X^k is eventually stationary under these homotopies, hence we have a well-defined homotopy f_t, $t \in [0, 1]$, with $f_1(X) \subset B$. □

Proof of Whitehead's Theorem: In the special case that f is the inclusion of a subcomplex, consider the long exact sequence of homotopy groups for the pair (Y, X). Since f induces isomorphisms on all homotopy groups, the relative groups $\pi_n(Y, X)$ are zero, so (Y, X) is ∞-connected. Applying the lemma to the identity map $(Y, X) \to (Y, X)$ then yields a deformation retraction of Y onto X.

The general case can be proved using mapping cylinders. Recall that the mapping cylinder M_f of a map $f : X \to Y$ is the quotient space of the disjoint union of $X \times I$ and Y under the identification $(x, 1) \sim f(x)$. Thus M_f contains both $X = X \times \{0\}$ and Y as subspaces, and M_f deformation retracts onto Y. The map f becomes the composition of the inclusion $X \hookrightarrow M_f$ with the retraction $M_f \to Y$. Since this retraction is a homotopy equivalence, it suffices to show that M_f deformation retracts onto X if f induces isomorphisms on homotopy groups, i.e., if the relative groups $\pi_n(M_f, X)$ are all zero.

To construct such a deformation retraction, first apply the preceding lemma to homotope the inclusion $(X \cup Y, X) \hookrightarrow (M_f, X)$ to a map into X, staying fixed on X. Since the pair $(M_f, X \cup Y)$ obviously satisfies the homotopy extension property, this homotopy extends to a homotopy from the identity map of M_f to a map $g : M_f \to M_f$ taking $X \cup Y$ into X. Then apply the lemma again to the composition $(X \times I \amalg Y, X \times \partial I \amalg Y) \longrightarrow (M_f, X) \xrightarrow{g} (M_f, X)$ to finish the construction of a deformation retraction of M_f onto X. □

Whitehead's theorem does not say that two CW complexes X and Y with isomorphic homotopy groups are homotopy equivalent, since there is a big difference between saying that X and Y have isomorphic homotopy groups and that there is a map $X \to Y$ inducing isomorphisms on homotopy groups. For example, consider $X = \mathbb{R}\mathrm{P}^2$ and $Y = S^2 \times \mathbb{R}\mathrm{P}^\infty$. These both have fundamental group $\mathbb{Z}_2$, and Proposition 4.1 implies that their higher homotopy groups are isomorphic since their universal covers S^2 and $S^2 \times S^\infty$ are homotopy

equivalent, S^∞ being contractible. But $\mathbb{R}\mathrm{P}^2$ and $S^2 \times \mathbb{R}\mathrm{P}^\infty$ are not homotopy equivalent since their homology groups are vastly different, $S^2 \times \mathbb{R}\mathrm{P}^\infty$ having nonvanishing homology in infinitely many dimensions since it retracts onto $\mathbb{R}\mathrm{P}^\infty$. Another pair of non-homotopy-equivalent CW complexes with isomorphic homotopy groups is S^2 and $S^3 \times \mathbb{C}\mathrm{P}^\infty$, as we shall see in §4.3.

The Hurewicz Theorem

To do our first real calculations of higher homotopy groups we will relate homotopy groups to homology groups via the Hurewicz theorem. For example, we will compute $\pi_n(S^n)$ this way, and show that $\pi_i(S^n) = 0$ for $i < n$. A different way of obtaining these calculations without using homology will be described in the next section.

There is a natural map from homotopy groups to homology groups. Thinking of $\pi_n(X, A, x_0)$ for $n > 0$ as homotopy classes of maps $f : (D^n, \partial D^n, s_0) \to (X, A, x_0)$, the **Hurewicz map** $h : \pi_n(X, A, x_0) \to H_n(X, A)$ is defined by $h([f]) = f_*(\alpha)$ where α is a fixed generator of $H_n(D^n, \partial D^n) \approx \mathbb{Z}$ and $f_* : H_n(D^n, \partial D^n) \to H_n(X, A)$ is induced by f. If $f \simeq g$ through maps $(D^n, \partial D^n, s_0) \to (X, A, x_0)$, or even through maps $(D^n, \partial D^n) \to (X, A)$, then $f_* = g_*$ so h is well-defined.

One can be more specific about the definition of h by identifying D^n with the standard simplex Δ^n. From the isomorphism between simplicial and singular homology it follows that the identity map $\Delta^n \to \Delta^n$, viewed as a singular n-simplex in Δ^n, represents a generator $\alpha \in H_n(\Delta^n, \partial\Delta^n)$. Hence if $f : (\Delta^n, \partial\Delta^n, v_0) \to (X, A, x_0)$ represents an element of $\pi_n(X, A, x_0)$, then $f_*(\alpha)$ is the relative homology class represented by $f : \Delta^n \to X$, regarded as a singular n-simplex which is a relative cycle since $f(\partial\Delta^n) \subset A$. Thus h sends the homotopy class represented by f to the homology class represented by f.

Proposition 4.6. *The Hurewicz map $h : \pi_n(X, A, x_0) \to H_n(X, A)$ is a homomorphism, assuming $n > 1$ so that $\pi_n(X, A, x_0)$ is a group.*

Proof: It suffices to show that for maps $f, g : (D^n, \partial D^n) \to (X, A)$, the induced maps on homology satisfy $(f + g)_* = f_* + g_*$, for then $h([f + g]) = (f + g)_*(\alpha) = f_*(\alpha) + g_*(\alpha) = h([f]) + h([g])$. Our proof that $(f + g)_* = f_* + g_*$ will in fact work for any homology theory, and the obvious modifications of the argument work for any cohomology theory.

Let $c : D^n \to D^n \vee D^n$ be the map collapsing the equatorial D^{n-1} to a point, and let $q_1, q_2 : D^n \vee D^n \to D^n$ be the two quotient maps collapsing the second and first wedge summands to a point, respectively. We then have a diagram

$$\begin{array}{ccccc} H_i(D^n,\partial D^n) & \xrightarrow{c_*} & H_i(D^n\vee D^n,\partial D^n\vee\partial D^n) & \xrightarrow{(f\vee g)_*} & H_i(X,A) \\ & & \approx\big\downarrow{\scriptstyle q_{1*}\oplus q_{2*}} & & \\ & & H_i(D^n,\partial D^n)\bigoplus H_i(D^n,\partial D^n) & & \end{array}$$

The map $q_{1*}\oplus q_{2*}$ is an isomorphism with inverse $i_{1*}+i_{2*}$ where i_1 and i_2 are the inclusions of the two summands of $D^n\vee D^n$. Since q_1c and q_2c are homotopic to the identity as maps $(D^n,\partial D^n)\to(D^n,\partial D^n)$, it follows that the composition $(q_{1*}\oplus q_{2*})c_*$ is the diagonal map $x\mapsto(x,x)$. The fact that $(f\vee g)i_1=f$ and $(f\vee g)i_2=g$ implies that $(f\vee g)_*(i_{1*}+i_{2*})$ sends $(x,0)$ to $f_*(x)$ and $(0,x)$ to $g_*(x)$, hence (x,x) to $f_*(x)+g_*(x)$. Thus the composition across the top of the diagram is $x\mapsto f_*(x)+g_*(x)$. On the other hand, $f+g=(f\vee g)c$, so this composition is also $(f+g)_*$. □

There is also an absolute Hurewicz map $h:\pi_n(X,x_0)\to H_n(X)$, defined by $h([f])=f_*(\alpha)$ for $f:(S^n,s_0)\to(X,x_0)$ and α a generator of $H_n(S^n)$. For example, if $X=S^n$ then $f_*(\alpha)$ is $\deg(f)\alpha$ by the definition of degree, so we can view h in this case as the degree map $\pi_n(S^n)\to\mathbb{Z}$.

An easy definition-check which we leave to the reader shows that the absolute Hurewicz map fits with the relative one to form a commutative diagram

$$\begin{array}{ccccccccc} \cdots\longrightarrow & \pi_n(A,x_0) & \longrightarrow & \pi_n(X,x_0) & \longrightarrow & \pi_n(X,A,x_0) & \longrightarrow & \pi_{n-1}(A,x_0) & \longrightarrow\cdots \\ & \downarrow h & & \downarrow h & & \downarrow h & & \downarrow h & \\ \cdots\longrightarrow & H_n(A) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X,A) & \longrightarrow & H_{n-1}(A) & \longrightarrow\cdots \end{array}$$

Taking $A=\{x_0\}$, the absolute h becomes a special case of the relative h. In particular, the absolute h is also a homomorphism.

Another elementary fact is that Hurewicz maps are natural, i.e., a map $f:(X,x_0)\to(Y,y_0)$ induces a commutative diagram

$$\begin{array}{ccc} \pi_n(X,x_0) & \xrightarrow{f_*} & \pi_n(Y,y_0) \\ \downarrow h & & \downarrow h \\ H_n(X) & \xrightarrow{f_*} & H_n(Y) \end{array}$$

and similarly in the relative case.

The kernel of the Hurewicz map $h:\pi_1(X,x_0)\to H_1(X)$ contains the commutator subgroup of $\pi_1(X)$ since $H_1(X)$ is abelian. Analogously, the kernel of $h:\pi_n(X,A,x_0)\to H_n(X,A)$ contains the elements of the form $[\gamma f-f]$ for γ a loop in A, since $\gamma f\simeq f$

as non-basepoint-preserving maps $(D^n, \partial D^n) \to (X, A)$, hence $(\gamma f)_*(\alpha) = f_*(\alpha)$. Let $\pi'_n(X, A, x_0)$ be the quotient group of $\pi_n(X, A, x_0)$ obtained by factoring out the subgroup generated by all elements of the form $[\gamma f - f]$, or the normal subgroup generated by such elements when $n = 2$ since $\pi_2(X, A, x_0)$ may not be abelian. Then h induces a quotient homomorphism $h' : \pi'_n(X, A, x_0) \to H_n(X, A)$.

We can now state the Hurewicz theorem in its general form:

Theorem 4.7. *For a pair (X, A) of path-connected spaces, suppose there is an integer $n \geq 2$ such that $\pi_i(X, A, x_0) = 0$ for $i < n$, i.e., the pair (X, A) is $(n-1)$-connected. Then $H_i(X, A) = 0$ for $i < n$ and $h' : \pi'_n(X, A, x_0) \to H_n(X, A)$ is an isomorphism.*

In particular, taking A to be a point, this says the absolute Hurewicz map $h : \pi_n(X) \to H_n(X)$ is an isomorphism if X is $(n-1)$-connected and $n \geq 2$. Thus if X is simply-connected we have $\pi_2(X) \approx H_2(X)$, and if the latter group is zero then $\pi_3(X) \approx H_3(X)$, etc. Hence for a simply-connected space, either all its homology groups are trivial, in which case all its homotopy groups also vanish, or its first nonzero homology group is isomorphic to its first nonzero homotopy group.

Corollary 4.8. *For all $n \geq 1$ we have $\pi_i(S^n) = 0$ if $i < n$ and $\pi_n(S^n) \approx \mathbb{Z}$, generated by the identity map, hence two maps $S^n \to S^n$ are homotopic iff they have the same degree.*

Proof: The preceding remarks give the calculation of $\pi_i(S^n)$ for $i \leq n$. For the assertion about a generator, recall that $h([f]) = f_*(\alpha)$ for $f : (S^n, s_0) \to (X, x_0)$ and α a generator of $H_n(S^n)$. Hence $h([\mathbb{1}]) = \alpha$, so $\pi_n(S^n)$ must be generated by the identity map. The last statement of the corollary follows since every map $S^n \to S^n$ can be deformed to be basepoint-preserving if $n \geq 1$. □

We will give two different proofs of the Hurewicz theorem in this chapter. The one we give now is a fairly direct proof by induction, using a small amount of homology theory. The other proof, in the next section, uses a little more homotopy theory to calculate $\pi_n(S^n)$ first and then deduce the Hurewicz theorem from this. Readers with an inclination toward homotopy theory might choose to skip the first proof.

The bulk of the proof we give now will be contained in three lemmas. The first is a general fact about the relative π_2:

Lemma 4.9. *For any (X, A, x_0), the formula $a + b - a = (\partial a)b$ holds for all $a, b \in \pi_2(X, A, x_0)$, where $\partial : \pi_2(X, A, x_0) \to \pi_1(A, x_0)$ is the usual boundary map and $(\partial a)b$ denotes the action of ∂a on b. Hence $\pi'_2(X, A, x_0)$ is abelian.*

Here the '+' and '−' in $a+b-a$ refer to the group operation in the nonabelian group $\pi_2(X, A, x_0)$.

Proof: The formula is obtained by constructing a homotopy from $a+b-a$ to $(\partial a)b$ as in the following pictures: □

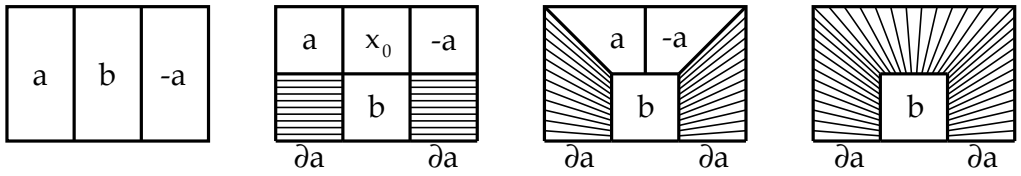

To state the second lemma we need some preliminary definitions. An (n, q)-*chain* in a pair (X, A) with basepoint $x_0 \in A$ is a linear combination $\sum_i n_i\sigma_i$, $n_i \in \mathbb{Z}$, of singular n-simplices $\sigma_i : \Delta^n \to X$ which take the vertices of Δ^n to x_0 and the q-skeleton of Δ^n to A. Let $C_{n,q}(X, A, x_0)$ denote the quotient of the group of (n, q)-chains in (X, A) by the subgroup of (n, q)-chains which lie in A. The usual boundary operator ∂ makes the groups $C_{n,q}(X, A, x_0)$ into a chain complex, with homology groups denoted $H_{n,q}(X, A, x_0)$, and the natural chain maps $C_{n,q}(X, A, x_0) \to C_n(X, A)$ induce homomorphisms $H_{n,q}(X, A, x_0) \to H_n(X, A)$.

Lemma 4.10. *If the pair (X, A) is q-connected and A is path-connected, then the maps $H_{n,q}(X, A, x_0) \to H_n(X, A)$ are isomorphisms for all n.*

Since $C_{n,q}(X, A, x_0) = 0$ if $n \le q$, this says in particular that $H_n(X, A) = 0$ for $n \le q$, which is part of the assertion of the Hurewicz theorem.

Proof: First we show surjectivity. Represent an element of $H_n(X, A)$ by a chain $\xi = \sum_i \varepsilon_i f_i$ with $f_i : \Delta^n \to X$. By repeating terms f_i with $|\varepsilon_i| > 1$ we may arrange that $\varepsilon_i = \pm 1$ for all i. Recall from §2.1 the construction of an associated Δ-complex $K = K_\xi$. This is obtained from a collection of disjoint n-simplices Δ^n_i, one for each f_i, by identifying certain $(n-1)$-dimensional faces in pairs, according to a pairing of terms in $\partial\xi$ which differ only in sign. Each such identification is by a linear homeomorphism preserving the given ordering of vertices. By the construction of K, the f_i's induce a map $f : K \to X$. If L is the subcomplex of K consisting of the faces of the Δ^n_i's which are not identified in pairs, then $f(L) \subset A$ since ξ is a relative cycle. Letting $q_i : \Delta^n \to K$ be the natural projection onto the image of Δ^n_i in K, the chain $\sum_i \varepsilon_i q_i$ represents a class $[K] \in H_n(K, L)$ such that $f_*([K])$ is the given element of $H_n(X, A)$.

We now construct a sequence of three homotopies of f through maps $(K, L) \to (X, A)$.

(1) Since A is path-connected, we can deform $f\big|L^0$ to send these vertices to x_0. By the homotopy extension property, we can extend this homotopy first to a homotopy defined on L, with image in A, then to a homotopy defined on all of K.

(2) Since (X, A) is 0-connected, the path-connectedness of A implies that X is path-connected, so for the new f produced by (1) there is a homotopy of $f\big|K^0 \cup L$, staying fixed on L, to a map sending all vertices of $K - L$ to x_0. Extend this homotopy to a homotopy defined on K.

(3) For the f produced in (2), Proposition 4.5 implies that there is a homotopy of $f\big|K^q \cup L$, staying fixed on $K^0 \cup L$, to a map with image in A, since (X, A) is q-connected. Finally, extend this homotopy to be defined on all of K.

For the new f, $f_*([K])$ still represents the same element of $H_n(X, A)$ and is in the image of $H_{n,q}(X, A, x_0) \to H_n(X, A)$ since the new $\xi = \sum_i \varepsilon_i f_i$, $f_i = fq_i$, is an (n, q)-chain.

Injectivity is proved by a relative form of the same argument. An element of the kernel of $H_{n,q}(X, A, x_0) \to H_n(X, A)$ is represented by an (n, q)-chain ζ which is a relative boundary mod A, i.e., $\zeta = \partial\xi + \eta$ for $\xi \in C_{n+1}(X)$ and $\eta \in C_n(A)$. As before, we construct from ξ a Δ-complex $K = K_\xi$, now of dimension $n+1$, with n-dimensional subcomplexes L_ζ and L_η corresponding to the terms ζ and $-\eta$ in $\partial\xi = \zeta - \eta$, and we have a map $f : K \to X$ with $f(L_\eta) \subset A$. As in steps (1)-(3) above, we homotope $f : (K, L_\eta) \to (X, A)$ fixing L_ζ so that vertices in L_η map to x_0, then so that all vertices map to x_0, then so that the q-skeleton maps to A. The result is that $\zeta - \eta$, with the original ζ but a new (n, q)-chain $\eta \in C_n(A)$, bounds the new $(n+1, q)$-chain ξ. Hence ζ is zero in $H_{n,q}(X, A, x_0)$. □

From this lemma it is easy to deduce the surjectivity of $h' : \pi'_n(X, A, x_0) \to H_n(X, A)$ under the hypotheses of the Hurewicz theorem. Namely, the lemma implies that the map $H_{n,n-1}(X, A, x_0) \to H_n(X, A)$ is an isomorphism, so every element of $H_n(X, A)$ is represented by an $(n, n-1)$-chain $\sum_i n_i\sigma_i$ with $\sigma_i : \Delta^n \to X$ taking $\partial\Delta^n$ to A and vertices of Δ^n to x_0. These σ_i's define elements $[\sigma_i] \in \pi_n(X, A, x_0)$, and $h\big(\sum_i n_i[\sigma_i]\big)$ is the homology class of $\sum_i n_i\sigma_i$, making h' surjective.

Injectivity of h' will be shown by embedding the isomorphism $H_{n,n-1}(X, A, x_0) \approx H_n(X, A)$ in a commutative diagram

$$\begin{array}{ccc} H_{n,n-1}(X, A, x_0) & \xrightarrow{\ \approx\ } & H_n(X, A) \\ & {\scriptstyle g'}\searrow \quad \nearrow{\scriptstyle h'} & \\ & \pi'_n(X, A, x_0) & \end{array}$$

Once we have constructed g', it will be obvious that it is surjective, so the diagram then gives injectivity of h'. Actually, the diagram automatically gives surjectivity of h' too, so the preceding paragraph is not logically necessary to the proof.

To begin the construction of g' we define a homomorphism $g: C_{n,n-1}(X,A,x_0) \to \pi_n(X,A,x_0)$ by setting $g\big(\sum_i n_i\sigma_i\big) = \sum_i n_i[\sigma_i]$ for $\sigma_i : (\Delta^n, \partial\Delta^n, v_0) \to (X,A,x_0)$, where $\Delta^n = [v_0,\cdots,v_n]$ and we identify $(\Delta^n,\partial\Delta^n,v_0)$ with $(D^n,\partial D^n,s_0)$. Since $[\sigma_i]=0$ if σ_i has image in A, g is well-defined on $C_{n,n-1}(X,A,x_0)$, at least if $n>2$. The case $n=2$ is special since $\pi_2(X,A,x_0)$ can be nonabelian, which might prevent g from being well-defined. But this difficulty disappears if we pass to the quotient $\pi_2'(X,A,x_0)$ since $\pi_2'(X,A,x_0)$ is always abelian, by the first lemma.

That g induces a homomorphism $g': H_{n,n-1}(X,A) \to \pi_n'(X,A,x_0)$ is a consequence of the next lemma.

Lemma 4.11. *The following composition is zero:*

$$C_{n+1,n-1}(X,A,x_0) \xrightarrow{\partial} C_{n,n-1}(X,A,x_0) \xrightarrow{g} \pi_n(X,A,x_0) \longrightarrow \pi_n'(X,A,x_0)$$

Proof: For $\sigma \in C_{n+1,n-1}(X,A,x_0)$ we have $\partial\sigma = \sum_i (-1)^i\sigma_i$ where σ_i is the restriction of σ to the face $\Delta_i = [v_0,\cdots,\widehat{v}_i,\cdots,v_{n+1}]$. The linear homeomorphism $\Delta^n \to \Delta_i$ preserving the ordering of vertices defines an element $\alpha_i \in \pi_n(\partial\Delta^{n+1}, B, v_0)$ where B is the union of the boundaries of the Δ_i's. When $i=0$ it is understood that the definition of α_0 involves a change of basepoint from v_1 to v_0 via the basepoint-change homomorphism associated to the edge $[v_0,v_1]$. Thus for $i>0$, $[\sigma_i] = \sigma_*(\alpha_i)$ in $\pi_n(X,A,x_0)$, while $\sigma_*(\alpha_0) = [\gamma\sigma_0]$ for γ the loop $\sigma\big|[v_0,v_1]$ in A.

The element $g(\partial\sigma) = \sum_i (-1)^i[\sigma_i]$ of $\pi_n(X,A,x_0)$ has the same image in $\pi_n'(X,A,x_0)$ as $[\gamma\sigma_0] + \sum_{i>0}(-1)^i[\sigma_i] = \sigma_*\big(\sum_i(-1)^i\alpha_i\big)$, so it suffices to show that $\sigma_*\big(\sum_i(-1)^i\alpha_i\big)=0$. We have a commutative diagram

$$\begin{array}{ccccc} \pi_n(\partial\Delta^{n+1},B,v_0) & \longrightarrow & \pi_n(\Delta^{n+1},B,v_0) & \xrightarrow[\approx]{\partial} & \pi_{n-1}(B,v_0) \\ & {\scriptstyle\sigma_*}\searrow \quad \swarrow{\scriptstyle\sigma_*} & & & \\ & \pi_n(X,A,x_0) & & & \end{array}$$

The boundary map ∂ comes from the long exact sequence of the pair (Δ^{n+1},B) and is an isomorphism since Δ^{n+1} is contractible. From this diagram we see that $\sigma_*\big(\sum_i(-1)^i\alpha_i\big) = 0$ if $\partial\big(\sum_i(-1)^i\alpha_i\big)=0$ in $\pi_{n-1}(B,v_0)$.

To show that $\partial\big(\sum_i(-1)^i\alpha_i\big)=0$ we begin with the nonabelian case $n=2$. Since we are only interested in the image of $g(\partial\sigma)$ in the abelian quotient $\pi_2'(X,A,x_0)$, it suffices to show that $\partial(\alpha_0+\alpha_2-\alpha_1-\alpha_3)=0$. Looking at a picture of Δ^3, we see that the loop

$\partial(\alpha_0+\alpha_2-\alpha_1-\alpha_3) = \partial\alpha_0+\partial\alpha_2-\partial\alpha_1-\partial\alpha_3$ passes successively through the vertices v_i for $i = 0,1,2,3,1,0,1,3,0,3,2,0,2,1,0$. This composite loop is obviously trivial in $\pi_1(B, v_0)$.

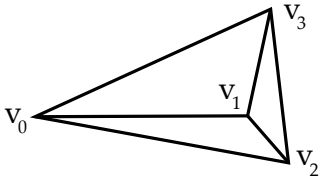

When $n > 2$ we consider the commutative diagram

$$\begin{array}{ccc} \pi_n(\Delta^{n+1}, B, v_0) & \xrightarrow[\approx]{\partial} & \pi_{n-1}(B, v_0) \\ \downarrow h & & \downarrow h \\ H_n(\Delta^{n+1}, B) & \xrightarrow[\approx]{\partial} & H_{n-1}(B) \end{array}$$

Since Δ^{n+1} is contractible, the two boundary maps are isomorphisms. Since B is the $(n-1)$-skeleton of the contractible complex Δ^{n+1} and $n > 2$, B is simply-connected and its homology groups are zero below dimension $n-1$. Assuming inductively that the Hurewicz theorem has already been proved in dimensions less than n, the right-hand h in the diagram is then an isomorphism. Hence the left-hand h is also an isomorphism by commutativity of the diagram. Thus to prove $\partial\big(\sum_i (-1)^i\alpha_i\big) = 0$ in $\pi_{n-1}(B, v_0)$, it suffices to prove that $h\big(\sum_i (-1)^i\alpha_i\big)$ is zero in $H_n(\Delta^{n+1}, B)$. But $h\big(\sum_i (-1)^i\alpha_i\big)$ is represented by the simplicial chain $\sum_i (-1)^i\Delta_i$, which is the simplicial boundary of the simplex Δ^{n+1} itself, so $\sum_i (-1)^i\Delta_i$ is zero in $H_n(\Delta^{n+1}, B)$. □

Proof of the Hurewicz Theorem: Note first that the composition $h'g'$ is the natural map $H_{n,n-1}(X, A, x_0) \to H_n(X, A)$ since g' takes the homology class of a singular $(n, n-1)$-simplex $\sigma : \Delta^n \to X$ to the element of $\pi'_n(X, A, x_0)$ represented by σ, and then h' takes this homotopy class to the homology class it represents in $H_n(X, A)$.

As we observed earlier, the proof will be complete once we show g' is surjective. We can represent each element of $\pi'_n(X, A, x_0)$ by a map $\sigma : (\Delta^n, \partial\Delta^n) \to (X, A)$ taking all but one $(n-1)$-dimensional face of Δ^n to x_0. Such a σ is a singular $(n, n-1)$-simplex, and defines an element of $H_{n,n-1}(X, A, x_0)$ since $C_{n-1,n-1}(X, A, x_0) = 0$. Hence g' is surjective. □

Example. Let $X = S^1 \vee S^n$ for $n \geq 2$. The universal cover $\widetilde{X}$ consists of a copy of $\mathbb{R}$ with an n-sphere S^n_k attached at each integer point k, so $\widetilde{X}$ is homotopy equivalent to $\bigvee_k S^n_k$. Thus $\pi_i(X)$ is zero for $1 < i < n$ and $\pi_n(X) \approx \pi_n(\widetilde{X}) \approx \mathbb{Z}^\infty$, the direct sum of infinitely many $\mathbb{Z}$'s, by the Hurewicz theorem. Moreover, if $\iota : S^n \to S^1 \vee S^n$ is the inclusion and $\gamma : I \to S^1 \vee S^n$ is the loop which wraps once around the S^1 summand, then the classes $[\gamma^k\iota]$, $k \in \mathbb{Z}$, form a basis for $\pi_n(X) \approx \mathbb{Z}^\infty$ since the lift of $\gamma^k\iota$ to $\widetilde{X}$ starting at the basepoint $0 \in \mathbb{R} \subset \widetilde{X}$ is the inclusion $S^n_k \hookrightarrow \widetilde{X}$ connected to 0 by a path. Thus $\pi_n(X)$ is a free $\mathbb{Z}[\pi_1(X)]$-module on the single basis element $[\iota]$. Writing $[\gamma] = t$,

$\mathbb{Z}[\pi_1(X)]$ becomes $\mathbb{Z}[t,t^{-1}]$, the Laurent polynomials in t and t^{-1} with $\mathbb{Z}$ coefficients, and we have $\pi_n(X) \approx \mathbb{Z}[t,t^{-1}]$.

This example shows that the homotopy groups of a finite CW complex need not be finitely generated, in contrast to the homology groups. However, if we restrict attention to spaces with trivial action of π_1 on all π_n's, then a theorem of Serre, proved in Volume II, says that the homotopy groups are finitely generated iff the homology groups are finitely generated.

Since homology groups are usually much easier to compute than homotopy groups, the following version of Whitehead's theorem is often more readily applicable.

Corollary 4.12. *A map $f: X \to Y$ between simply-connected CW complexes is a homotopy equivalence if the induced maps $f_*: H_n(X) \to H_n(Y)$ are isomorphisms for all n.*

Proof: After replacing Y by the mapping cylinder M_f we may take f to be an inclusion $X \hookrightarrow Y$. Via the long exact sequences of homology and homotopy groups for the pair (Y,X), the corollary reduces, by Whitehead's theorem, to the statement that $\pi_n(Y,X) = 0$ for all n iff $H_n(Y,X) = 0$ for all n. Since $\pi_1(Y)$ and $\pi_0(X)$ are trivial by hypothesis, we have $\pi_1(Y,X) = 0$. The Hurewicz theorem then says that the first nonzero $\pi_n(Y,X)$ is isomorphic to the first nonzero $H_n(Y,X)$. □

In the nonsimply-connected case it is certainly possible for a map to induce isomorphisms on all homology groups but not on homotopy groups. Nonsimply-connected acyclic spaces, for which the inclusion of a point induces an isomorphism on homology, exhibit this phenomenon in its purest form. Perhaps the simplest nontrivial acyclic space is the 2-dimensional complex constructed in Example 3 of §2.2 with fundamental group $\langle\, a, b \mid a^5 = b^3 = (ab)^2 \,\rangle$ of order 120.

It is also possible for a map between spaces with abelian fundamental groups to induce isomorphisms on homology but not on higher homotopy groups, as the next example shows.

Example. We construct a space $X = (S^1 \vee S^n) \cup e^{n+1}$, for arbitrary $n > 1$, such that the inclusion $S^1 \hookrightarrow X$ induces an isomorphism on all homology groups and on π_i for $i < n$, but not on π_n. From the preceding Example we have $\pi_n(S^1 \vee S^n) \approx \mathbb{Z}[t,t^{-1}]$. Let X be obtained from $S^1 \vee S^n$ by attaching a cell e^{n+1} via a map $S^n \to S^1 \vee S^n$ corresponding to $2t-1 \in \mathbb{Z}[t,t^{-1}]$. By looking in the universal cover we see that $\pi_n(X) \approx \mathbb{Z}[t,t^{-1}]/(2t-1)$. Note that setting $t = \frac{1}{2}$ embeds $\mathbb{Z}[t,t^{-1}]/(2t-1)$ in $\mathbb{Q}$ as the subring $\mathbb{Z}[\frac{1}{2}]$ consisting of rationals with denominator a power of 2. From the long exact sequence of homotopy groups for the $(n-1)$-connected pair (X, S^1) we see that the inclusion $S^1 \hookrightarrow X$ induces an isomorphism on π_i for $i < n$. The fact that this inclusion also induces isomorphisms

on all homology groups can be deduced from cellular homology. The key point is that the cellular boundary map $H_{n+1}(X^{n+1}, X^n) \to H_n(X^n, X^{n-1})$ is an isomorphism since the degree of the composition of the attaching map $S^n \to S^1 \vee S^n$ of e^{n+1} with the collapse $S^1 \vee S^n \to S^n$ is $2 - 1 = 1$.

This example relies heavily on the nontriviality of the action of $\pi_1(X)$ on $\pi_n(X)$, so one might guess that the simple-connectivity assumption in Corollary 4.12 can be weakened to trivial action of π_1 on all π_n's. This is indeed the case, as we will show in Proposition 4.47.

We can apply the Hurewicz theorem to prove one of the basic homotopy-theoretic properties of CW complexes, the Cellular Approximation Theorem:

Theorem 4.13. *Every map $f : X \to Y$ of CW complexes is homotopic to a cellular map, i.e., one for which $f(X^n) \subset Y^n$ for all n. If f is already cellular on a subcomplex $A \subset X$, the homotopy may be taken to be stationary on A.*

It follows that maps of CW pairs $f : (X, A) \to (Y, B)$ can also be deformed to be cellular. One can first deform $f : A \to B$ to be cellular, then extend this to a homotopy of f on all of X, then deform the resulting map to be cellular staying fixed on A. This produces a homotopy $f_t : (X, A) \to (Y, B)$ from f to a cellular map. An easy consequence of this relative cellular approximation is:

Corollary 4.14. *If X is a CW complex with n-skeleton X^n, then the pair (X, X^n) is n-connected. Hence the inclusion $X^n \hookrightarrow X$ induces isomorphisms on π_i for $i < n$ and a surjection on π_n.*

Proof: The first statement follows by applying relative cellular approximation to maps $(D^i, \partial D^i) \to (X, X^n)$ for $i \leq n$. The second statement then comes from the long exact sequence of the pair (X, X^n). □

Proof of 4.13: Using the Hurewicz theorem we will prove the first assertion of the corollary, that (X, X^n) is n-connected, then we will deduce the cellular approximation theorem from this.

To show that (X, X^n) is n-connected it suffices to prove this for each component of X separately, so we may assume X is connected. The pair (X, X^n) is 0-connected since every point in X can be joined by a path to a 0-cell, and X^n contains all the 0-cells of X. The fact that (X, X^n) is 1-connected if $n \geq 1$ follows immediately from the long exact sequence of homotopy groups for this pair, since $\pi_1(X^n) \to \pi_1(X)$ is surjective when $n \geq 1$. To show $\pi_i(X, X^n) = 0$ for $i \leq n$ and $i \geq 2$, consider a map

$f:(D^i,\partial D^i)\to(X,X^n)$. This lifts to the universal cover $\widetilde{X}$, which contains the universal cover $\widetilde{X}^n$ since the map $\pi_1(X^n)\to\pi_1(X)$ is an isomorphism if $n\geq 2$. We can apply the Hurewicz theorem to $(\widetilde{X},\widetilde{X}^n)$ since $\widetilde{X}^n$ is simply-connected, to conclude that $(\widetilde{X},\widetilde{X}^n)$ is n-connected since its relative homology groups vanish through dimension n, $\widetilde{X}^n$ being the n-skeleton of $\widetilde{X}$. Thus the lift of f is homotopic rel ∂D^i to a map into $\widetilde{X}^n$, and this homotopy projects down to a homotopy of f which shows that $\pi_i(X,X^n)=0$.

Now to prove the cellular approximation theorem suppose inductively that we have already homotoped f so that $f(X^{n-1})\subset Y^{n-1}$. The composition of f with a characteristic map of an n-cell of X gives a map $(D^n,\partial D^n)\to(Y,Y^{n-1})$, and since (Y,Y^n) is n-connected by the preceding proposition, this map can be homotoped rel ∂D^n to a map with image in Y^n. Since this homotopy is fixed on ∂D^n it induces a homotopy of f on the n-cell. Doing this for all n-cells which do not already map to Y^n, and taking the constant homotopy on $X^{n-1}\cup A$, gives a homotopy of $f\big|(X^n\cup A)$. Extending this to a homotopy defined on all of X finishes the induction step. The infinite concatenation of these homotopies for varying n then produces a cellular map. □

At the beginning of §4.2 we will give a direct geometric proof of the cellular approximation theorem without using any homology theory. One can think of cellular approximation as an analog for CW complexes of simplicial approximation for simplicial complexes, Theorem 2.34, since simplicial maps are cellular. However, simplicial maps are much more rigid than cellular maps, which perhaps accounts for why subdivision of the domain is required for simplicial approximation but not for cellular approximation.

Eilenberg-MacLane Spaces

A space X having just one nontrivial homotopy group $\pi_n(X)\approx G$ is called an **Eilenberg-MacLane space**, or more concisely, a $K(G,n)$. The case $n=1$ was considered in §1.4b, where the condition that $\pi_i(X)=0$ for $i>1$ was replaced by the condition that X have a contractible universal cover, which is equivalent for spaces having the homotopy type of CW complexes.

We can build a CW complex $K(G,n)$ for arbitrary G and n, with G abelian if $n>1$, in the following way. To begin, let X^{n+1} be an $(n-1)$-connected CW complex of dimension $n+1$ such that $\pi_n(X^{n+1})\approx G$. When $n=1$ such a complex was constructed in Corollary 1.15, and when $n>1$ we can appeal to the Hurewicz theorem and take X^{n+1} to be a Moore space $M(G,n)$ as constructed in §2.2. The rest of the construction consists of attaching cells to X^{n+1} so as to make all the higher homotopy groups zero. This can be done inductively by attaching $(n+k+1)$-cells to X^{n+k} via cellular maps representing a

set of generators for $\pi_{n+k}(X^{n+k})$. The resulting CW complex X is then a $K(G,n)$, we claim. Since the attached cells all have dimension greater than $n+1$, they don't affect π_i for $i \le n$, so $\pi_i(X)$ is correct for $i \le n$. For $i > n$, every element of $\pi_i(X)$ is represented by a cellular map $S^i \to X^i$, and this is nullhomotopic in X^{i+1} and hence in X, since the induced map $\pi_i(X^i) \to \pi_i(X^{i+1})$ is zero by construction.

By taking products of $K(G,n)$'s for varying n we can then realize any sequence of groups G_n, abelian for $n > 1$, as the homotopy groups π_n of a space.

A fair number of $K(G,1)$'s arise naturally in a variety of contexts, and a few of these are mentioned in §1.4b. By contrast, naturally-occurring $K(G,n)$'s for $n \ge 2$ are rare. It seems the only real example is $\mathbb{C}\mathrm{P}^\infty$, which is a $K(\mathbb{Z},2)$ as we shall see in §4.3. One could of course trivially generalize this example by taking a product of $\mathbb{C}\mathrm{P}^\infty$'s to get a $K(G,2)$ with G a product of $\mathbb{Z}$'s.

Actually there is a fairly natural construction of a $K(\mathbb{Z},n)$, namely the infinite symmetric product $SP(S^n)$ defined in §3.4c where we showed that $SP(S^2) = \mathbb{C}\mathrm{P}^\infty$. However, it requires a fair bit of work to show that $SP(S^n)$ is a $K(\mathbb{Z},n)$ for $n > 2$. In §4.4j we prove a more general statement, a theorem of [Dold-Thom] asserting that the functor SP has the surprising property of converting homology groups into homotopy groups, namely $\pi_i\big(SP(X)\big) \approx H_i(X;\mathbb{Z})$ for all $i > 0$ and all connected CW complexes X. Taking X to be a sphere, we deduce that $SP(S^n)$ is a $K(\mathbb{Z},n)$, and more generally, $SP\big(M(G,n)\big)$ is a $K(G,n)$ for each Moore space $M(G,n)$.

We have observed that the homotopy groups of a space are not generally sufficient to determine its homotopy type. But this does hold for Eilenberg-MacLane spaces, and there is an analogous assertion that simply-connected Moore spaces are determined by their homology groups:

Proposition 4.15.

(a) The homotopy type of a CW complex $K(G,n)$ is uniquely determined by G and n.

(b) The homotopy type of a CW complex $M(G,n)$ is uniquely determined by G and n if $n > 1$.

The proof will be based on a more technical statement:

Lemma 4.16. *Let Z be a CW complex of the form $\big(\bigvee_\alpha S^n_\alpha\big) \bigcup_\beta e^{n+1}_\beta$ for some $n \ge 1$, and let X be any path-connected space. Then for every homomorphism $\psi : \pi_n(Z) \to \pi_n(X)$ there exists a map $f : Z \to X$ with $f_* = \psi$.*

Proof: To begin, let f send the natural basepoint of $\bigvee_\alpha S^n_\alpha$ to a chosen basepoint $x_0 \in X$. Extend f over each sphere S^n_α via a map representing $\psi([i_\alpha])$ where i_α is the inclusion

$S^n_\alpha \hookrightarrow Z$. Thus for the map $f : Z^n \to X$ constructed so far we have $f_*([i_\alpha]) = \psi([i_\alpha])$ for all α, hence $f_*([\varphi]) = \psi([\varphi])$ for all basepoint-preserving maps $\varphi : S^n \to Z^n$ since the i_α's generate $\pi_n(Z^n)$. To extend f over a cell e^{n+1}_β all we need is that the composition of the attaching map $\varphi_\beta : S^n \to Z^n$ for this cell with f be nullhomotopic in X. But this composition $f\varphi_\beta$ represents $f_*([\varphi_\beta]) = \psi([\varphi_\beta])$, and $[\varphi_\beta]$ is zero in $\pi_n(Z)$ since φ_β is nullhomotopic in Z via the characteristic map of e^{n+1}_β. Thus we obtain an extension $f : Z \to Y$. By cellular approximation, the elements $[i_\alpha]$ generate $\pi_n(Z)$ since they generate $\pi_n(Z^n)$. Hence $f_* = \psi$ on $\pi_n(Z)$. □

Proof of 4.15: Suppose X and Y are $K(G, n)$ CW complexes. Since homotopy equivalence is an equivalence relation, there is no loss of generality if we assume X is a particular $K(G, n)$, namely one constructed from a space Z as in the lemma by attaching cells of dimension $n + 2$ and greater. By the lemma there is a map $f : Z \to Y$ inducing an isomorphism on π_n. To extend this f over X we proceed inductively, first extending over the $(n+2)$-cells, then the $(n+3)$-cells, etc. For a cell e^{n+2} for example, the composition of its attaching map with f is nullhomotopic in Y since $\pi_{n+1}(Y) = 0$, so f extends over this cell. And similarly for the higher-dimensional cells. The resulting $f : X \to Y$ is a homotopy equivalence since it induces isomorphisms on all homotopy groups, π_n being the only nontrivial one.

For part (b), let Z be an $M(G, n)$ of the form considered in the lemma, and let Y be any other $M(G, n)$ CW complex. The lemma gives a map $f : Z \to Y$ inducing an isomorphism on π_n. We can enlarge Z and Y to $K(G, n)$'s by attaching cells of dimension $n+2$ and greater, and we can extend the map f to a homotopy equivalence between these $K(G, n)$'s as in the preceding paragraph. This extended f then induces isomorphisms on all homology groups, hence the original $f : Z \to Y$ induces an isomorphism on H_n since cells of dimension $n + 2$ and greater have no effect on H_n. The map $f : Z \to Y$ now induces an isomorphism on all homology groups, and is therefore a homotopy equivalence by the homology form of Whitehead's theorem since we are dealing with simply-connected spaces. □

The Hurewicz theorem gives a connection between homology and homotopy groups, and so, indirectly, a connection between cohomology and homotopy groups. But there is a more direct connection of a quite different sort. We shall show that for every CW complex X there is a natural bijection between $H^n(X; G)$ and the set $[X, K(G, n)]$ of homotopy classes of maps from X to a $K(G, n)$. That there should be any connection at all between cohomology and homotopy classes of maps is the first surprise here, and the second is

that Eilenberg-MacLane spaces are involved, since their definition is entirely in terms of homotopy groups, which on the face of it have nothing to do with cohomology.

It is easy enough to define maps $[X, K(G,n)] \to H^n(X;G)$. Just take any element $\iota \in H^n(K(G,n);G)$ and send a map $f : X \to K(G,n)$ to the element $f^*(\iota) \in H^n(X;G)$. The main theorem is then:

Theorem 4.17. *There is a class $\iota \in H^n(K;G)$ such that the map $[X, K(G,n)] \to H^n(X;G)$, $[f] \mapsto f^*(\iota)$, is a bijection for all CW complexes X.*

Such a class ι is called a **fundamental class**. Obviously a fundamental class is rather special, and an easy way to construct a special element $\iota \in H^n(K;G)$ is the following. Let us abbreviate notation by setting $K = K(G,n)$. Because K is $(n-1)$-connected, $H_{n-1}(K;\mathbb{Z}) = 0$. The universal coefficient theorem then gives an isomorphism $H^n(K;G) \approx Hom(H_n(K;\mathbb{Z}), G)$. Let ι be the element of $H^n(K;G)$ corresponding under this isomorphism to the inverse of the Hurewicz isomorphism $G = \pi_n(K) \xrightarrow{\approx} H_n(K)$. The proof of the theorem given below will show that this ι is a fundamental class.

There are several different ways of proving the theorem, in fact. The proof we give here has the advantage of being fairly direct, and it gives good geometric intuition into why the result should be true. However, the technical details of this proof are perhaps not as simple as one would like. A rather different proof which deduces the theorem as a corollary of a general result about arbitrary cohomology theories is given in §4.4e. Some readers may prefer this more abstract proof, with general constructions replacing explicit formulas.

Proof: We may assume the CW complex $K = K(G,n)$ has $(n-1)$-skeleton a point since $[X,K]$ depends only on the homotopy type of K. The class ι defined above is then represented by the cellular cochain assigning to each n-cell of K the element of $\pi_n(K)$ represented by the characteristic map of the cell. This can be seen from the diagram

$$\begin{array}{ccccccc}
 & & \pi_n(K^n, K^{n-1}) & \longrightarrow & \pi_n(K^{n+1}, K^{n-1}) & \approx & \pi_n(K) \\
 & & \Big\downarrow{\scriptstyle h} & & \Big\downarrow{\scriptstyle h}\approx & & \Big\downarrow{\scriptstyle h}\approx \\
H_{n+1}(K^{n+1}, K^n) & \xrightarrow{\partial} & H_n(K^n, K^{n-1}) & \longrightarrow & H_n(K^{n+1}, K^{n-1}) & \approx & H_n(K)
\end{array}$$

where the maps h are Hurewicz homomorphisms. By definition, ι is the inverse of the right-hand vertical isomorphism, so as a cellular cohomology class this is represented by the cochain which is the composition $H_n(K^n, K^{n-1}) \to \pi_n(K)$ across the bottom of the diagram. A basis for $H_n(K^n, K^{n-1})$ is represented by the Hurewicz images of the elements of $\pi_n(K^n, K^{n-1})$ represented by the characteristic maps of the n-cells, so commutativity of the diagram gives the result.

If $f: X \to K$ is cellular, it takes the $(n-1)$-skeleton X^{n-1} to a point. Then the result in the previous paragraph implies that $f^*(\iota)$ is represented by the cellular cochain c_f assigning to each n-cell e^n_α of X the element of $\pi_n(K)$ determined by $f\big|e^n_\alpha$, i.e., the composition $f\Phi_\alpha$ where Φ_α is the characteristic map of e^n_α. Note that c_f can also be viewed as the composition

$$H_n(X^n/X^{n-1}) \xrightarrow{h^{-1}} \pi_n(X^n/X^{n-1}) \xrightarrow{\overline{f}_*} \pi_n(K)$$

where $\overline{f}: X^n/X^{n-1} \to K$ is the map induced by f. When $n = 1$ one should replace $\pi_n(X^n/X^{n-1})$ here by its abelianization in order to get the Hurewicz isomorphism $H_n(X^n/X^{n-1}) \approx \pi_n(X^n/X^{n-1})$.

In particular, c_f is defined when $X = X^n$, or in other words, when we are only given a cellular map $f: X^n \to K$. The key observation in this situation is:

$(*)$ A cellular map $f: X^n \to K$ extends over X^{n+1}, and hence over X, iff the cellular cochain c_f is a cellular cocycle.

To show this, observe first that f extends over a cell e^{n+1}_β iff the composition $f\varphi_\beta$ is nullhomotopic, where φ_β is the attaching map of e^{n+1}_β. Now let us compute the value of the cellular coboundary δc_f on the cell e^{n+1}_β:

$$\begin{aligned}
\delta c_f(e^{n+1}_\beta) &= c_f(\partial e^{n+1}_\beta) \\
&= c_f\big(h([q\varphi_\beta])\big) \quad \text{where } q: X^n \to X^n/X^{n-1} \text{ is the quotient map} \\
&= [\overline{f} q\varphi_\beta] \quad \text{since } c_f = \overline{f}_* h^{-1} \\
&= [f\varphi_\beta]
\end{aligned}$$

Thus c_f is a cocycle iff all the obstructions $[f\varphi_\beta]$ to extending f over $(n+1)$-cells vanish. Once f has been extended over X^{n+1} there are no further obstructions to extending inductively over higher cells since $\pi_i(K) = 0$ for $i > n$.

Surjectivity of the map $[X, K(G,n)] \to H^n(X;G)$, $[f] \mapsto f^*(\iota)$, can now be deduced easily from $(*)$. Represent a class in $H^n(X;G)$ by a cellular cocycle c. Define a cellular map $f: X^n \to K$ sending X^{n-1} to the basepoint and mapping each cell e^n_α according to the value $c(e^n_\alpha) \in \pi_n(K)$. Since c is a cocycle, the resulting map f extends over X by $(*)$, and $c_f = c$ by construction.

To show injectivity, suppose we have maps $f, g: X \to K$ with $c_f - c_g = \delta\eta$ for some cellular $(n-1)$-cochain η. We will use this data to construct a homotopy $F: X \times I \to K$ from f to g. On $X \times \{0\}$ and $X \times \{1\}$, F is given by f and g, so it remains to extend

F over cells $e^i_\gamma \times (0,1)$. For $i < n-1$ we map such cells to the basepoint, to make F cellular. On a cell $e^{n-1}_\gamma \times (0,1)$, whose boundary lies in the $(n-1)$-skeleton of $X \times I$ and hence maps to the basepoint, we let F represent the class $\varepsilon\eta(e^{n-1}_\gamma) \in \pi_n(K)$ where $\varepsilon = \pm 1$ is a sign to be determined later. To extend F over cells $e^n_\alpha \times (0,1)$ we shall apply $(*)$, and for this we need the formula

$$\partial\big(e^n_\alpha \times (0,1)\big) = \partial e^n_\alpha \times (0,1) \pm (e^n_\alpha \times \{1\} - e^n_\alpha \times \{0\})$$

This special case of Proposition 3.35 can be proved directly by considering the geometric decomposition $\partial(D^n \times I) = \partial D^n \times I \cup D^n \times \partial I$. The subspace $D^n \times \partial I$ clearly gives the terms $e^n_\alpha \times \{1\} - e^n_\alpha \times \{0\}$ in $\partial\big(e^n_\alpha \times (0,1)\big)$. To obtain the remaining term $\partial e^n_\alpha \times (0,1)$, recall how ∂e^n_α is defined, as a sum $\sum_\gamma n_{\alpha\gamma} e^{n-1}_\gamma$ with $n_{\alpha\gamma}$ the degree of a certain map $\Delta_{\alpha\gamma}$, the composition $S^{n-1} \xrightarrow{\varphi_\alpha} X^{n-1} \longrightarrow X^{n-1}/X^{n-2} \longrightarrow S^{n-1}_\gamma$ where the last map collapses all $(n-1)$-cells except e^{n-1}_γ. For $e^n_\alpha \times (0,1)$ the corresponding map, after we collapse the subspaces $D^n \times \{0\}$ and $D^n \times \{1\}$ of $\partial(D^n \times I)$ to points, is the suspension of $\Delta_{\alpha\gamma}$, having the same degree $n_{\alpha\gamma}$, so from $\partial D^n \times I$ we get a contribution of $\sum_\gamma n_{\alpha\gamma}\big(e^{n-1}_\gamma \times (0,1)\big)$ to $\partial(e^n_\alpha \times (0,1))$.

Now to apply $(*)$ we compute:

$$\begin{aligned}
\delta c_F(e^n_\alpha \times (0,1)\big) &= c_F\big(\partial(e^n_\alpha \times (0,1))\big) \\
&= \sum_\beta n_{\alpha\beta} c_F\big(e^{n-1}_\beta \times (0,1)\big) \pm \big(c_F(e^n_\alpha \times \{1\}) - c_F(e^n_\alpha \times \{0\})\big) \\
&= \sum_\beta n_{\alpha\beta} \varepsilon\eta(e^{n-1}_\beta) \pm \big(c_g(e^n_\alpha) - c_f(e^n_\alpha)\big) \\
&= \varepsilon\eta(\partial e^n_\alpha) \pm \big(c_g(e^n_\alpha) - c_f(e^n_\alpha)\big) \\
&= \varepsilon\delta\eta(e^n_\alpha) \pm \big(c_g(e^n_\alpha) - c_f(e^n_\alpha)\big)
\end{aligned}$$

So with the right choice of the sign ε this is 0 and we can extend F over $e^n_\alpha \times (0,1)$. Thus F extends over $X^n \times I$, and then there are no further obstructions to extending over $X \times I$. □

For connected CW complexes X the theorem also holds when $[X, K(G,n)]$ is replaced by $\langle X, K(G,n)\rangle$, the basepoint-preserving homotopy classes, since $[X, K(G,n)] = \langle X, K(G,n)\rangle$ by Proposition 4.50 when $n > 1$, and the case $n = 1$ is an exercise for §4.4a, using the assumption that G is abelian. Alternatively, the proof of the theorem can easily be adjusted to use $\langle X, K(G,n)\rangle$ instead of $[X, K(G,n)]$.

This theorem leads to a basic principle which reappears a number of times in algebraic topology, the idea that the occurence or non-occurence of a certain phenomenon is governed

by what happens in a single special case, the *universal example*. To illustrate, let us prove the following special fact:

(∗) For all spaces X the map $H^1(X;\mathbb{Z}) \to H^2(X;\mathbb{Z})$, $\alpha \mapsto \alpha^2$, is identically zero.

By taking a CW approximation to X, as described in the next section, this reduces to the case that X is a CW complex. Then every $\alpha \in H^1(X;\mathbb{Z})$ is $f^*(\iota)$ for some $f : X \to K(\mathbb{Z},1)$, further reducing us to proving the result just for the element $\iota \in H^1(K(\mathbb{Z},1);\mathbb{Z})$, the 'universal example.' And this special case is trivial since S^1 is a $K(\mathbb{Z},1)$ and $H^2(S^1;\mathbb{Z}) = 0$. Note that this argument uses only the surjectivity of $[X, K(G,n)] \to H^n(X;G)$, which was quite a bit easier to prove than injectivity. (The surjectivity proof could be further simplified in the case $n = 1$ by using the fact that $H^1(X;G) = Hom(H_1(X), G) = Hom(\pi_1(X), G)$ when X is path-connected.)

Does (∗) generalize? It certainly does not hold if we replace $\mathbb{Z}$ coefficients by $\mathbb{Z}_2$ since $H^*(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2) = \mathbb{Z}_2[\alpha]$. Indeed, the example of $\mathbb{R}\mathrm{P}^\infty$ shows more generally that the fundamental class $\iota \in H^n(K(\mathbb{Z}_2,n);\mathbb{Z}_2)$ generates a polynomial subalgebra $\mathbb{Z}_2[\iota] \subset H^*(K(\mathbb{Z}_2,n);\mathbb{Z}_2)$ for each $n \geq 1$, since there is a map $f : \mathbb{R}\mathrm{P}^\infty \to K(\mathbb{Z}_2,n)$ with $f^*(\iota) = \alpha^n$ and all the powers of α^n are nonzero, hence also all the powers of ι. By the same reasoning, the example of $\mathbb{C}\mathrm{P}^\infty$ shows that the class $\iota \in H^{2n}(K(\mathbb{Z},2n);\mathbb{Z})$ generates a polynomial subalgebra $\mathbb{Z}[\iota]$ in $H^*(K(\mathbb{Z},2n);\mathbb{Z})$. As we shall see in Volume II, $H^*(K(\mathbb{Z},2n);\mathbb{Z})/torsion$ is exactly this polynomial algebra $\mathbb{Z}[\iota]$.

A little more subtle is the question of identifying the subalgebra of $H^*(K(\mathbb{Z},n);\mathbb{Z})$ generated by ι for odd $n \geq 3$. By the commutativity property of cup products we know that ι^2 is either zero or of order two. To see that ι^2 is nonzero it suffices to find a single space X with an element $\gamma \in H^n(X;\mathbb{Z})$ such that $\gamma^2 \neq 0$. The first place to look might be $\mathbb{R}\mathrm{P}^\infty$, but its cohomology with $\mathbb{Z}$ coefficients is concentrated in even dimensions. Instead, consider $X = \mathbb{R}\mathrm{P}^\infty \times \mathbb{R}\mathrm{P}^\infty$. This has $\mathbb{Z}_2$-cohomology $\mathbb{Z}_2[\alpha,\beta]$ and we show in an example in §3.4e that its $\mathbb{Z}$-cohomology is the $\mathbb{Z}_2[\alpha^2,\beta^2]$-submodule generated by 1 and $\alpha^2\beta + \alpha\beta^2$, except in dimension zero of course, where 1 generates a $\mathbb{Z}$ rather than a $\mathbb{Z}_2$. In particular we can take $\gamma = \alpha^{2k}(\alpha^2\beta + \alpha\beta^2)$ for any $k \geq 0$, and then all powers γ^m are nonzero since we are inside the polynomial ring $\mathbb{Z}_2[\alpha,\beta]$. It follows that the subalgebra of $H^*(K(\mathbb{Z},n);\mathbb{Z})$ generated by ι is $\mathbb{Z}[\iota]/(2\iota^2)$ for odd $n \geq 3$.

These examples lead one to wonder just how complicated the cohomology of $K(G,n)$'s is. If one works directly from the general construction of a $K(G,n)$ one doesn't get very far. Consider the case $G = \mathbb{Z}$ for example. Here one would start with S^n and attach $(n+2)$-cells to kill $\pi_{n+1}(S^n)$. Since $\pi_{n+1}(S^n)$ happens to be cyclic, only one $(n+2)$-cell

is needed. To continue, one would have to compute generators for π_{n+2} of the resulting space $S^n \cup e^{n+2}$, use these to attach $(n+3)$-cells, then compute the resulting π_{n+3}, etc. When $n = 2$ this procedure happens to work out very neatly, and the resulting $K(\mathbb{Z}, 2)$ is $\mathbb{C}\mathrm{P}^\infty$ with its usual CW structure having one cell in each even dimension, according to an exercise at the end of the section. However, for larger n it quickly becomes impractical to make this procedure explicit since homotopy groups are so hard to compute. One can get some idea of the difficulties of the next case $n = 3$ by considering the homology groups of $K(\mathbb{Z}, 3)$. Using techniques in Volume II, the groups $H_i(K(\mathbb{Z}, 3); \mathbb{Z})$ for $0 \le i \le 12$ can be computed to be:

$$\mathbb{Z},\ 0,\ 0,\ \mathbb{Z},\ 0,\ \mathbb{Z}_2,\ 0,\ \mathbb{Z}_3,\ \mathbb{Z}_2,\ \mathbb{Z}_2,\ \mathbb{Z}_3,\ \mathbb{Z}_{10},\ \mathbb{Z}_2$$

To get this sequence of homology groups would require quite a few cells, and the situation only gets worse in higher dimensions, where the homology groups are not always cyclic.

Indeed, one might guess that computing the homology groups of $K(\mathbb{Z}, n)$'s would be of the same order of difficulty as computing the homotopy groups of spheres, but by some miracle this is not the case. The calculations are indeed complicated, but they were completely done by Serre and Cartan in the 1950's, not just for $K(\mathbb{Z}, n)$'s, but for all $K(G, n)$'s with G finitely generated abelian. For example, $H^*(K(\mathbb{Z}, 3); \mathbb{Z}_2)$ is the polynomial algebra $\mathbb{Z}_2[x_3, x_5, x_9, x_{17}, x_{33}, \cdots]$ with generators of dimensions $2^i + 1$, indicated by the subscripts. And in general, for G finitely generated abelian, $H^*(K(G, n); \mathbb{Z}_p)$ is a polynomial algebra on generators of specified dimensions if $p = 2$, while for p an odd prime one gets the tensor product of a polynomial ring on generators of specified even dimensions and an exterior algebra on generators of specified odd dimensions. With $\mathbb{Z}$ coefficients the description of the cohomology is not nearly so neat, however. We will study these questions in some detail in Volume II.

Exercises

1. Suppose a sum $f +' g$ of maps $f, g : (I^n, \partial I^n) \to (X, x_0)$ is defined using a coordinate of I^n other than the first coordinate as in the usual sum $f + g$. Show that $(f + g) +' (h + k) = (f +' h) + (g +' k)$, and deduce that $f +' k \simeq f + k$ so the two sums agree on $\pi_n(X, x_0)$, and also that $g +' h \simeq h + g$ so the addition is abelian.

2. Show that a homotopy equivalence $\varphi : X \to Y$ induces isomorphisms $\varphi_* : \pi_n(X, x_0) \to \pi_n(Y, \varphi(x_0))$ for all n. [See the exercises for §1.1 for the case $n = 1$.]

3. For an H-space (X, x_0) with multiplication $\mu : X \times X \to X$, show that the group operation in $\pi_n(X, x_0)$ can also be defined by the rule $(f + g)(x) = \mu\big(f(x), g(x)\big)$.

4. Let $p:\widetilde{X}\to X$ be the universal cover of a path-connected space X. Show that under the isomorphism $\pi_n(X)\approx\pi_n(\widetilde{X})$, which holds for $n\geq 2$, the action of $\pi_1(X)$ on $\pi_n(X)$ corresponds to the action of $\pi_1(X)$ on $\pi_n(\widetilde{X})$ induced by the action of $\pi_1(X)$ on $\widetilde{X}$ as deck transformations. More precisely, prove a formula like $\gamma p_*(\alpha)=p_*\big(\beta_\gamma(\gamma_*(\alpha))\big)$ where $\gamma\in\pi_1(X,x_0)$, $\alpha\in\pi_n(\widetilde{X},\widetilde{x}_0)$, and γ_* denotes the homomorphism induced by the action of γ on $\widetilde{X}$.

5. For a pair (X,A) of path-connected spaces, show that $\pi_1(X,A,x_0)$ can be identified in a natural way with the set of cosets αH of the subgroup $H\subset\pi_1(X,x_0)$ represented by loops in A at x_0.

6. If $p:(\widetilde{X},\widetilde{A},\widetilde{x}_0)\to(X,A,x_0)$ is a covering space with $\widetilde{A}=p^{-1}(A)$, show that the map $p_*:\pi_n(\widetilde{X},\widetilde{A},\widetilde{x}_0)\to\pi_n(X,A,x_0)$ is an isomorphism for all $n>1$.

7. Extend the results proved near the beginning of this section for the change-of-basepoint maps β_γ for absolute homotopy groups to the case of relative homotopy groups.

8. From Lemma 4.9 deduce that the image of the map $\pi_2(X,x_0)\to\pi_2(X,A,x_0)$ lies in the center of $\pi_2(X,A,x_0)$.

9. Construct a CW pair (X,A) such that $\pi_1(A)\approx\pi_1(X)$ acts trivially on $\pi_2(A)$ and on $\pi_2(X,A)$ but not on $\pi_2(X)$.

10. Suppose we define $\pi_0(X,A,x_0)$ to be the quotient set $\pi_0(X,x_0)/\pi_0(A,x_0)$, so that the long exact sequence for the pair (X,A) extends to $\cdots\to\pi_0(X,x_0)\to\pi_0(X,A,x_0)\to 0$.

(a) Show that with this extension, the five-lemma holds for the map of long exact sequences induced by a map $(X,A,x_0)\to(Y,B,y_0)$, in the following form: Each map between the two sequences is a bijection if the two adjacent maps on each side of the given map are bijections for all choices of x_0.

(b) Show that the long exact sequence of a triple (X,A,B,x_0) can be extended only to the term $\pi_0(A,B,x_0)$ in general, and that the five-lemma holds for this extension.

11. Show the 'quasi-circle' described in an exercise in §1.3 has trivial homotopy groups but is not contractible, hence does not have the homotopy type of a CW complex.

12. Show that a CW complex is contractible if it is the union of an increasing sequence of subcomplexes $X_1\subset X_2\subset\cdots$ such that each X_i is contractible in X_{i+1}, i.e., the inclusion $X_i\hookrightarrow X_{i+1}$ is nullhomotopic. This applies for example to S^∞, or more generally to the infinite suspension $S^\infty X$ of any CW complex X.

13. Show that an n-connected, n-dimensional CW complex is contractible.

14. Given a CW pair (X,A) and a map $f:A\to Y$ with Y path-connected, show that there is an extension of f to a map $X\to Y$ provided that $\pi_n(Y)=0$ for all n such

that $X - A$ has cells of dimension $n+1$. Deduce that a CW complex retracts onto any contractible subcomplex.

15. Use homotopy groups to show there is no retraction $\mathbb{R}\mathrm{P}^n \to \mathbb{R}\mathrm{P}^k$ if $n > k > 0$.

16. Show the action of $\pi_1(\mathbb{R}\mathrm{P}^n)$ on $\pi_n(\mathbb{R}\mathrm{P}^n) \approx \mathbb{Z}$ is trivial for n odd and nontrivial for n even.

17. Let $X \subset \mathbb{R}^{n+1}$ be the union of the infinite sequence of spheres S^n_k of radius $1/k$ and center $(1/k, 0, \cdots, 0)$. Show that $\pi_i(X) = 0$ for $i < n$ and construct a homomorphism from $\pi_n(X)$ onto the direct product (not direct sum) of the groups $\pi_n(S^n_k) \approx \mathbb{Z}$. [Is this surjection an isomorphism if $n > 1$?]

18. Show the relative Hurewicz theorem in dimension n implies the absolute Hurewicz theorem in dimension $n-1$ by considering the pair (CX, X) where CX is the cone on X.

19. Show the suspension of an acyclic CW complex is contractible.

20. Show that a map between simply-connected CW complexes is a homotopy equivalence if its mapping cone is contractible. Use the preceding exercise to give an example where this fails in the nonsimply-connected case.

21. Let the CW complex X be obtained from $S^1 \vee S^n$, $n \geq 2$, by attaching a cell e^{n+1} by a map representing the polynomial $p(t) \in \mathbb{Z}[t, t^{-1}] \approx \pi_n(S^1 \vee S^n)$, so $\pi_n(X) \approx \mathbb{Z}[t, t^{-1}]/\big(p(t)\big)$. Show $\pi'_n(X)$ is cyclic and compute its order in terms of $p(t)$. Give examples showing that the group $\pi_n(X)$ can be finitely generated or not, independently of whether $\pi'_n(X)$ is finite or infinite.

22. Let X be a connected CW complex with 1-skeleton X^1. Show that $\pi_2(X, X^1) \approx \pi_2(X) \times K$ where K the kernel of $\pi_1(X^1) \to \pi_1(X)$, a free group. Also, show that $\pi'_2(X, X^1) \approx \pi'_2(X) \times K'$ where K' denotes the quotient of K obtained by factoring out the conjugation action of $\pi_1(X^1)$. Note that K' is abelian.

23. Show that if X and Y are homotopy equivalent CW complexes, then their n-skeleta X^n and Y^n are also homotopy equivalent if X and Y have no cells of dimension $n+1$. [Use cellular approximation.]

24. Show that a map $f : X \to Y$ of connected CW complexes is a homotopy equivalence if it induces an isomorphism on π_1 and if a lift $\widetilde{f} : \widetilde{X} \to \widetilde{Y}$ to the universal covers induces an isomorphism on homology. [The latter condition can be rephrased in terms of homology with local coefficients as saying that $f_* : H_*(X; \mathbb{Z}[\pi_1 X]) \to H_*(Y; \mathbb{Z}[\pi_1 Y])$ is an isomorphism; see §3.4i.]

25. Show that a map between connected n-dimensional CW complexes is a homotopy equivalence if it induces an isomorphism on π_i for $i \le n$. [Pass to universal covers and use homology.]

26. If an n-dimensional CW complex X contains a subcomplex Y homotopy equivalent to S^n, show that the map $\pi_n(Y) \to \pi_n(X)$ induced by inclusion is injective. [Hint: Hurewicz]

27. Show that a closed simply-connected 3-manifold is homotopy equivalent to S^3. [Use Poincaré duality, and also the fact that closed manifolds are homotopy equivalent to CW complexes, from Corollary A.11 in the Appendix. The much stronger statement that a closed simply-connected 3-manifold is homeomorphic to S^3 is still unproved. This is the **Poincaré conjecture**, without doubt the most famous open problem in topology. The higher-dimensional analog, that a closed n-manifold homotopy equivalent to S^n is homeomorphic to S^n, has been proved for all $n \ge 4$. This is one of the high points of manifold theory, with major contributions by many topologists.]

28. Show that the closed surfaces with infinite fundamental group are $K(\pi, 1)$'s by showing that their universal covers are contractible, via the Hurewicz theorem and results of §3.3.

29. If X and Y are simply-connected CW complexes such that $\widetilde{H}_i(X)$ and $\widetilde{H}_j(Y)$ are finite and of relatively prime orders for all pairs (i, j), show that the inclusion $X \vee Y \hookrightarrow X \times Y$ is a homotopy equivalence and $X \wedge Y$ is contractible. [Use the Künneth formula.]

30. If X is a $K(G, 1)$ CW complex, show that $\pi_n(X^n)$ is free abelian for $n \ge 2$.

31. Let G be a group and X a simply-connected space. Show that for the product $K(G, 1) \times X$ the action of π_1 on π_n is trivial for all $n > 1$.

32. Given a sequence of CW complexes $K(G_n, n)$, $n = 1, 2, \cdots$, let X_n be the CW product of the first n of these complexes, and let X be the CW complex which is the union of subcomplexes X_n via the inclusions $X_{n-1} \subset X_n$ as the subcomplex with n^{th} coordinate a basepoint 0-cell of $K(G_n, n)$. Show $\pi_n(X) \approx G_n$ for all n.

33. Show that $H_{n+1}(K(G, n); \mathbb{Z}) = 0$ if $n > 1$. [Build a $K(G, n)$ from a Moore space $M(G, n)$ by attaching cells of dimension $> n + 1$.]

34. Extend the Hurewicz theorem by showing that if X is an $(n-1)$-connected CW complex, then the Hurewicz homomorphism $h : \pi_{n+1}(X) \to H_{n+1}(X)$ is surjective when $n > 1$, and when $n = 1$ show there is an isomorphism $H_2(X)/h\big(\pi_2(X)\big) \approx H_2\big(K(\pi_1(X), 1)\big)$. [Build a $K(\pi_n(X), n)$ from X by attaching cells of dimension greater than $n + 1$ and consider the homology sequence of the pair (Y, X) where Y is X with the $(n+2)$-cells of $K(\pi_n(X), n)$ attached. Note that the image of the boundary map $H_{n+2}(Y, X) \to H_n(X)$

coincides with the image of h, and $H_{n+1}(Y) \approx H_{n+1}\big(K(\pi_n(X), n)\big)$. The previous exercise is needed for the case $n > 1$.]

35. Show there is a Moore space $M(G, 1)$ with $\pi_1\big(M(G, 1)\big) \approx G$ iff $H_2(K(G, 1); \mathbb{Z}) = 0$. [Use the preceding problem. Build such an $M(G, 1)$ from the 2-skeleton K^2 of a $K(G, 1)$ by attaching 3-cells according to a basis for the free group $H_2(K^2; \mathbb{Z})$.] In particular, there is no $M(\mathbb{Z}^n, 1)$ with fundamental group $\mathbb{Z}^n$, free abelian of rank n, if $n \geq 2$.

36. Let X be a CW complex with $\pi_i(X) = 0$ for $1 < i < n$ for some $n \geq 2$. Show that $H_n(X)/h\big(\pi_n(X)\big) \approx H_n\big(K(\pi_1(X), 1)\big)$, where h is the Hurewicz homomorphism.

37. Generalizing the example of $\mathbb{R}\mathrm{P}^2$ and $S^2 \times \mathbb{R}\mathrm{P}^\infty$, show that if X is a finite-dimensional connected CW complex with universal cover $\widetilde{X}$, then X and $\widetilde{X} \times K(\pi_1(X), 1)$ have isomorphic homotopy groups but are not homotopy equivalent if $\pi_1(X)$ contains elements of finite order.

38. Show that the group structure on S^1 (multiplication in $\mathbb{C}$ restricted to the unit circle) induces a group structure on $[X, S^1]$, such that the bijection $[X, S^1] \to H^1(X; \mathbb{Z})$ of Theorem 4.17 is an isomorphism.

39. Suppose that a CW complex X contains a subcomplex S^1 such that the inclusion $S^1 \hookrightarrow X$ induces an injection $H_1(S^1; \mathbb{Z}) \to H_1(X; \mathbb{Z})$ with image a direct summand of $H_1(X; \mathbb{Z})$. Show that S^1 is a retract of X.

40. Show that, given abelian groups G and H and CW complexes $K(G, n)$ and $K(H, n)$, the map $[K(G, n), K(H, n)] \to Hom(G, H)$, $[f] \mapsto f_* : \pi_n(K(G, n)) \to \pi_n(K(H, n))$, is a bijection.

41. Show that $[X, S^n] \approx H^n(X; \mathbb{Z})$ if X is an n-dimensional CW complex. [Build a $K(\mathbb{Z}, n)$ from S^n by attaching cells of dimension $\geq n + 2$.]

42. For an abelian group G, Theorem 4.17 gives a map $\mu : K(G, n) \times K(G, n) \to K(G, n)$ with $\mu^*(\iota) = \iota \otimes 1 + 1 \otimes \iota$ where ι is the fundamental class for $K(G, n)$. Show that μ defines an H-space structure on $K(G, n)$ which is, up to homotopy, commutative and associative. Show also that the H-space multiplication μ is unique up to homotopy.

43. Using an H-space multiplication μ on $K(G, n)$, define an addition in $[X, K(G, n)]$ by $[f] + [g] = [\mu(f, g)]$ and show that under the bijection $H^n(X; G) \approx [X, K(G, n)]$ this addition corresponds to the usual addition in cohomology.

44. Show that the group $\mathbb{Z}_p \times \mathbb{Z}_p$ with p prime cannot act freely on any sphere S^n, by filling in details of the following argument. Such an action would define a covering space $S^n \to M$ with M a closed manifold. When $n > 1$, build a $K(\mathbb{Z}_p \times \mathbb{Z}_p, 1)$ from

M by attaching a single $(n+1)$-cell and then cells of higher dimension. Deduce that $H^{n+1}(K(\mathbb{Z}_p \times \mathbb{Z}_p, 1); \mathbb{Z}_p)$ is $\mathbb{Z}_p$ or 0, a contradiction. (The case $n = 1$ is more elementary.)

45. Finish the homotopy classification of lens spaces which was begun in exercise 2 of §3.4e by showing that $L_m(\ell_1, \cdots, \ell_n) \simeq L_m(\ell'_1, \cdots, \ell'_n)$ if $\ell_1 \cdots \ell_n \equiv \pm \ell'_1 \cdots \ell'_n k^n \mod m$ for some integer k, via the following steps:

(a) Reduce to the case $k = 1$ by showing that $L_m(\ell'_1, \cdots, \ell'_n) = L_m(k\ell'_1, \cdots, k\ell'_n)$ if k is relatively prime to m. [Rechoose the generator of the $\mathbb{Z}_m$ action on S^{2n-1}.]

(b) Let $f : L \to L'$ be a map constructed as in part (b) of the exercise in §3.4e. Construct a map $g : L \to L'$ as a composition $L \longrightarrow L \vee S^{2n-1} \longrightarrow L \vee S^{2n-1} \longrightarrow L'$ where the first map collapses the boundary of a small ball to a point, the second map is the wedge of the identity on L and a map of some degree d on S^{2n-1}, and the third map is f on L and the projection $S^{2n-1} \to L'$ on S^{2n-1}. Show that g has degree $k_1 \cdots k_n + dm$, i.e., g induces multiplication by $k_1 \cdots k_n + dm$ on $H_{2n-1}(-; \mathbb{Z})$. [Show first that a lift of g to the universal cover S^{2n-1} has this degree.]

(c) If $\ell_1 \cdots \ell_n \equiv \pm \ell'_1 \cdots \ell'_n \mod m$, choose d so that $k_1 \cdots k_n + dm = \pm 1$ and show this implies that g induces an isomorphism on all homotopy groups, hence is a homotopy equivalence. [For π_i with $i > 1$, consider a lift of g to the universal cover.]

2. Homotopy Properties of CW Complexes

In this section we use direct geometric arguments to prove some basic homotopy-theoretic properties of CW complexes. In particular we prove an excision property of homotopy groups which leads to the concept of stable homotopy groups and stable homotopy theory. Along the way we digress to give a second proof of the Hurewicz theorem.

Cellular Approximation

In the previous section we gave a proof of the cellular approximation theorem, Theorem 4.13, as an application of the Hurewicz theorem and basic properties of homology. Now we shall give a direct geometric proof which perhaps explains better why this fundamental result is true. In addition, the main lemma in this proof will be a key step in the proof of the homotopy excision property later in the section.

Recall that map $f: X \to Y$ between CW complexes is cellular if $f(X^n) \subset Y^n$ for all n. The cellular approximation theorem then says that every map between CW complexes is homotopic to a cellular map. Moreover, if the map is already cellular on a subcomplex, the homotopy may be taken to be stationary on this subcomplex. As pointed out following Theorem 4.13, it follows that maps of CW pairs can also be homotoped to cellular maps.

Note that the cellular approximation theorem immediately imples that $\pi_i(S^n) = 0$ for $i < n$. For if S^i and S^n are given their usual CW structures, with the 0-cells as basepoints, then every basepoint-preserving map $S^i \to S^n$ can be homotoped, fixing the basepoint, to be cellular, and hence constant if $i < n$.

The direct geometric proof of cellular approximation proceeds in the following way. Suppose inductively that $f: X \to Y$ is already cellular on the skeleton X^{n-1}, and let e^n be an n-cell of X. Since e^n has compact closure in X, $f(e^n)$ has compact closure in Y, hence meets only finitely many cells of Y. Let $e^m \subset Y$ be a cell of highest dimension meeting $f(e^n)$. We may assume $m > n$, otherwise f is already cellular on e^n. The main lemma below will imply that we can deform $f \big| X^{n-1} \cup e^n$, staying fixed on X^{n-1}, so that $f(e^n)$ misses some point $p \in e^m$. Once this is accomplished, we can then deform $f \big| X^{n-1} \cup e^n$ so that $f(e^n)$ misses the whole cell e^m by composing with a deformation retraction of $Y^m - \{p\}$ onto $Y^m - e^m$. By finitely many iterations of this process we eventually make $f(e^n)$ miss all cells of dimension greater than n. Doing this for all n-cells simultaneously, staying fixed on n-cells in A where f is already cellular, we obtain a homotopy of $f \big| X^n$ rel $X^{n-1} \cup A^n$ to a cellular map. The induction step is then completed by extending this homotopy, together with the constant homotopy on A, to a homotopy defined on all of X. Letting n go to ∞, the resulting possibly infinite string of

homotopies can be realized as a single homotopy by performing the n^{th} homotopy during the t-interval $[1-1/2^n, 1-1/2^{n+1}]$. This makes sense since each point of X lies in some X^n, which is eventually stationary in the infinite chain of homotopies.

The missing step in this argument is filled in using:

Lemma 4.18. *Let $f: I^n \to Z$ be a map, where Z is obtained from a subspace W by attaching a cell e^m. Then there is a homotopy f_t of f, stationary on $f^{-1}(W)$, to a map f_1 such that there is a simplex $\Delta^m \subset e^m$ with $f_1^{-1}(\Delta^m)$ a union of finitely many convex polyhedra, on each of which f_1 is the restriction of a surjective linear map $\mathbb{R}^n \to \mathbb{R}^m$.*

Here a convex polyhedron in $I^n \subset \mathbb{R}^n$ is a subspace defined by a finite number of linear inequalities of the form $\sum_i a_i x_i \le b$.

Before proving the lemma, let us see how it fills in the missing step in the proof of cellular approximation. Composing the given map $f: X^{n-1} \cup e^n \to Y^m$ with a characteristic map $I^n \to X$ for e^n, we obtain a map f as in the lemma, with $Z = Y^m$ and $W = Y^m - e^m$. The homotopy given by the lemma is fixed on ∂I^n, hence induces a homotopy f_t of $f \,|\, X^{n-1} \cup e^n$ fixed on X^{n-1}. If $m > n$, there are no surjective linear maps $\mathbb{R}^n \to \mathbb{R}^m$, so $f_1^{-1}(\Delta^m)$ must be empty, and we can choose p to be any point of Δ^m.

Proof of 4.18: Identifying e^m with $\mathbb{R}^m$, let $B_r \subset e^m$ be the closed ball of radius r centered at the origin. Then for each finite r, $f^{-1}(B_r)$ is compact. Since f is uniformly continuous on $f^{-1}(B_2)$, we may choose $\varepsilon > 0$ so that $|x-y| < \varepsilon$ implies $|f(x)-f(y)| < 1/4$ for $x, y \in f^{-1}(B_2)$. Subdivide the interval I so that the induced subdivision of I^n into cubes has each cube lying in a ball of diameter less than ε. Let K_1 be the union of all the closed cubes meeting $f^{-1}(B_1)$, and let K_2 be the union of all closed cubes meeting K_1, so $f^{-1}(B_1) \subset K_1 \subset K_2 \subset f^{-1}(B_2)$.

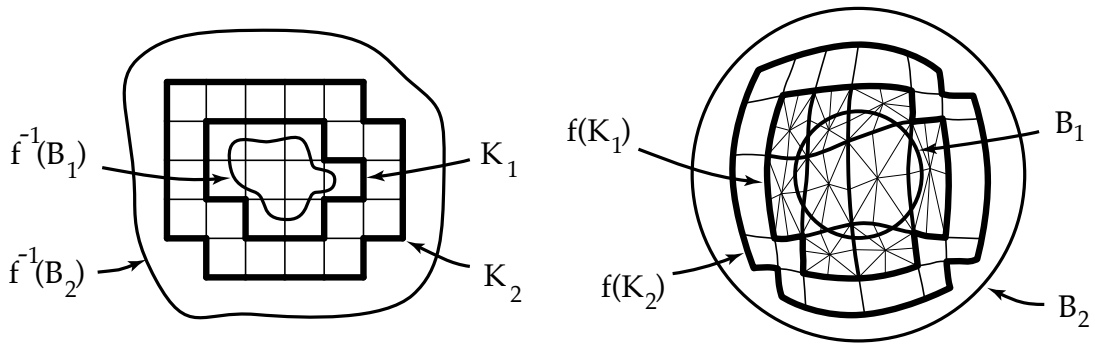

We can view K_2 as a CW complex whose k-cells are k-dimensional open cubes. The barycentric subdivision of this cubical structure is a simplicial complex structure on

K_2 whose vertices are the center points of the cells. Let $g: K_2 \to e^m \approx \mathbb{R}^m$ be the map which equals f on vertices and is linear on simplices. Then $|f(x) - g(x)| < 1/2$ for all $x \in K_2$ since every point of K_2 is within distance ε of a vertex. Let $f_t: K_2 \to e^m$ be the homotopy $(1 - t\varphi)f + (t\varphi)g$ where $\varphi: K_2 \to [0,1]$ is a map which is 0 on the frontier of K_2 and 1 on K_1. Thus $f_0 = f$ and $f_1 \big| K_1 = g \big| K_1$. Since f_t is the constant homotopy on the frontier of K_2, we may extend f_t to be the constant homotopy of f on the rest of I^n.

The image of each point of K_2 moves a distance less than $1/2$ during the homotopy f_t, along the segment from $f(x)$ to $g(x)$. Since f maps the exterior of K_1 outside B_1, it follows that f_1 maps the exterior of K_1 outside $B_{1/2}$. Thus if we choose a small simplex Δ^m in $B_{1/2}$, then for each simplex Δ^n of K_1, f_1 will map the convex polyhedron $\Delta^n \cap f^{-1}(\Delta^m)$ linearly to Δ^m. Some of these linear maps $f_1: \Delta^n \cap f^{-1}(\Delta^m) \to \Delta^m$ may be the restrictions of nonsurjective linear maps $\mathbb{R}^n \to \mathbb{R}^m$, but we can choose Δ^m in the complement of the images of such nonsurjective linear maps. □

Cellular Models

The theme in this subsection is to build CW complexes modelled in one way or another on a given space. Here are some examples of what we shall show:

(1) Given an arbitrary path-connected space Y, there is a connected CW complex X and a map $f: X \to Y$ for which $f_*: \pi_n(X) \to \pi_n(Y)$ is an isomorphism for all n. Such an X is called a **CW approximation** to Y. As we shall see, the map f also induces isomorphisms on homology and cohomology, so CW approximations allow many general statements in algebraic topology to be reduced to the case of CW complexes, where one can often make cell-by-cell arguments.

(2) An n-connected CW complex is homotopy equivalent to a CW complex with n-skeleton a point. More generally, if (X, A) is an n-connected CW pair with $A \neq \varnothing$, then there is a CW pair $(Y, A) \simeq (X, A)$ rel A such that all cells of $Y - A$ have dimension greater than n. Thus the homotopy-theoretic notion of connectivity has a geometric interpretation.

(3) Given a connected CW complex X and an integer $n \geq 1$, there is an n-connected CW complex Z and a map $Z \to X$ inducing an isomorphism on π_i for all $i > n$. If $n = 1$, the universal cover of X has these properties, so Z can be viewed as an ***n*-connected cover** of X.

(4) As a sort of mirror reflection of (3), given a connected CW complex X and an integer $n \geq 1$, there is a CW complex Y with $\pi_i(Y) = 0$ for $i > n$ and a map $X \to Y$

inducing an isomorphism on π_i for $i \le n$. Such a Y is called an **n-approximation** to X. Letting n go from 1 to ∞ one has then a sequence of spaces which form better and better approximations to X, each successive space agreeing with X in one more homotopy group. This is the essential idea of Postnikov towers, which will be studied in §4.3.

We shall do all four of these constructions by a technique which in effect splices together two given spaces X and Y to form a new space Z with some homotopy properties of both X and Y. This is formulated in the following definition. An ***n*-factorization** of a map $f: X \to Y$ of path-connected spaces is a composition $X \to Z_n \to Y$ homotopic to f, such that:

(i) The map $X \to Z_n$ induces an isomorphism on π_i for $i < n$ and a surjection for $i = n$.
(ii) The map $Z_n \to Y$ induces an isomorphism on π_i for $i > n$ and an injection for $i = n$.

Thus $f_*: \pi_n(X) \to \pi_n(Y)$ is factored as a composition $\pi_n(X) \to \pi_n(Z_n) \to \pi_n(Y)$ of a surjection followed by an injection, just as any group homomorphism $\varphi: A \to B$ can be factored as a surjection $f: A \to Im(\varphi)$ followed by an injection $Im(\varphi) \hookrightarrow B$.

One can think of Z_n as a sort of hybrid of X and Y, having the same homotopy groups as X below dimension n and the same homotopy groups as Y above dimension n. As n increases, these hybrids look more and more like X, and less and less like Y.

Proposition 4.19. *Every map $f: X \to Y$ from a connected CW complex X to a path-connected space Y has n-factorizations $X \hookrightarrow Z_n \to Y$ for all n, with Z_n a CW complex obtained from the subcomplex X by attaching cells of dimension greater than n.*

Proof: Suppose we have $X \hookrightarrow Z_n \to Y$ with Z_n obtained from X by attaching cells of dimension greater than n. Then (Z_n, X) is n-connected, so from the long exact sequence of homotopy groups for the pair (Z_n, X) we see that condition (i) will automatically hold. Thus it is condition (ii) that we need to focus on. By taking the mapping cylinder of $Z_n \to Y$ we may convert this map into an inclusion, and then condition (ii) is equivalent to $\pi_i(Y, Z_n)$ being 0 for $i > n$. Our strategy will be to construct Z_n as a union of subspaces $X = X_n \subset X_{n+1} \subset \cdots$ where X_k is obtained from X_{k-1} by attaching k-cells and $\pi_i(Y, X_k) = 0$ for $n < i \le k$.

Suppose inductively that we have already constructed X_{k-1} with a map $X_{k-1} \to Y$ extending the given f. Replacing Y by the mapping cylinder, we may assume this map is an inclusion $X_{k-1} \hookrightarrow Y$. Choose maps $\varphi_\alpha: (D^k, S^{k-1}) \to (Y, X_{k-1})$ representing some set of generators for $\pi_k(Y, X_{k-1})$. We may assume the maps φ_α have been homotoped to be cellular on S^{k-1}, and then we use the restrictions $\varphi_\alpha \big| S^{k-1}$ to attach cells e^k_α to X_{k-1}, forming X_k. The maps φ_α themselves give an extension of the inclusion $X_{k-1} \hookrightarrow Y$ to

a map $X_k \to Y$, which once again we turn into an inclusion via a mapping cylinder. The long exact sequence of the triple (Y, X_k, X_{k-1}) contains the segment

$$\pi_k(X_k, X_{k-1}) \longrightarrow \pi_k(Y, X_{k-1}) \longrightarrow \pi_k(Y, X_k) \longrightarrow 0$$

where the final 0 is $\pi_{k-1}(X_k, X_{k-1})$, which vanishes by cellular approximation. The first map in the segment is surjective since characteristic maps for the cells e^k_α represent elements of $\pi_k(X_k, X_{k-1})$ mapping onto the chosen generators for $\pi_k(Y, X_{k-1})$. Exactness then implies that $\pi_k(Y, X_k) = 0$. Since $\pi_i(X_k, X_{k-1}) = 0$ for $i < k$, the subsequent terms of the exact sequence gives isomorphisms $\pi_i(Y, X_k) \approx \pi_i(Y, X_{k-1})$ for $i < k$, completing the induction step.

We let $Z_n = \bigcup_k X_k$, with the map $Z_n \to Y$ given by the inductively constructed extensions $X_k \to Y$ of f. We now have an n-factorization of f since the maps $\pi_i(X) \to \pi_i(Z_n) \to \pi_i(Y)$ depend only on the $(i+1)$-skeleton of Z_n, by cellular approximation. □

The applications (1)-(4) listed earlier can now be derived as special cases:

(1) If we take X to be a point and Y to be any path-connected space, then a 0-factorization of any map $X \to Y$ is a CW approximation to Y. Note that the construction of Z_0 gives in this case a CW complex with a single 0-cell, and all the higher cells are attached by basepoint-preserving maps. In particular, any connected CW complex is homotopy equivalent to a CW complex with these properties.

One can also obtain a CW approximation to an arbitrary pair (Y, B) of path-connected spaces by first constructing a CW approximation $X \to B$, then forming a 0-factorization of the composition $X \to B \hookrightarrow Y$. The result is a map $(Z, X) \to (Y, B)$ with (Z, X) a CW pair, such that $Z \to Y$ and its restriction $X \to B$ are CW approximations to Y and B. It follows from the five-lemma that the map $(Z, X) \to (Y, B)$ induces isomorphisms on relative as well as absolute homotopy groups.

The restriction to path-connected spaces can be removed by forming CW approximations to each path-component separately.

(2) Given an n-connected CW pair (X, A), form an n-factorization $A \hookrightarrow Z \xrightarrow{f} X$ of the inclusion $A \hookrightarrow X$, where Z is obtained from A by attaching cells of dimension greater than n. It follows that f induces isomorphisms on π_i for all i. Namely, this is true for $i > n$ by the definition of an n-factorization, and for $i < n$ both inclusions $A \hookrightarrow Z$ and $A \hookrightarrow X$ induce isomorphisms on π_i, hence also f. In the remaining case $i = n$, f induces an injection and the inclusion $A \hookrightarrow X$ induces a surjection, hence f does also. Since f induces isomorphisms on all homotopy groups, it is a homotopy equivalence. To

make this a homotopy equivalence rel A we can appeal to Proposition 0.4, or we can give a direct argument as follows. Form a quotient space W of M_f by collapsing each segment $\{a\} \times I$ to a point, for $a \in A$. Assuming f has been made cellular, W is a CW complex containing X and Z as subcomplexes, and W deformation retracts to X just as M_f does. Also, $\pi_i(W, Z) = 0$ for all i since f induces isomorphisms on all homotopy groups, so W deformation retracts onto Z. These two deformation retractions of W onto X and Z are stationary on A, hence give a homotopy equivalence $X \simeq Z$ rel A.

(3) This is easy: Just take an n-factorization of the inclusion of a point into X.

(4) Equally simply, take an $(n-1)$-factorization of a map of X to a point.

A map $f : X \to Y$ is called a **weak homotopy equivalence** if it induces isomorphisms $\pi_n(X, x_0) \to \pi_n\big(Y, f(x_0)\big)$ for all n and all choices of basepoint x_0. Whitehead's theorem can be restated as saying that a weak homotopy equivalence between CW complexes is a homotopy equivalence. But in general, weak homotopy equivalence is strictly weaker than homotopy equivalence. For example, an exercise at the end of the previous section is to show that the 'quasi-circle' is a noncontractible space whose homotopy groups are all trivial.

Often a weak homotopy equivalence is as good as a homotopy equivalence. For example:

Proposition 4.20. *A weak homotopy equivalence $f : X \to Y$ induces isomorphisms $f_* : H_n(X; G) \to H_n(Y; G)$ and $f^* : H^n(Y; G) \to H^n(X; G)$ for all n and all coefficient groups G.*

Proof: By the universal coefficient theorems it suffices to prove the result for homology with $\mathbb{Z}$ coefficients. Furthermore, replacing Y by the mapping cylinder M_f and looking at the long exact sequences of homology and homotopy groups for (M_f, X), we see that it suffices to show:

(∗) If a pair of path-connected spaces (Z, X) has $\pi_i(Z, X) = 0$ for $i \le n$, then $H_i(Z, X; \mathbb{Z}) = 0$ for $i \le n$.

Recall from §2.1 the singular complex $S(Z)$, a Δ-complex whose n-simplices are the maps $\Delta^n \to Z$. For the subspace $X \subset Z$, $S(X)$ is a subcomplex of $S(Z)$. The natural evaluation map $e : \big(S(Z), S(X)\big) \to (Z, X)$ induces an isomorphism on relative homology groups since the simplicial homology of $\big(S(Z), S(X)\big)$ is identical with the singular homology of (Z, X) and the simplicial homology of a Δ-complex pair is equivalent to its singular homology by Theorem 2.20. If $\pi_i(Z, X) = 0$ for $i \le n$, Lemma 4.5 implies that

the map e is homotopic, staying fixed on $S(X)$, to a map sending the n-skeleton of $S(Z)$ to X. Hence the induced map $e_*: H_i(S(Z), S(X); \mathbb{Z}) \to H_i(Z, X; \mathbb{Z})$ is zero for $i \le n$. But since e_* is an isomorphism, this implies $H_i(Z, X; \mathbb{Z}) = 0$ for $i \le n$. □

CW approximations can be used to reduce many statements about general spaces to the special case of CW complexes. For example, the cup product version of the Künneth formula in Theorem 3.9, asserting that $H^*(X \times Y; R) \approx H^*(X; R) \otimes H^*(Y; R)$ under certain conditions, can now be extended to non-CW spaces since if X and Y are CW approximations to spaces Z and W, respectively, then $X \times Y$, with the CW topology, is a CW approximation to $Z \times W$. (Changing the product topology to the CW topology on $X \times Y$ has no effect on homotopy groups since these two topologies have the same compact sets.) Similarly, the general Künneth formula for homology in §3.4b holds for arbitrary products $X \times Y$.

The condition for a map $Y \to Z$ to be a weak homotopy equivalence involves only maps of spheres into Y and Z, but since spheres are the building blocks of CW complexes it should not be surprising that weak homotopy equivalences $Y \to Z$ behave nicely with respect to maps of arbitrary CW complexes into Y and Z, not just spheres. The following proposition gives a precise statement, using the notations $[X, Y]$ for the set of homotopy classes of maps $X \to Y$ and $\langle X, Y \rangle$ for the set of basepoint-preserving-homotopy classes of basepoint-preserving maps $X \to Y$. (The notation $\langle X, Y \rangle$ is not standard, but is intended to suggest 'pointed homotopy classes.')

Proposition 4.21. *A weak homotopy equivalence $f: Y \to Z$ induces bijections $[X, Y] \to [X, Z]$ and $\langle X, Y \rangle \to \langle X, Z \rangle$ for all CW complexes X.*

Proof: Consider first $[X, Y] \to [X, Z]$. We may assume f is an inclusion by replacing Z by the mapping cylinder M_f as usual. The groups $\pi_n(Z, Y, y_0)$ are then zero for all n and all basepoints $y_0 \in Y$, so Lemma 4.5 implies that any map $X \to Z$ can be homotoped to have image in Y. This gives surjectivity of $[X, Y] \to [X, Z]$. A relative version of this argument shows injectivity since we can deform a homotopy $(X \times I, X \times \partial I) \to (Z, Y)$ to have image in Y.

In the case of $\langle X, Y \rangle \to \langle X, Z \rangle$ the same argument applies if M_f is replaced by the reduced mapping cylinder, which, for a basepoint $y_0 \in Y$, is the quotient of M_f obtained by collapsing the segment $\{y_0\} \times I$ to a point. This point then serves as the common basepoint of Y, Z, and the reduced mapping cylinder. The reduced mapping cylinder deformation retracts to Z just as the unreduced one does, but with the advantage that the basepoint does not move. □

To finish off our discussion of n-factorizations we show that they behave nicely with respect to maps.

Proposition 4.22. *Suppose we have a diagram*

$$\begin{array}{ccccc} X & \longrightarrow & Z_n & \longrightarrow & Y \\ \downarrow & & & & \downarrow \\ X' & \longrightarrow & Z'_m & \longrightarrow & Y' \end{array}$$

which is commutative up to homotopy, with the first row an n-factorization and the second row an m-factorization. Then if $m \leq n$ and Z_n has the homotopy type of a CW complex, the diagram can be enlarged by a map $Z_n \to Z'_m$ so as to be still homotopy-commutative, and such a map $Z_n \to Z'_m$ is unique up to homotopy.

Proof: By composing the maps to and from Y and X' we get a diagram

$$\begin{array}{ccc} X & \longrightarrow & Z_n \\ \downarrow & & \downarrow \\ Z'_m & \longrightarrow & Y' \end{array}$$

Replacing the map $X \to Z_n$ by an inclusion into its mapping cylinder and then appying (2) above, we may assume that Z_n is obtained from X by attaching cells of dimension greater than n. We may also assume the map $Z'_m \to Y'$ is an inclusion, via the mapping cylinder. Since $\pi_i(Y', Z_m) = 0$ for $i > m$, the assumption $n \geq m$ implies that we may apply Lemma 4.5 to deform the map $Z_n \to Y'$ to have image in Z'_m, staying fixed on X. This gives the desired map $Z_n \to Z'_m$. In similar fashion, given two such maps $Z_n \to Z'_m$, one can construct a homotopy $Z_n \times I \to Z'_m$ between them which is stationary on X, since $Z_n \times I$ is obtained from $X \times I \cup Z_n \times \partial I$ by attaching cells of dimension greater than $n+1$. □

This proposition implies for example that a map between spaces induces a map between CW approximations to the two spaces, unique up to homotopy. It follows formally that CW approximations are unique up to homotopy equivalence. By similar reasoning, n-approximations and n-connected coverings are unique up to homotopy equivalence. Another important application of the proposition will be to construct Postnikov towers, in §4.3.

There are quite a few situations in algebraic topology where having a nontrivial fundamental group complicates matters considerably. We shall next describe a construction

which in certain circumstances allows one to modify a space so as to eliminate its fundamental group, or at least simplify it, without affecting homology or cohomology. Here is the simplest case:

Proposition 4.23. *Let X be a connected CW complex with $H_1(X) = 0$. Then there is a simply-connected CW complex X^+ and a map $X \to X^+$ inducing isomorphisms on all homology groups.*

Proof: Choose loops $\varphi_\alpha : S^1 \to X^1$ generating $\pi_1(X)$ and use these to attach cells e^2_α to X to form a simply-connected CW complex X'. For the pair (X', X) the homology exact sequence

$$0 \longrightarrow H_2(X) \longrightarrow H_2(X') \longrightarrow H_2(X', X) \longrightarrow 0 = H_1(X)$$

splits since $H_2(X', X)$ is free with basis the cells e^2_α. Thus we have $H_2(X') \approx H_2(X) \oplus H_2(X', X)$. Since X' is simply-connected, the Hurewicz theorem gives an isomorphism $H_2(X') \approx \pi_2(X')$, and so we may represent the generators $[e^2_\alpha]$ of the summand $H_2(X', X)$ by maps $\psi_\alpha : S^2 \to X'$. We may assume these are cellular maps, and then use them to attach cells e^3_α to X' forming a simply-connected CW complex X^+, with the inclusion $X \hookrightarrow X^+$ an isomorphism on all homology groups. □

In the preceding proposition, the condition $H_1(X) = 0$ means that $\pi_1(X)$ is equal to its commutator subgroup, i.e., $\pi_1(X)$ is a **perfect** group. Suppose more generally that X is a connected CW complex and $H \subset \pi_1(X)$ is a perfect subgroup. Let $p : \widetilde{X} \to X$ be the covering space corresponding to H, so $\pi_1(\widetilde{X}) \approx H$ is perfect, hence $H_1(\widetilde{X}) = 0$. From the previous proposition we get an inclusion $\widetilde{X} \hookrightarrow \widetilde{X}^+$. Let X^+ be obtained from the disjoint union of $\widetilde{X}^+$ and the mapping cylinder M_p by identifying the copies of $\widetilde{X}$ in these two spaces. Thus we have a commutative diagram of inclusion maps

$$\begin{array}{ccc} \widetilde{X} & \longrightarrow & \widetilde{X}^+ \\ \downarrow & & \downarrow \\ X \simeq M_p & \longrightarrow & X^+ \end{array}$$

From the van Kampen theorem, the induced map $\pi_1(X) \to \pi_1(X^+)$ is surjective with kernel the normal subgroup generated by H. Further, since X^+/M_p is homeomorphic to $\widetilde{X}^+/\widetilde{X}$ we have $H_*(X^+, M_p) = H_*(\widetilde{X}^+, \widetilde{X}) = 0$, so the map $X \to X^+$ induces an isomorphism on homology.

This construction $X \to X^+$, killing a perfect subgroup of $\pi_1(X)$ while preserving homology, is known as the Quillen **plus construction**. In some of the main applications X is a $K(G, 1)$ where G has perfect commutator subgroup, so the map $X \to X^+$ abelianizes

π_1 while preserving homology. The space X^+ need no longer be a $K(\pi,1)$, and in fact its homotopy groups can be quite interesting. The most striking example is $G = \Sigma_\infty$, the infinite symmetric group consisting of permutations of $1, 2, \cdots$ fixing all but finitely many n's, with commutator subgroup the infinite alternating group A_∞ which is perfect. In this case a famous theorem of Barrett-Kahn-Priddy says that the homotopy groups $\pi_i(K(\Sigma_\infty, 1)^+)$ are the stable homotopy groups of spheres!

There are limits, however, on which subgroups of $\pi_1(X)$ can be killed without affecting the homology of X. For example, for $X = S^1 \vee S^1$ it is impossible to kill the commutator subgroup of $\pi_1(X)$ while preserving homology. In fact, by exercise 33 at the end of the previous section, every space with fundamental group $\mathbb{Z} \times \mathbb{Z}$ must have nontrivial H_2.

Excision for Homotopy Groups

What makes homotopy groups so much harder to compute than homology groups is the failure of excision. As a concrete example of this failure, consider S^2 decomposed by S^1 into hemispheres D^2_+ and D^2_-. If excision held we would have $\pi_i(S^2, D^2_-) \approx \pi_i(D^2_+, S^1)$, and from the long exact sequences for these pairs we would then deduce isomorphisms $\pi_i(S^2) \approx \pi_{i-1}(S^1)$ for all i. In particular, for $i = 3$ we would have $\pi_3(S^2) \approx \pi_2(S^1)$. The latter group is 0 by Proposition 4.1, but $\pi_3(S^2) \approx \mathbb{Z}$, as we will show in §4.3.

However, there is a certain dimension range, depending on connectivities, in which excision does hold for homotopy groups:

Theorem 4.24. *Let X be a CW complex which is the union of subcomplexes A and B whose intersection $C = A \cap B$ is path-connected. If (A, C) is m-connected and (B, C) is n-connected, $m, n \geq 1$, then the map $\pi_i(A, C) \to \pi_i(X, B)$ induced by inclusion is an isomorphism for $i < m + n$ and a surjection for $i = m + n$.*

This yields the Freudenthal suspension theorem:

Corollary 4.25. *The suspension map $\pi_i(S^n) \to \pi_{i+1}(S^{n+1})$ is an isomorphism for $i < 2n - 1$ and a surjection for $i = 2n - 1$.*

Proof: Decompose S^{n+1} as the union of two hemispheres D^{n+1}_+ and D^{n+1}_- intersecting in S^n. The suspension map is the same as the map

$$\pi_i(S^n) \approx \pi_{i+1}(D^{n+1}_+, S^n) \longrightarrow \pi_{i+1}(S^{n+1}, D^{n+1}_-) \approx \pi_{i+1}(S^{n+1})$$

where the two isomorphisms come from long exact sequences of pairs and the middle map is induced by inclusion. Since (D^{n+1}, S^n) is n-connected, the theorem says the middle map is an isomorphism for $i + 1 < 2n$ and surjective for $i + 1 = 2n$. □

Corollary 4.26. *$\pi_n(S^n) \approx \mathbb{Z}$, generated by the identity map, for all $n \geq 1$. In particular, the degree map $\pi_n(S^n) \to \mathbb{Z}$ is an isomorphism.*

Proof: We know that $\pi_1(S^1) \approx \mathbb{Z}$ from §1.1, so the preceding corollary implies that $\pi_n(S^n)$ is a cyclic group independent of n for $n \geq 2$. This cyclic group is infinite since there exist basepoint-preserving maps $S^n \to S^n$ of arbitrary degree if $n \geq 1$, and degree is a homotopy invariant. The degree map $\pi_n(S^n) \to \mathbb{Z}$ is an isomorphism since the the map $z \mapsto z^k$ of S^1 has degree k, as do its iterated suspensions.

Alternatively, if one wants a purely homotopy-theoretic proof that $\pi_n(S^n) \approx \mathbb{Z}$, one can use the Hopf bundle $S^1 \to S^3 \to S^2$ described in §4.3, whose long exact sequence of homotopy groups gives an isomorphism $\pi_1(S^1) \approx \pi_2(S^2)$. □

Proof of 4.24: We proceed by proving successively more general cases. The first case contains the heart of the argument, and suffices for the calculation of $\pi_n(S^n)$.

Case 1: A is obtained from C by attaching cells e^{m+1}_α and B is obtained from C by attaching a cell e^{n+1}. To show surjectivity of $\pi_i(A, C) \to \pi_i(X, B)$, let $f : (I^i, \partial I^i, J^{i-1}) \to (X, B, x_0)$ be given. The image of f is compact, meeting only finitely many of these cells e^{m+1}_α and e^{n+1}, so by repeated applications of Lemma 4.18 we may homotope f, through maps $(I^i, \partial I^i, J^{i-1}) \to (X, B, x_0)$, so that the preimages $f^{-1}(\Delta^{m+1}_\alpha)$ and $f^{-1}(\Delta^{n+1})$ of simplices in e^{m+1}_α and e^{n+1} are finite unions of convex polyhedra, on each of which f is the restriction of a linear surjection.

Claim: If $i \leq m + n$, then there exist points $p_\alpha \in \Delta^{m+1}_\alpha$, $q \in \Delta^{n+1}$, and a map $\varphi : I^{i-1} \to [0, 1)$ such that:

(a) $f^{-1}(q)$ lies below the graph of φ in $I^{i-1} \times I = I^i$.
(b) $f^{-1}(p_\alpha)$ lies above the graph of φ for each α.
(c) $\varphi = 0$ on ∂I^{i-1}.

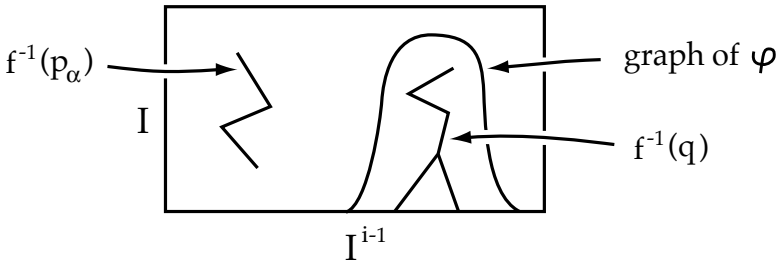

Granting this, let f_t be a homotopy of f excising the region under the graph of φ by restricting f to the region above the graph of $t\varphi$ for $0 \leq t \leq 1$. By (b), $f_t(I^{i-1})$ is disjoint from p_α for all t and α, and by (a), $f_1(I^i)$ is disjoint from q. So, replacing $\pi_i(A, C) \to$

$\pi_i(X,B)$ by the equivalent map $\pi_i(X-q, X-q-\bigcup_\alpha\{p_\alpha\}) \to \pi_i(X, X-\bigcup_\alpha\{p_\alpha\})$, we see that $[f]=[f_1]$ lies in the image of this map. This proves the surjectivity statement.

Now we prove the Claim. For any $q \in \Delta^{n+1}$, $f^{-1}(q)$ is a finite union of convex polyhedra of dimension $\le i-n-1$ since $f^{-1}(\Delta^{n+1})$ is a finite union of convex polyhedra on each of which f is the restriction of a linear surjection $\mathbb{R}^i \to \mathbb{R}^{n+1}$. We wish to choose the points $p_\alpha \in \Delta^{m+1}_\alpha$ so that not only is $f^{-1}(q)$ disjoint from $f^{-1}(p_\alpha)$ for each α, but also so that $f^{-1}(q)$ and $f^{-1}(p_\alpha)$ have disjoint images under the projection $\pi: I^i \to I^{i-1}$. This is equivalent to saying that $f^{-1}(p_\alpha)$ is disjoint from $T = \pi^{-1}\big(\pi(f^{-1}(q))\big)$, the union of all segments $\{x\}\times I$ meeting $f^{-1}(q)$. This set T is a finite union of convex polyhedra of dimension $\le i-n$ since $f^{-1}(q)$ is a finite union of convex polyhedra of one lower dimension. Since linear maps cannot increase dimension, $f(T)\cap\Delta^{m+1}_\alpha$ is also a finite union of convex polyhedra of dimension $\le i-n$. Thus if $m+1 > i-n$, there is a point $p_\alpha \in \Delta^{m+1}_\alpha$ not in $f(T)$. This gives $f^{-1}(p_\alpha)\cap T = \varnothing$ if $i \le m+n$. Hence we can choose a neighborhood U of $\pi\big(f^{-1}(q)\big)$ in I^{n-1} disjoint from $\pi\big(f^{-1}(p_\alpha)\big)$ for all α. Then there exists $\varphi: I^{i-1} \to [0,1)$ having support in U, with $f^{-1}(q)$ lying under the graph of φ. This verifies the Claim, and so finishes the proof of surjectivity in Case 1.

For injectivity in Case 1 the argument is the same, with $f:(I^i,\partial I^i, J^{i-1})\times[0,1] \to (X,B,x_0)$ now a nullhomotopy of a map $(I^i,\partial I^i,J^{i-1}) \to (A,C,x_0)$. One constructs a function $\varphi: I^{i-1}\times I \to [0,1)$ separating $f^{-1}(q)$ from sets $f^{-1}(p_\alpha)$ as before, and this allows one to excise $f^{-1}(q)$ from the domain of f. Since $I^i\times I$ now plays the role of I^i, the dimension i is replaced by $i+1$ and the dimension restriction $i \le m+n$ becomes $i < m+n$.

Case 2: A is obtained from C by attaching $(m+1)$-cells as in Case 1, and B is obtained from C by attaching cells of dimension $\ge n+1$. To show surjectivity of $\pi_i(A,C) \to \pi_i(X,B)$, consider an element of $\pi_i(X,B)$ represented by a map $f:(I^i,\partial I^i,J^{i-1}) \to (X,B,x_0)$. The image of f is compact, meeting only finitely many cells, and by repeated applications of Case 1 we can push f off the cells of $B-C$ one at a time, in order of decreasing dimension. Injectivity is similar, taking f to be a nullhomotopy $(I^i,\partial I^i,J^{i-1})\times[0,1] \to (X,B,x_0)$ of an element of the kernel of $\pi_i(A,C) \to \pi_i(X,B)$.

Case 3: A is obtained from C by attaching cells of dimension $\ge m+1$ and B is as in Case 2. We may assume all cells of $A-C$ have dimension $\le m+n+1$ since higher-dimensional cells have no influence on π_i for $i \le m+n$, by cellular approximation. Let $A_k \subset A$ be the union of C with the cells of A of dimension $\le k$ and let $X_k = A_k \cup B$. We prove the result for $\pi_i(A_k,C) \to \pi_i(X_k,B)$ by induction on k. The induction starts with $k=m+1$ which is Case 2. For the induction step consider the commutative diagram formed by the

exact sequences of the triples (A_k, A_{k-1}, C) and (X_k, X_{k-1}, B):

$$\begin{array}{ccccccccc}
\pi_{i+1}(A_k, A_{k-1}) & \longrightarrow & \pi_i(A_{k-1}, C) & \longrightarrow & \pi_i(A_k, C) & \longrightarrow & \pi_i(A_k, A_{k-1}) & \longrightarrow & \pi_{i-1}(A_{k-1}, C) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
\pi_{i+1}(X_k, X_{k-1}) & \longrightarrow & \pi_i(X_{k-1}, B) & \longrightarrow & \pi_i(X_k, B) & \longrightarrow & \pi_i(X_k, X_{k-1}) & \longrightarrow & \pi_{i-1}(X_{k-1}, B)
\end{array}$$

When $i < m + n$ the first and fourth vertical maps are isomorphisms by Case 2, while by induction the second and fifth maps are isomorphisms, so the middle map is an isomorphism by the five-lemma. Similarly, when $i = m + n$ the second and fourth maps are surjective and the fifth map is injective, which is enough to imply the middle map is surjective by one half of the five-lemma.

Case 4: The general case. The connectivity assumptions on the pairs (A, C) and (B, C) imply that they are homotopy equivalent to pairs (A', C) and (B', C) as in Case 3, via homotopy equivalences fixed on C, so these homotopy equivalences fit together to give a homotopy equivalence $A \cup B \simeq A' \cup B'$. Thus the general case reduces to Case 3. □

Excision for homology implies that the relative homology groups of a CW pair (X, A) are equivalent to the absolute homology groups of X/A. For homotopy groups this is also true in a certain dimension range:

Proposition 4.27. *If a CW pair (X, A) is r-connected, $r \geq 1$, and A is s-connected, $s \geq 0$, then the map $\pi_i(X, A) \to \pi_i(X/A)$ induced by the quotient map $X \to X/A$ is an isomorphism for $i \leq r + s$ and a surjection for $i = r + s + 1$.*

Proof: Let CA be the cone $(A \times I)/(A \times \{1\})$. Identifying $A \subset X$ with $A = A \times \{0\} \subset CA$ produces the space $X \cup CA$. Since CA is a contractible subcomplex of $X \cup CA$, the quotient map $X \cup CA \to (X \cup CA)/CA = X/A$ is a homotopy equivalence. So we have a diagram

$$\begin{array}{ccccc}
\pi_i(X, A) & \longrightarrow & \pi_i(X \cup CA, CA) & \longrightarrow & \pi_i(X \cup CA/CA) = \pi_i(X/A) \\
 & & \uparrow \approx & \nearrow \approx & \\
 & & \pi_i(X \cup CA) & &
\end{array}$$

where the vertical isomorphism comes from a long exact sequence. Now apply the theorem to the first map in the diagram, using the fact that (CA, A) is $(s + 1)$-connected if A is s-connected, which comes from the exact sequence for the pair (CA, A). □

Our next goal is to show how the preceding results can be used to give a different derivation of the Hurewicz theorem. We begin with two special cases:

Example 1. We show that for $n \geq 2$, $\pi_n(\bigvee_\alpha S^n_\alpha)$ is free abelian with basis the homotopy classes of the inclusions $S^n_\alpha \hookrightarrow \bigvee_\alpha S^n_\alpha$. Suppose first that there are only finitely many summands S^n_α. Then we can regard $\bigvee_\alpha S^n_\alpha$ as the n-skeleton of the product $\prod_\alpha S^n_\alpha$, where S^n_α is given its usual CW structure and $\prod_\alpha S^n_\alpha$ has the product CW structure. The pair $(\prod_\alpha S^n_\alpha, \bigvee_\alpha S^n_\alpha)$ is $(2n-1)$-connected since $\prod_\alpha S^n_\alpha$ has cells only in dimensions a multiple of n. Hence from the long exact sequence of this pair we see that the inclusion $\bigvee_\alpha S^n_\alpha \hookrightarrow \prod_\alpha S^n_\alpha$ induces an isomorphism on π_n if $n \geq 2$. By Proposition 4.2 we have $\pi_n(\prod_\alpha S^n_\alpha) \approx \bigoplus_\alpha \pi_n(S^n_\alpha)$, a free abelian group with basis the inclusions $S^n_\alpha \hookrightarrow \prod_\alpha S^n_\alpha$, so this takes care of the case of finitely many S^n_α's.

The case of infinitely many summands S^n_α reduces to the finite case by the following argument. The inclusions $S^n_\alpha \hookrightarrow \bigvee_\alpha S^n_\alpha$ induce a homomorphism $\Phi : \bigoplus_\alpha \pi_n(S^n_\alpha) \to \pi_n(\bigvee_\alpha S^n_\alpha)$. This is surjective since any map $f : S^n \to \bigvee_\alpha S^n_\alpha$ has compact image contained in the wedge sum of finitely many S^n_α's, so by the finite case already proved, $[f]$ is in the image of Φ. Similarly, a nullhomotopy of f has compact image contained in a finite wedge sum of S^n_α's, so the special case also implies that Φ is injective.

Example 2. Suppose X is obtained from a wedge of spheres $\bigvee_\alpha S^n_\alpha$ by attaching cells e^{n+1}_β via basepoint-preserving maps $\varphi_\beta : S^n \to \bigvee_\alpha S^n_\alpha$, with $n \geq 2$. By cellular approximation we know that $\pi_i(X) = 0$ for $i < n$, and we shall show that $\pi_n(X)$ is the quotient of the free abelian group $\pi_n\big(\bigvee_\alpha S^n_\alpha\big) \approx \bigoplus_\alpha \mathbb{Z}$ by the subgroup generated by the classes $[\varphi_\beta]$. Any subgroup can be realized in this way, by choosing maps φ_β to represent a set of generators for the subgroup, so it follows that every abelian group can be realized as $\pi_n(X)$ for such a space $X = \big(\bigvee_\alpha S^n_\alpha\big) \bigcup_\beta e^{n+1}_\beta$. This is the higher-dimensional analog of the construction in Chapter 1 of a 2-dimensional CW complex with prescribed fundamental group.

To see that $\pi_n(X)$ is as claimed, consider the following portion of the long exact sequence of the pair (X, X^n):

$$\pi_{n+1}(X, X^n) \xrightarrow{\partial} \pi_n(X^n) \longrightarrow \pi_n(X) \longrightarrow 0$$

The quotient X/X^n is a wedge of spheres S^{n+1}_β, so the preceding proposition implies that $\pi_{n+1}(X, X^n)$ is free with basis the characteristic maps of the cells e^{n+1}_β. The boundary map ∂ takes these to the classes $[\varphi_\beta]$, and the result follows.

Using these calculations we can easily prove the most often used cases of the Hurewicz theorem:

Theorem 4.28. *If a space X is $(n-1)$-connected, $n \geq 2$, then $\widetilde{H}_i(X) = 0$ for $i < n$ and $\pi_n(X) \approx H_n(X)$. If a pair (X, A) is $(n-1)$-connected, $n \geq 2$, with A simply-connected and nonempty, then $H_i(X, A) = 0$ for $i < n$ and $\pi_n(X, A) \approx H_n(X, A)$.*

Proof: Taking CW approximations to X and to (X, A) reduces us to the cases that X is a CW complex and (X, A) is a CW pair. The relative case then reduces to the absolute case since $\pi_i(X, A) \approx \pi_i(X/A)$ for $i \leq n$ by Proposition 4.27, while $H_i(X, A) \approx \widetilde{H}_i(X/A)$ for all i.

In the absolute case we can replace X by a homotopy equivalent CW complex with $(n-1)$-skeleton a point, hence $\widetilde{H}_i(X) = 0$ for $i < n$. To show that $\pi_n(X) \approx H_n(X)$, we can further simplify by throwing away cells of dimension greater than $n+1$ since these have no effect on π_n or H_n. Thus X has the form $\big(\bigvee_\alpha S^n_\alpha\big) \bigcup_\beta e^{n+1}_\beta$. We may assume the attaching maps φ_β of the cells e^{n+1}_β are cellular since this is what the proof of Proposition 4.19 gives. Example 2 above then applies to compute $\pi_n(X)$ as the quotient of $\bigoplus_\alpha \mathbb{Z}$ by the subgroup generated by the $[\varphi_\beta]$'s. If one computes $\widetilde{H}_n(X)$ via cellular homology one gets the same calculation identically. This is because the cellular boundary map $H_{n+1}(X^{n+1}, X^n) \to H_n(X^n, X^{n-1})$ is determined by the degrees of the compositions $q_\alpha \varphi_\beta$ where q_α collapses all n-cells except e^n_α to a point, and the isomorphism $\pi_n(S^n) \approx \mathbb{Z}$ in Corollary 4.26 is given by degree. □

To use this sort of argument to get the general version of the relative Hurewicz theorem which allows the subspace A to be nonsimply-connected takes more work, and may or may not be worth the trouble. In any case, here is how it can be done. First a preliminary step:

Lemma 4.29. *If X is a connected CW complex to which cells e^n_α are attached for a fixed $n \geq 2$, forming a CW complex $W = X \bigcup_\alpha e^n_\alpha$, then $\pi_n(W, X)$ is a free $\pi_1(X)$-module with basis the homotopy classes of the characteristic maps Φ_α of the cells e^n_α, provided the map $\pi_1(X) \to \pi_1(W)$ induced by inclusion is an isomorphism. In particular, this is always the case if $n \geq 3$. In the general $n = 2$ case, $\pi_2(W, X)$ is generated by the classes of the characteristic maps of the cells e^2_α together with their images under the action of $\pi_1(X)$.*

If the characteristic maps $\Phi_\alpha : (D^n, \partial D^n) \to (W, X)$ do not take a basepoint $s_0 \in \partial D^n$ to the basepoint $x_0 \in X$, then in order for them to define elements of $\pi_n(W, X, x_0)$ one must choose basepoint-change paths from the points $\Phi_\alpha(s_0)$ to x_0. Different choices of such paths yield elements of $\pi_n(W, X, x_0)$ related by the action of $\pi_1(X, x_0)$, so the basis for $\pi_n(W, X, x_0)$ is well-defined up to multiplication by invertible elements of $\mathbb{Z}[\pi_1(X)]$.

The situation when $n = 2$ and the map $\pi_1(X) \to \pi_1(W)$ is not an isomorphism is more complicated because the relative π_2 can be nonabelian in this case. Whitehead analyzed what happens here and showed that $\pi_2(W, X)$ is a 'free crossed $\pi_1(X)$-module.' See [Hog-Angeloni,Metzler,Sieradski].

Proof: Since $W/X = \bigvee_\alpha S^n_\alpha$, we have $\pi_n(W,X) \approx \pi_n(\bigvee_\alpha S^n_\alpha)$ when X is simply-connected, by Proposition 4.27. The conclusion of the lemma in this case is then immediate from the calculation in Example 1 following Proposition 4.27.

When X is not simply-connected but the inclusion $X \hookrightarrow W$ induces an isomorphism on π_1, then the universal cover of W is obtained from the universal cover of X by attaching n-cells lifting the cells e^n_α. If we choose one such lift $\tilde e^n_\alpha$ of e^n_α, then all the other lifts are the images $\gamma\tilde e^n_\alpha$ of $\tilde e^n_\alpha$ under the deck transformations $\gamma \in \pi_1(X)$. The special case proved in the preceding paragraph shows that the relative π_n for the universal cover is the free abelian group with basis corresponding to the cells $\gamma\tilde e^n_\alpha$. By the relative version of Proposition 4.1, the projection of the universal cover of W onto W induces an isomorphism on relative π_n's, so $\pi_n(W,X)$ is free abelian with basis the classes $[\gamma e^n_\alpha]$ as γ ranges over $\pi_1(X)$, or in other words the free $\pi_1(X)$-module with basis the cells e^n_α.

It remains to consider the $n = 2$ case in general. Since both of the pairs (W, X^1) and $(X^1 \bigcup_\alpha e^2_\alpha, X^1)$ are 1-connected, the homotopy excision theorem implies that the map $\pi_2(X^1 \bigcup_\alpha e^2_\alpha, X^1) \to \pi_2(W,X)$ is surjective, reducing us to the case that X is 1-dimensional. We may also assume the 2-cells e^2_α are attached along loops passing through the basepoint 0-cell x_0, since this can be achieved by homotopy of the attaching maps, which does not affect the homotopy type of the pair (W,X).

In the closure of each 2-cell e^2_α choose an embedded disk D^2_α which contains x_0 but is otherwise contained entirely in the interior of e^2_α. Let $Y = X \bigcup_\alpha D^2_\alpha = X \bigvee_\alpha D^2_\alpha$ and let $Z = W - \bigcup_\alpha int(D^2_\alpha)$, so Y and Z are 2-dimensional CW complexes with a common 1-skeleton $Y^1 = Z^1 = Y \cap Z = X \bigvee_\alpha \partial D^2_\alpha$. The inclusion $(W,X) \hookrightarrow (W,Z)$ is a homotopy equivalence of pairs. Homotopy excision gives a surjection $\pi_2(Y,Y^1) \to \pi_2(W,Z)$. The universal cover $\widetilde Y$ of Y is obtained from the universal cover $\widetilde X$ of X by taking the wedge sum with lifts $\widetilde D^2_{\alpha\beta}$ of the disks D^2_α. Hence we have isomorphisms

$$\begin{aligned}
\pi_2(Y,Y^1) &\approx \pi_2(\widetilde Y, \widetilde Y^1) && \text{where } \widetilde Y^1 \text{ is the 1-skeleton of } \widetilde Y\\
&\approx \pi_2(\textstyle\bigvee_{\alpha\beta} \widetilde D^2_{\alpha\beta}, \bigvee_{\alpha\beta} \partial\widetilde D^2_{\alpha\beta}) && \text{since } \widetilde X \text{ is contractible}\\
&\approx \pi_1(\textstyle\bigvee_{\alpha\beta} \partial\widetilde D^2_{\alpha\beta}) && \text{since } \bigvee_{\alpha\beta} \widetilde D^2_{\alpha\beta} \text{ is contractible}
\end{aligned}$$

This last group is free with basis the loops $\partial\widetilde D^2_{\alpha\beta}$, so the inclusions $\widetilde D^2_{\alpha\beta} \hookrightarrow \bigvee_{\alpha\beta} \widetilde D^2_{\alpha\beta}$ form a basis for $\pi_2(\bigvee_{\alpha\beta} \widetilde D^2_{\alpha\beta}, \bigvee_{\alpha\beta} \partial\widetilde D^2_{\alpha\beta})$. This implies that $\pi_2(Y,Y^1)$ is generated by the inclusions $D^2_\alpha \hookrightarrow Y$ and their images under the action of loops in X. The same is true for $\pi_2(W,Z)$ via the surjection $\pi_2(Y,Y^1) \to \pi_2(W,Z)$. Using the isomorphism $\pi_2(W,Z) \approx \pi_2(W,X)$, we conclude that $\pi_2(W,X)$ is generated by the characteristic maps of the cells e^2_α and their images under the action of $\pi_1(X)$. □

Proof of the Hurewicz Theorem: We may assume (X, A) is a CW pair such that the cells of $X - A$ have dimension $\geq n$, by taking a CW approximation to the given (X, A).

We first prove the result assuming that $\pi_1(A) \to \pi_1(X)$ is an isomorphism. This is always the case if $n \geq 3$, so this case will finish the proof except when $n = 2$. We may assume also that $X = X^{n+1}$ since higher-dimensional cells have no effect on π_n or H_n. Consider the commutative diagram:

$$\begin{array}{ccccccc}
\pi_{n+1}(X, X^n \cup A) & \xrightarrow{\partial} & \pi_n(X^n \cup A, A) & \xrightarrow{i_*} & \pi_n(X, A) & \longrightarrow & 0 \\
\downarrow & & \downarrow & & \downarrow & & \\
\pi'_{n+1}(X, X^n \cup A) & \xrightarrow{\partial'} & \pi'_n(X^n \cup A, A) & \xrightarrow{i'_*} & \pi'_n(X, A) & \longrightarrow & 0 \\
\downarrow h' & & \downarrow h' & & \downarrow h' & & \\
H_{n+1}(X, X^n \cup A) & \longrightarrow & H_n(X^n \cup A, A) & \longrightarrow & H_n(X, A) & \longrightarrow & 0
\end{array}$$

The first and third rows are exact sequences for the triple $(X, X^n \cup A, A)$. The left-hand h' is an isomorphism since by the preceding lemma, $\pi_{n+1}(X, X^n \cup A)$ is a free π_1-module with basis the characteristic maps of the $(n+1)$-cells of $X - A$, so $\pi'_{n+1}(X, X^n \cup A)$ is a free abelian group with the same basis, and $H_{n+1}(X, X^n \cup A)$ is also free with basis the $(n+1)$-cells of $X - A$. Similarly, the lemma implies that the middle h' is an isomorphism since the assumption that $\pi_1(A) \to \pi_1(X)$ is an isomorphism implies that $\pi_1(A) \to \pi_1(X^n \cup A)$ is injective, hence an isomorphism if $n \geq 2$.

A simple diagram-chase now shows that the right-hand h' is an isomorphism. Namely, surjectivity follows since $H_n(X^n \cup A, A) \to H_n(X, A)$ is surjective and the middle h' is an isomorphism. For injectivity, take an element $x \in \pi'_n(X, A)$ with $h'(x) = 0$. The map i'_* is surjective since i_* is, so $x = i'_*(y)$ for some $y \in \pi'_n(X^n \cup A, A)$. Since the first two maps h' are isomorphisms and the bottom row is exact, there is a $z \in \pi'_{n+1}(X, X^n \cup A)$ with $\partial'(z) = y$. Hence $x = 0$ since $i_*\partial = 0$ implies $i'_*\partial' = 0$.

It remains to prove the theorem when $n = 2$ and $\pi_1(A) \to \pi_1(X)$ is not an isomorphism. The proof above will apply once we show that the middle h' in the preceding diagram is an isomorphism. The preceding lemma implies that $\pi'_2(X^2 \cup A, A)$ is generated by characteristic maps of the 2-cells of $X - A$. The images of these generators under h' form a basis for $H_2(X^2 \cup A, A)$. Thus h' is a homomorphism from an abelian group (by Lemma 4.9) to a free abelian group taking a set of generators to a basis, hence h' is an isomorphism. □

Stable Homotopy Groups

The suspension SX of a space X is the quotient of $X \times I$ in which $X \times \{0\}$ and $X \times \{1\}$ are each collapsed to points. A map $f : X \to Y$ suspends to $Sf : SX \to SY$, the quotient

map of $f \times \mathbb{1} : X \times I \to Y \times I$, and a homotopy $f \simeq g$ suspends to a homotopy $Sf \simeq Sg$. However, if one wishes to keep track of basepoints, then there is a slight awkwardness: If x_0 is the basepoint of X, then which point of $\{x_0\} \times I \subset SX$ should be the basepoint of SX? To avoid making any choice it is better to consider the reduced suspension ΣX, which is the quotient of SX in which $\{x_0\} \times I$ is collapsed to a point. Since $\{x_0\} \times I$ is contractible, the quotient map $SX \to \Sigma X$ is a homotopy equivalence, assuming X is a CW complex and x_0 is a 0-cell. So there is no loss of homotopy information in replacing SX by ΣX. Taking the image of $\{x_0\} \times I$ in ΣX as the natural basepoint, then a basepoint-preserving map $f : X \to Y$ suspends to a basepoint-preserving map $\Sigma f : \Sigma X \to \Sigma Y$, and similarly for homotopies.

Another way of looking at the reduced suspension ΣX is as the smash product $X \wedge S^1$, where $X \wedge Y$ is the quotient $X \times Y / X \vee Y$ and $X \vee Y$ is embedded in $X \times Y$ as $X \times \{y_0\} \cup \{x_0\} \times Y$ for y_0 a basepoint of Y. From this viewpoint, $\Sigma f : \Sigma X \to \Sigma Y$ becomes $f \wedge \mathbb{1} : X \wedge S^1 \to Y \wedge S^1$.

Though SX and ΣX are usually only homotopy equivalent, when $X = S^n$ they are both homeomorphic to S^{n+1}. Thus, reduced suspension defines a map $\Sigma : \pi_n(X) \to \pi_{n+1}(\Sigma X)$. This is a homomorphism for $n > 0$ since $\Sigma(f + g) = \Sigma f + \Sigma g$.

Theorem 4.30. *For an r-connected CW complex X, the suspension map $\Sigma : \pi_i(X) \to \pi_{i+1}(\Sigma X)$ is an isomorphism when $i < 2r + 1$ and a surjection when $i = 2r + 1$.*

The case $X = S^n$ is the Freudenthal suspension theorem, proved earlier.

Proof: Let CX denote the reduced cone $X \times I / (X \times \{1\} \cup \{x_0\} \times I)$, so $CX/X = \Sigma X$. Reduced cones, like ordinary cones, are contractible by flowing along the lines $\{x\} \times I$. From the long exact sequence for the pair (CX, X) we have $\pi_i(X) \approx \pi_{i+1}(CX, X)$ for all i. Under this isomorphism, $f : S^i \to X$ corresponds to $Cf : D^{i+1} = CS^i \to CX$, the quotient of $f \times \mathbb{1} : S^i \times I \to X \times I$; this is immediate from the definition of the boundary map in the long exact sequence of homotopy groups. Applying Proposition 4.27 to (CX, X) then gives the result since the isomorphisms $\pi_{i+1}(CX, X) \approx \pi_i(X)$ imply that (CX, X) is $(r + 1)$-connected if X is r-connected. □

If X is an r-connected CW complex then ΣX is $(r + 1)$-connected, since we may assume X has r-skeleton a point and then ΣX has $(r + 1)$-skeleton a point. Since suspension increases connectivity, it follows that if X is any CW complex, then iterating the suspension map

$$\pi_i(X) \longrightarrow \pi_{i+1}(\Sigma X) \longrightarrow \pi_{i+2}(\Sigma^2 X) \longrightarrow \cdots$$

eventually produces isomorphisms. The resulting **stable homotopy group** is denoted $\pi_i^s(X)$.

Stable homotopy groups are nicer than ordinary homotopy groups in that they form a reduced homology theory. Namely, for a CW pair (X, A) there is a long exact sequence

$$\cdots \longrightarrow \pi_i^s(A) \longrightarrow \pi_i^s(X) \longrightarrow \pi_i^s(X/A) \longrightarrow \pi_{i-1}^s(A) \longrightarrow \cdots$$

This is obtained by repeated suspension of the usual long exact sequence involving relative groups $\pi_i(X, A)$, since Proposition 4.27 implies that $\pi_{i+n}(\Sigma^n X, \Sigma^n A) \approx \pi_{i+n}(\Sigma^n X/\Sigma^n A)$ for n sufficiently large, and $\Sigma^n X/\Sigma^n A = \Sigma^n(X/A)$.

The difficulty with this homology theory, unfortunately, is that its coefficient groups $\pi_i^s(S^0)$ are so hard to compute. The group $\pi_i^s(S^0)$, which equals $\pi_{i+n}(S^n)$ for $n > i+1$, is often abbreviated to π_i^s and called the **stable i-stem**. As we will see in Volume II, π_i^s is always finite for $i > 0$. These stable homotopy groups of spheres are among the most fundamental objects in topology, and much effort has gone into their calculation. At the present time, complete calculations are known only for i up to around 60 or so. Here is a table for $i \le 19$, taken from [Toda62]:

i	0	1	2	3	4	5	6	7	8	9	10	11	12
π_i^s	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	0	0	$\mathbb{Z}_2$	$\mathbb{Z}_{240}$	$\mathbb{Z}_2\times\mathbb{Z}_2$	$\mathbb{Z}_2\times\mathbb{Z}_2\times\mathbb{Z}_2$	$\mathbb{Z}_6$	$\mathbb{Z}_{504}$	0

13	14	15	16	17	18	19
$\mathbb{Z}_3$	$\mathbb{Z}_2\times\mathbb{Z}_2$	$\mathbb{Z}_{480}\times\mathbb{Z}_2$	$\mathbb{Z}_2\times\mathbb{Z}_2$	$\mathbb{Z}_2\times\mathbb{Z}_2\times\mathbb{Z}_2\times\mathbb{Z}_2$	$\mathbb{Z}_8\times\mathbb{Z}_2$	$\mathbb{Z}_{264}\times\mathbb{Z}_2$

Finding patterns here is no easy matter!

Besides the additive structure of the groups π_i^s, there are products $\pi_i^s \times \pi_j^s \to \pi_{i+j}^s$ defined by compositions $S^{i+j+k} \to S^{j+k} \to S^k$. These are formally somewhat similar to cup products:

Proposition 4.31. *The composition product induces a graded ring structure on $\pi_*^s = \bigoplus_i \pi_i^s$ satisfying the commutativity relation $\alpha\beta = (-1)^{ij}\beta\alpha$ for $\alpha \in \pi_i^s$ and $\beta \in \pi_j^s$.*

Proof: Only distributivity and commutativity need to be checked. One distributivity law is easy: Given $f, g : S^{i+j+k} \to S^{j+k}$ and $h : S^{j+k} \to S^k$, then $h(f+g) = hf+hg$ since both expressions equal hf and hg on the two hemispheres of S^{i+j+k}. The other distributivity law will follow from this one and the commutativity relation. For commutativity, let $f : S^{i+k} \to S^k$ and $g : S^{j+k} \to S^k$ be given. We may assume k is even. Consider the commutative diagram

$$\begin{array}{ccccc} S^{i+k} \wedge S^{j+k} & \xrightarrow{f\wedge 1\!\!1} & S^k \wedge S^{j+k} & \xrightarrow{1\!\!1\wedge g} & S^k \wedge S^k \\ & & \downarrow\sigma & & \downarrow\tau \\ & & S^{j+k} \wedge S^k & \xrightarrow{g\wedge 1\!\!1} & S^k \wedge S^k \end{array}$$

where σ and τ transpose the two factors. Thinking of S^{j+k} and S^k as smash products of circles, σ is the composition of $k(j+k)$ transpositions of adjacent circle factors. Such a transposition has degree -1 since it is realized as a reflection of the $S^2 = S^1 \wedge S^1$ involved. Hence σ has degree $(-1)^{k(j+k)}$, which is $+1$ since k is even. Thus σ is homotopic to the identity. Similarly, τ is homotopic to the identity. Hence $f \wedge g = (\mathbb{1} \wedge g)(f \wedge \mathbb{1})$ is homotopic to the composition $(g \wedge \mathbb{1})(f \wedge \mathbb{1})$, which is stably equivalent to the composition gf. Symmetrically, fg is stably homotopic to $g \wedge f$. So it suffices to show $f \wedge g \simeq (-1)^{ij} g \wedge f$. This we do by the commutative diagram

$$\begin{array}{ccc} S^{i+k} \wedge S^{j+k} & \xrightarrow{f \wedge g} & S^k \wedge S^k \\ \downarrow \sigma & & \downarrow \tau \\ S^{j+k} \wedge S^{i+k} & \xrightarrow{g \wedge f} & S^k \wedge S^k \end{array}$$

where σ and τ are again the transpositions of the two factors. As before, τ is homotopic to the identity, but now σ has degree $(-1)^{(i+k)(j+k)}$, which equals $(-1)^{ij}$ since k is even. The composition $(g \wedge f)\sigma$ is homotopic to $(-1)^{ij}(g \wedge f)$ since additive inverses in homotopy groups are obtained by pre-composing with a reflection, and a reflection has degree -1. Thus from the commutativity of the diagram we obtain the desired relation $f \wedge g \simeq (-1)^{ij} g \wedge f$. □

Patterns in the stable homotopy groups of spheres begin to emerge only when one projects π_i^s onto its ***p*-components**, the quotient groups obtained by factoring out all elements of order relatively prime to the prime p. For $i > 0$ the p-component ${}_p\pi_i^s$ is of course isomorphic to the subgroup of π_i^s consisting of elements of order a power of p, but the quotient viewpoint is in some ways preferable. The ring structure on π_*^s passes down to a quotient ring structure on ${}_p\pi_*^s$, and in positive dimensions the isomorphism $\pi_i^s \approx \bigoplus_p ({}_p\pi_i^s)$ is a ring isomorphism.

The figure below is a schematic diagram of the 2-components of π_i^s for $i \le 60$. A vertical chain of n dots in the i^{th} column represents a $\mathbb{Z}_{2^n}$ summand of π_i^s. The bottom dot of such a chain denotes a generator of this summand, and the vertical segments denote multiplication by 2. The three generators η, ν, and σ in dimensions 1, 3, and 7 are represented by the Hopf maps $S^3 \to S^2$, $S^7 \to S^4$, $S^{15} \to S^8$ defined in the next section. Multiplication by η or ν is indicated by segments extending 1 or 3 units to the right. Thus for example we see the relation $\eta^3 = 4\nu$ in ${}_2\pi_3^s$. Remember that ${}_2\pi_3^s \approx \mathbb{Z}_8$ is a quotient of $\pi_3^s \approx \mathbb{Z}_{24}$, where the actual relation is $\eta^3 = 12\nu$ since $2\eta = 0$ implies $2\eta^3 = 0$.

Across the bottom of the diagram there is a repeated pattern of pairs of 'teeth.' This pattern continues to infinity, though with the spikes in dimensions $8k-1$ not all of the same

height, namely, the spike in dimension $2^m(2n+1)-1$ has height $m+1$. In the upper part of the diagram, however, there is considerably less regularity, and this complexity seems to persist in higher dimensions as well.

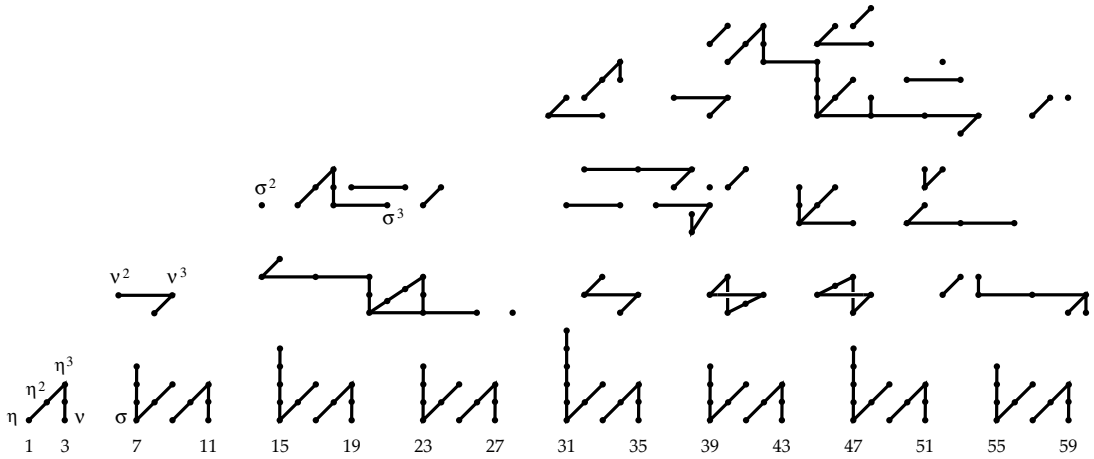

The next diagram shows the 3-components of π_i^s for $i \leq 100$, and the increase in regularity is quite noticeable. Here vertical segments denote multiplication by 3 and the other solid segments denote composition with elements $\alpha_1 \in {}_3\pi_3^s$ and $\beta_1 \in {}_3\pi_{10}^s$. The meaning of the dotted lines will be explained below. The most regular part of the diagram is the 'triadic ruler' across the bottom. This continues in the same pattern forever, with spikes of height $m+1$ in dimension $4k-1$ for 3^m the highest power of 3 dividing $4k$. Looking back at the $p=2$ diagram, one can see that the vertical sections of the 'teeth' form a 'dyadic ruler.'

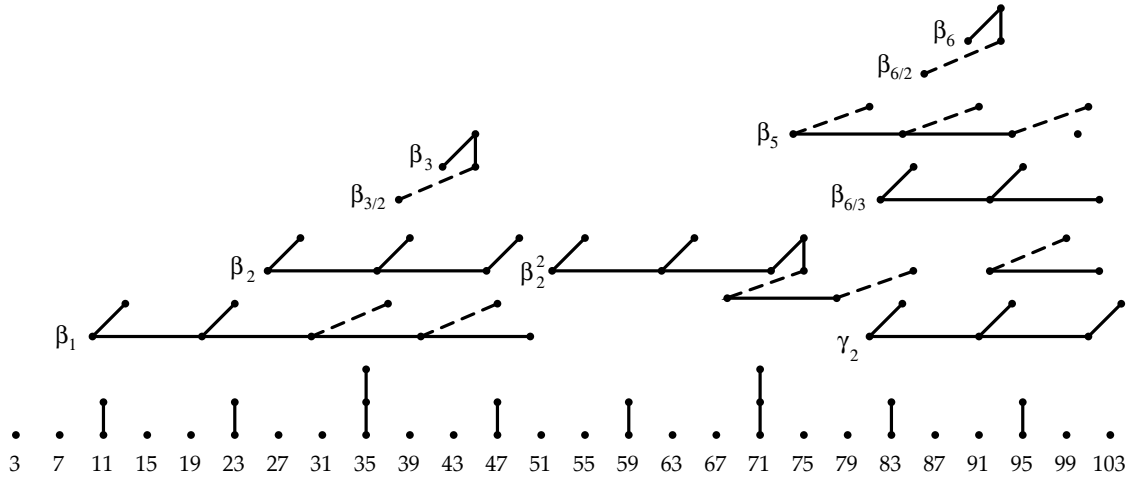

Even more regularity appears with larger primes, beginning with the case $p=5$ shown

in the next diagram. Again one has the infinite ruler, this time a 'pentadic' ruler, but there is also much regularity in the rest of the diagram.

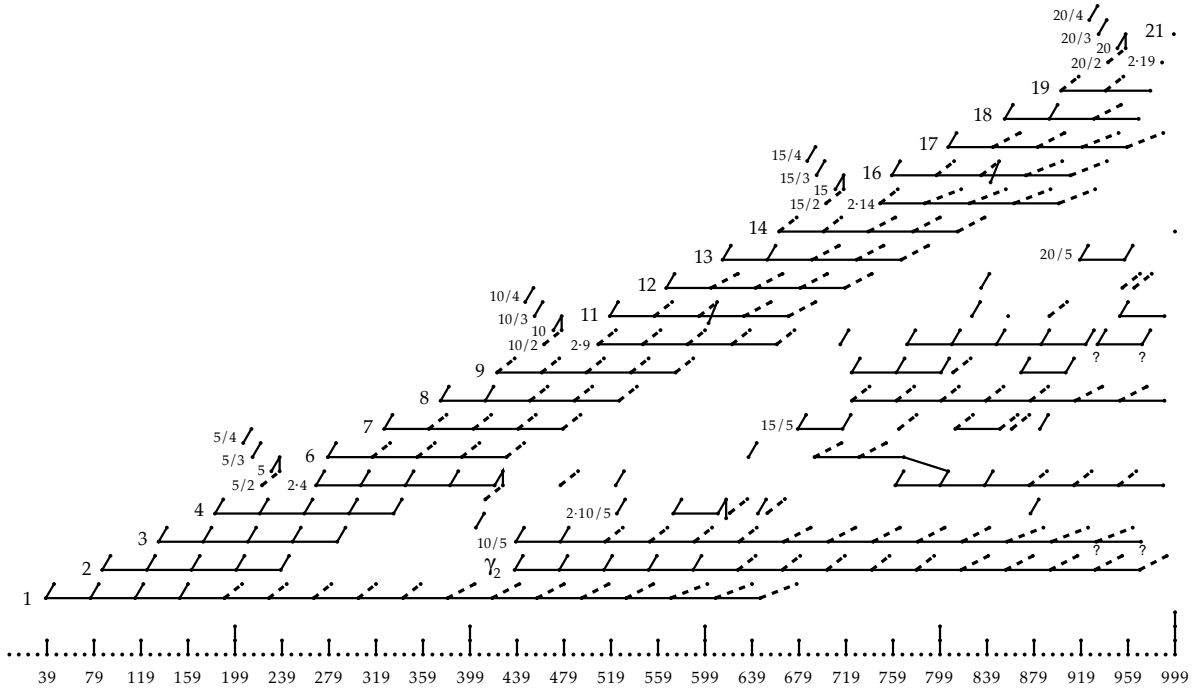

The four dots labeled with a question mark near the right edge of the diagram are hypothetical: the calculations in [Ravenel] do not decide whether these potential elements of ${}_5\pi_i^s$ for $i = 932, 933, 970$, and 971 actually exist.

These diagrams are drawn from tables in [Kochman] and [Kochman-Mahowald] for $p = 2$, and [Ravenel] for $p = 3, 5$.

For each p there is a similar infinite 'p-adic ruler,' corresponding to cyclic subgroups of order p^{m+1} in ${}_p\pi^s_{2j(p-1)-1}$ for all j, where p^m is the highest power of p dividing j. These subgroups are the p-components of a certain cyclic subgroup of π^s_{4k-1} known as $Im\,J$, the image of a homomorphism $J : \pi_{4k-1}(O) \to \pi^s_{4k-1}$. There are also $\mathbb{Z}_2$ subgroups of π_i^s for $i = 8k, 8k+1$ which comprise $Im\,J$ in these dimensions. In the diagram of ${}_2\pi_i^s$ these are the parts of the teeth connected to the spike in dimension $8k-1$. The J homomorphism will be studied in some detail in [VBKT].

There are other known infinite families in π^s_*, notably elements $\beta_n \in {}_p\pi^s_{2(p^2-1)n-2p}$ for $p \geq 5$ and $\gamma_n \in {}_p\pi^s_{2(p^3-1)n-2p^2-2p+1}$ for $p \geq 7$. The β_n's appear in the diagram for $p = 5$ as generators for the 'combs' along the upward diagonal, except when n is a multiple of 5 and the comb is generated by $\beta_2\beta_{n-1}$ rather than β_n. The element γ_2 generates the long comb starting in dimension 437, but $\gamma_3 = 0$. The element γ_4 in dimension 933

is one of the question marks. The theory behind these families β_n and γ_n and possible generalizations, as explained in [Ravenel], is one of the more esoteric branches of algebraic topology.

In π_*^s there are many compositions which are zero. One can get some idea of this from the diagrams above, where all sequences of edges break off after a short time. As a special instance of the vanishing of products, the commutativity formula implies that the square of an odd-dimensional element of odd order is zero. Much more generally, a theorem of Nishida says that every positive-dimensional element $\alpha \in \pi_*^s$ is nilpotent, i.e., $\alpha^n = 0$ for some n. For example, for the element $\beta_1 \in {}_5\pi_{38}^s$ the smallest such n is 18.

The widespread vanishing of products in π_*^s can be seen as limiting their usefulness in describing the structure of π_*^s. But it can also be used to construct new elements of π_*^s. Suppose one has maps $W \xrightarrow{f} X \xrightarrow{g} Y \xrightarrow{h} Z$ such that the compositions gf and hg are both homotopic to constant maps. A nullhomotopy of gf gives an extension of gf to a map $F : CW \to Y$, and a nullhomotopy of hg gives an extension of hg to a map $G : CX \to Z$. Regarding the suspension SW as the union of two cones CW, define the **Toda bracket** $\langle f, g, h\rangle : SW \to Z$ to be the composition $G(Cf)$ on one cone and hF on the other.

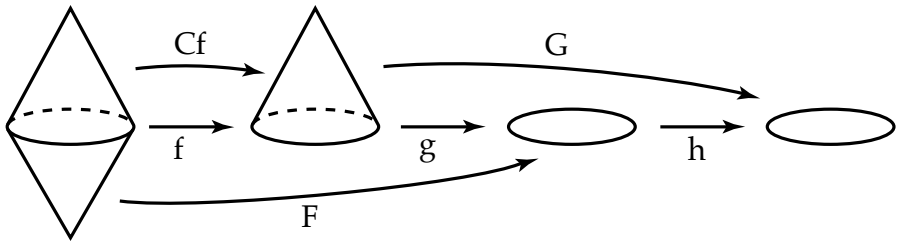

The map $\langle f, g, h\rangle$ is not uniquely determined by f, g, and h since it depends on the choices of the nullhomotopies. In the case of π_*^s, the various choices of $\langle f, g, h\rangle$ range over a coset of a certain subgroup, described in an exercise at the end of the section.

There are also higher-order Toda brackets $\langle f_1, \cdots, f_n\rangle$ defined in somewhat similar fashion. The dotted lines in the diagrams of ${}_3\pi_*^s$ and ${}_5\pi_*^s$ join an element x to a bracket element $\langle \alpha_1, \cdots, \alpha_1, x\rangle$. Most of the unlabelled elements above the rulers in all three diagrams are obtained from the labelled elements by compositions and brackets. For example, in ${}_2\pi_*^s$ the 8-dimensional element is $\langle \nu, \eta, \nu\rangle$ and the 14-dimensional elements are σ^2 and $\langle \nu, \langle \nu, \eta, \nu\rangle, 2, \eta\rangle$.

Exercises

1. Use Lemma 4.18 (or simplicial approximation, Theorem 2.34) to give a direct proof that $\pi_n(S^n) \approx \mathbb{Z}$ for $n \geq 2$, by the following steps.

(a) Via the degree homomorphism $\pi_n(S^n) \to \mathbb{Z}$, reduce the problem to showing every map $f : S^n \to S^n$ is homotopic to a multiple of the identity map.

(b) Use 4.18 to reduce to the case that there exists a point $q \in S^n$ with $f^{-1}(q) = \{p_1, \cdots, p_k\}$ and f is an invertible linear map near each p_i.

(c) For f as in (b), consider the composition gf where $g : S^n \to S^n$ collapses the complement of a small ball about q to the basepoint. Use this to reduce (b) further to the case $k = 1$.

(d) Finish the argument by showing that an invertible $n \times n$ matrix can be joined by a path of such matrices to either the identity matrix or the matrix of a reflection. (Use Gaussian elimination, for example.)

2. Show that the smash product $X \wedge Y$ of CW complexes X and Y is $(m+n+1)$-connected if X is m-connected and Y is n-connected, $m, n \geq 0$.

3. Give an example of a weak homotopy equivalence $X \to Y$ for which there does not exist a weak homotopy equivalence $Y \to X$.

4. Consider the equivalence relation '$\simeq_w$' generated by weak homotopy equivalence: $X \simeq_w Y$ if there are spaces $X = X_1, X_2, \cdots, X_n = Y$ with weak homotopy equivalences $X_i \to X_{i+1}$ or $X_i \leftarrow X_{i+1}$ for each i. Show that $X \simeq_w Y$ iff X and Y have a common CW approximation.

5. Show that $[X, Y]$ is finite if X is a finite connected CW complex and $\pi_i(Y)$ is finite for $i \leq dim\, X$.

6. For this problem it is convenient to use the notation X^n for an n-approximation to X and X_n for an $(n-1)$-connected covering of X, where X is a connected CW complex. Show that the notation X^n_m is unambiguous, i.e., that the operations of forming n-approximations and $(m-1)$-connected coverings commute, up to homotopy equivalence, so there is a diagram

$$\begin{array}{ccc} X_m & \longrightarrow & X^n_m \\ \downarrow & & \downarrow \\ X & \longrightarrow & X^n \end{array}$$

Thus $\pi_i(X^n_m) \approx \pi_i(X)$ for $m \leq i \leq n$ and all other homotopy groups of X^n_m are zero.

7. Show that if X is an $(n-1)$-connected CW complex, $n > 0$, and Y is a path-connected space with $\pi_i(Y) = 0$ for $i > n$, then the map $\langle X, Y \rangle \to Hom\big(\pi_n(X), \pi_n(Y)\big)$, $[f] \mapsto f_*$, is a bijection. Deduce that CW complex $K(G, n)$'s are uniquely determined, up to homotopy type, by G and n.

8. Construct a CW complex X with prescribed homotopy groups $\pi_i(X)$ and prescribed actions of $\pi_1(X)$ on the $\pi_i(X)$'s, by building X as a union of subcomplexes $X^1 \subset X_1 \subset$

$X^2 \subset X_2 \subset \cdots$ where X^n is obtained from X_{n-1} by attaching n-cells to realize prescribed generators for π_n, and X_n is obtained from X^n by attaching $(n+1)$-cells to realize prescribed relations for π_n. (When $n > 1$ these can be generators and relations for π_n as a π_1-module.)

9. Show that the indeterminacy of a Toda bracket $\langle f, g, h\rangle$ with $f \in \pi_i^s$, $g \in \pi_j^s$, $h \in \pi_k^s$ is the subgroup $f \cdot \pi_{j+k+1}^s + h \cdot \pi_{i+j+1}^s$ of $\pi_{i+j+k+1}^s$.

3. Fibrations

A 'short exact sequence of spaces' $A \hookrightarrow X \to X/A$ gives rise to a long exact sequence of homology groups, but not to a long exact sequence of homotopy groups due to the failure of excision. However, there is a different sort of 'short exact sequence of spaces' which does give a long exact sequence of homotopy groups. This sort of short exact sequence $F \longrightarrow E \xrightarrow{p} B$, called a fibration, is distinguished from the type $A \hookrightarrow X \to X/A$ in that it has more homogeneity: All the subspaces $p^{-1}(b) \subset E$, which are called **fibers**, are homotopy equivalent to F. In some of the most important cases, called fiber bundles, the fibers are all actually homeomorphic. Such homotopy-theoretic or topological fiber-homogeneity is rather like the algebraic homogeneity in a short exact sequence of groups $0 \to K \longrightarrow G \xrightarrow{p} H \to 0$ where the 'fibers' $p^{-1}(h)$ are the cosets of K in G.

In a few fiber bundles $F \to E \to B$ the space E is actually a group, F is a subgroup (though seldom a normal subgroup), and B is the space of left or right cosets. One of the nicest such examples is the Hopf bundle $S^1 \to S^3 \to S^2$ where S^3 is the group of quaternions of unit norm and S^1 is the subgroup of unit complex numbers. For this bundle, the long exact sequence of homotopy groups takes the form

$$\cdots \longrightarrow \pi_i(S^1) \longrightarrow \pi_i(S^3) \longrightarrow \pi_i(S^2) \longrightarrow \pi_{i-1}(S^1) \longrightarrow \pi_{i-1}(S^3) \longrightarrow \cdots$$

Since $\pi_i(S^1) = 0$ for $i > 1$ by Proposition 4.1, exactness of this sequence implies that there are isomorphisms $\pi_i(S^3) \approx \pi_i(S^2)$ for all $i \geq 3$, so in particular $\pi_3(S^2) \approx \pi_3(S^3) \approx \mathbb{Z}$. Also, the exact sequence gives an isomorphism $\pi_2(S^2) \approx \pi_1(S^1)$ since the two adjacent terms $\pi_2(S^3)$ and $\pi_1(S^3)$ are zero by cellular approximation. Thus we have a direct homotopy-theoretic proof that $\pi_2(S^2) \approx \mathbb{Z}$, and the excision theorem for homotopy groups then implies that $\pi_n(S^n) \approx \mathbb{Z}$ for all n, as we saw in §4.2.

The Homotopy Lifting Property

After these preliminary remarks, let us begin by defining the property which leads to a long exact sequence of homotopy groups. A map $p : E \to B$ is said to have the **homotopy lifting property** with respect to a space X if, given a homotopy $g_t : X \to B$ and a map $\widetilde{g}_0 : X \to E$ which lifts g_0, i.e., $p\widetilde{g}_0 = g_0$, then there exists a homotopy $\widetilde{g}_t : X \to E$ lifting g_t. A **fibration** is a map $p : E \to B$ having the homotopy lifting property with respect to all spaces X. This is stronger than what is needed to get a long exact sequence:

Theorem 4.32. *Suppose $p : E \to B$ has the homotopy lifting property with respect to disks D^k for all $k \geq 0$. Choose basepoints $b_0 \in B$ and $x_0 \in F = p^{-1}(b_0)$. Then*

the map $p_*:\pi_n(E,F,x_0)\to\pi_n(B,b_0)$ *is an isomorphism for all* $n\geq 1$. *Hence if* B *is path-connected there is a long exact sequence*

$$\cdots\to\pi_n(F,x_0)\to\pi_n(E,x_0)\xrightarrow{p_*}\pi_n(B,b_0)\to\pi_{n-1}(F,x_0)\to\cdots\to\pi_0(E,x_0)\to 0$$

In the proof we shall use a relative form of the homotopy lifting property. The map $p:E\to B$ is said to have the **homotopy lifting property for a pair** (X,A) if each homotopy $f_t:X\to B$ lifts to a homotopy $\widetilde{g}_t:X\to E$ starting with a given lift $\widetilde{g}_0$ and extending a given lift $\widetilde{g}_t:A\to E$. Observe that the homotopy lifting property for D^k is equivalent to the homotopy lifting property for $(D^k,\partial D^k)$ since the pairs $(D^k\times I, D^k\times\{0\})$ and $(D^k\times I, D^k\times\{0\}\cup\partial D^k\times I)$ are homeomorphic. This implies that the homotopy lifting property for disks is equivalent to the homotopy lifting property for all CW pairs (X,A). For by induction over the skeleta of X it suffices to construct a lifting $\widetilde{g}_t$ one cell of $X-A$ at a time. Composing with the characteristic map $\Phi:D^k\to X$ of a cell then reduces us to the case $(X,A)=(D^k,\partial D^k)$.

Proof: To show p_* is onto, represent an element of $\pi_n(B,b_0)$ by a map $f:(I^n,\partial I^n)\to(B,b_0)$. This can be viewed as a homotopy $I^{n-1}\times I\to B$ from the constant map to the constant map, taking ∂I^{n-1} to b_0 at all times. We will apply the homotopy lifting property to this homotopy, but with the last I coordinate, the time parameter, reversed. The constant map to x_0 is a lift $J^{n-1}\to E$, so the homotopy lifting property extends this to a lift $\widetilde{f}:I^n\to E$ with $\widetilde{f}(\partial I^n)\subset F$ since $f(\partial I^n)=b_0$. Then $\widetilde{f}$ represents an element of $\pi_n(E,F,x_0)$ with $p_*([\widetilde{f}])=[f]$.

Injectivity of p_* is similar. Given $\widetilde{f}_0,\widetilde{f}_1:(I^n,\partial I^n,J^{n-1})\to(E,F,x_0)$ such that $p_*([f_0])=p_*([f_1])$, let $G:I^n\times I\to B$ be a homotopy from $p\widetilde{f}_0$ to $p\widetilde{f}_1$ with $G(\partial I^n\times I)=b_0$. Changing viewpoint, we regard the n^{th} rather than the $(n+1)^{st}$ I factor of $I^n\times I$ as a homotopy parameter, with the ends of this I factor reversed as before. We have a partial lift $\widetilde{G}$ given by $\widetilde{f}_0$ on $I^n\times\{0\}$, $\widetilde{f}_1$ on $I^n\times\{1\}$, and the constant map to x_0 on $J^{n-1}\times I$. The homotopy lifting property extends this to a lift $\widetilde{G}:I^n\times I\to E$, giving a homotopy $\widetilde{f}_t:(I^n,\partial I^n,J^{n-1})\to(E,F,x_0)$ from $\widetilde{f}_0$ to $\widetilde{f}_1$. So p_* is injective.

When $\pi_n(B,b_0)$ is plugged in for $\pi_n(E,F,x_0)$ in the long exact sequence for the pair (E,F), the map $\pi_n(E,x_0)\to\pi_n(E,F,x_0)$ becomes the composition $\pi_n(E,x_0)\to\pi_n(E,F,x_0)\xrightarrow{p_*}\pi_n(B,b_0)$, which is just $p_*:\pi_n(E,x_0)\to\pi_n(B,b_0)$.

The 0 at the end of the long exact sequence, i.e., surjectivity of $\pi_0(F,x_0)\to\pi_0(E,x_0)$, comes from the hypothesis that B is path-connected, since a path in E from an arbitrary point $x\in E$ to F can be obtained by lifting a path in B from $p(x)$ to b_0. □

Proposition 4.33. *For a fibration $p:E \to B$, the fibers $F_b = p^{-1}(b)$ over each path component of B are all homotopy equivalent.*

Proof: A path $\gamma:I \to B$ can be viewed as a homotopy of the projection $F_{\gamma(0)} \to B$, with $F_{\gamma(0)}$ mapping to the point $\gamma(t)$ at time t. The homotopy lifting property then gives a homotopy $L_{\gamma,t}:F_{\gamma(0)} \to F_{\gamma(t)} \subset E$ starting with $L_{\gamma,0}$ the identity map of $F_{\gamma(0)}$ and ending with a map $L_\gamma = L_{\gamma,1:}F_{\gamma(0)} \to F_{\gamma(1)}$. The association $\gamma \mapsto L_\gamma$ has the following basic properties:

(a) If $\gamma \simeq \gamma' \text{rel } \partial I$, then $L_\gamma \simeq L_{\gamma'}$. In particular the homotopy class of L_γ is independent of the choice of lifting.

(b) For a composition $\gamma\eta$ of paths, the composition $L_\eta L_\gamma$ can be chosen for $L_{\gamma\eta}$.

To see the first of these, suppose $\gamma(s,t)$ is a homotopy from $\gamma(t)$ to $\gamma'(t)$, $(s,t) \in I \times I$, and we are given lifts $L_{\gamma,t}$ and $L_{\gamma',t}$. Define $L_{\gamma,s,t:}F_{\gamma(0)} \to F_{\gamma(s,t)}$ for $(s,t) \in I \times \{0\} \cup \partial I \times I$, by $L_{\gamma,0,t} = L_{\gamma,t}$, $L_{\gamma,1,t} = L_{\gamma',t}$, and $L_{\gamma,s,0} = \mathbb{1}$. The homotopy lifting property allows us to extend $L_{\gamma,s,t}$ to be defined for all $(s,t) \in I \times I$ since the pair $(I \times I, I \times \{0\} \cup \partial I \times I)$ is homeomorphic to the pair $(I \times I, I \times \{0\})$. The resulting $L_{\gamma,s,1}$ is a homotopy from L_γ to $L_{\gamma'}$, verifying (a). Property (b) holds since we can let $L_{\gamma\eta,t}$ be $L_{\gamma,2t}$ for $0 \le t \le \frac{1}{2}$ and $L_{\eta,2t-1}L_\gamma$ for $\frac{1}{2} \le t \le 1$.

From (a) and (b) it follows that L_γ is a homotopy equivalence with homotopy inverse $L_{\overline{\gamma}}$, where $\overline{\gamma}$ is the inverse path of γ. □

Fiber Bundles

A product space $B \times F$ can be viewed, asymmetrically, as a union of 'parallel' copies $F_b = \{b\} \times F$ of F, parametrized by $b \in B$. The notion of a fiber bundle is a natural generalization of this idea to spaces which are not products, but still consist of 'parallel' copies of one space F parametrized by another space B. Here is the precise definition: A **fiber bundle** structure on a space E, with fiber F, consists of a projection map $p:E \to B$ such that each point of B has a neighborhood U for which there is a homeomorphism $h:p^{-1}(U) \to U \times F$ making the following diagram commute

$$\begin{array}{ccc} p^{-1}(U) & \xrightarrow{h} & U \times F \\ & {\scriptstyle p}\searrow \quad \swarrow & \\ & U & \end{array}$$

where the unlabelled map is projection onto the first factor. Commutativity of the diagram means that h carries each fiber $F_b = p^{-1}(b)$ homeomorphically onto the copy $\{b\} \times F$ of

F. Thus the fibers F_b are arranged locally as in the product $B \times F$, though not necessarily globally. An h as above is called a **local trivialization** of the bundle. Since the first coordinate of h is just p, h is determined by its second coordinate, a map $p^{-1}(U) \to F$ which is a homeomorphism on each fiber F_b.

The fiber bundle structure is determined by the projection map $p: E \to B$, but to indicate what the fiber is we sometimes write a fiber bundle as $F \to E \to B$, a 'short exact sequence of spaces.' The space B is called the **base space** of the bundle, and E is the **total space**.

Examples.

(1) A fiber bundle with fiber a discrete space is a covering space. Conversely, a covering space whose fibers all have the same cardinality, for example a covering space over a connected base space, is a fiber bundle with discrete fiber.

(2) One of the simplest nontrivial bundles is the Mobius band $E = I \times [-1,1]/(0,v) \sim (1,-v)$, with $p: E \to S^1$ induced by the projection $I \times [-1,1] \to I$, so the fiber is $[-1,1]$. Glueing two copies of E together by the identity map between their boundary circles produces a Klein bottle, a bundle over S^1 with fiber S^1.

(3) Projective spaces yield interesting fiber bundles. In the real case we have the familiar covering spaces $S^n \to \mathbb{R}\mathrm{P}^n$ with fiber S^0. Over the complex numbers the analog of this is a fiber bundle $S^1 \to S^{2n+1} \to \mathbb{C}\mathrm{P}^n$. Here S^{2n+1} is the unit sphere in $\mathbb{C}^{n+1}$ and $\mathbb{C}\mathrm{P}^n$ is viewed as the quotient space of S^{2n+1} under the equivalence relation $(z_0, \cdots, z_n) \sim \lambda(z_0, \cdots, z_n)$ for $\lambda \in S^1$, the unit circle in $\mathbb{C}$. The projection $p: S^{2n+1} \to \mathbb{C}\mathrm{P}^n$ sends $(z_0, \cdots, z_n)$ to its equivalence class $[z_0, \cdots, z_n]$, so the fibers are copies of S^1. To see that the local triviality condition for fiber bundles is satisfied, let $U_i \subset \mathbb{C}\mathrm{P}^n$ be the open set of equivalence classes $[z_0, \cdots, z_n]$ with $z_i \neq 0$. Define $h_i: p^{-1}(U_i) \to U_i \times S^1$ by $h_i(z_0, \cdots, z_n) = ([z_0, \cdots, z_n], z_i/|z_i|)$. This takes fibers to fibers, and is a homeomorphism since its inverse is the map $([z_0, \cdots, z_n], \lambda) \mapsto \lambda |z_i| z_i^{-1}(z_0, \cdots, z_n)$, as one checks by calculation.

(4) The case $n = 1$ is particularly interesting since $\mathbb{C}\mathrm{P}^1 = S^2$ and the bundle becomes $S^1 \to S^3 \to S^2$ with fiber, total space, and base all spheres. This is known as the **Hopf bundle**, and is of a low enough dimension to be seen explicitly. The projection $S^3 \to S^2$ can be taken to be $(z_0, z_1) \mapsto z_0/z_1 \in \mathbb{C} \cup \{\infty\} = S^2$. In polar coordinates we have $p(r_0 e^{i\theta_0}, r_1 e^{i\theta_1}) = (r_0/r_1) e^{i(\theta_0 - \theta_1)}$ where $r_0^2 + r_1^2 = 1$. For a fixed ratio $\rho = r_0/r_1 \in (0, \infty)$ the angles θ_0 and θ_1 vary independently over S^1, so the points $(r_0 e^{i\theta_0}, r_1 e^{i\theta_1})$ form a torus $T_\rho \subset S^3$. Letting ρ vary, these disjoint tori T_ρ fill up S^3, if we include the limiting cases T_0 and T_∞ where the radii r_0 and r_1 are zero, making the tori T_0 and T_∞ degenerate

to circles. These two circles are the unit circles in the two $\mathbb{C}$ factors of $\mathbb{C}^2$, so under stereographic projection of S^3 from the point $(0,1)$ onto $\mathbb{R}^3$ they correspond to the unit circle in the xy-plane and the z-axis. The concentric tori T_ρ are then arranged as in the figure.

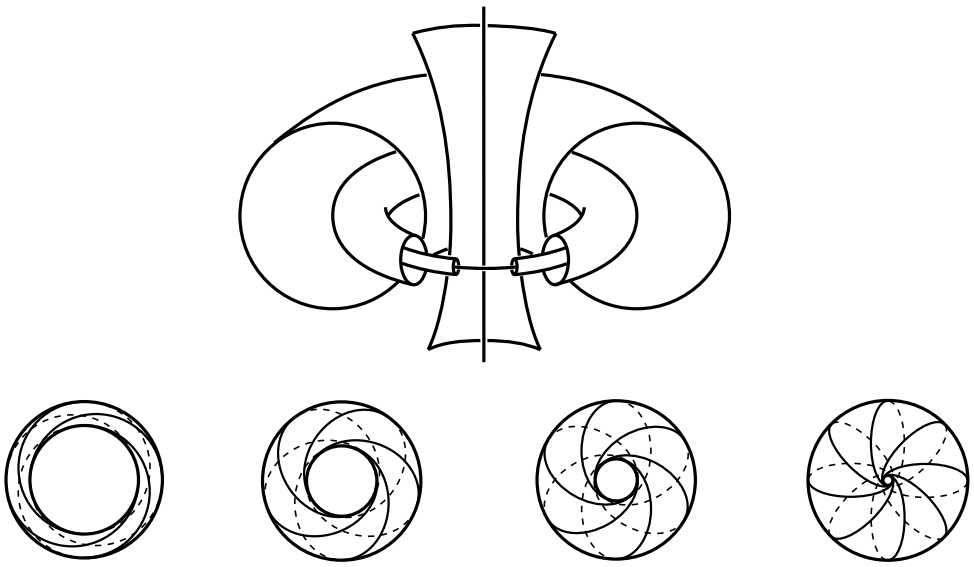

Each torus T_ρ is a union of circle fibers, the pairs (θ_0, θ_1) with $\theta_0 - \theta_1$ fixed. These fiber circles have slope 1 on the torus, winding around once longitudinally and once meridionally. As ρ goes to 0 or ∞ these fibers approach the circles T_0 and T_∞, which are also fibers.

(5) The construction of the bundle $S^1 \to S^{2n+1} \to \mathbb{C}\mathrm{P}^n$ also works when $n = \infty$, so there is a fiber bundle $S^1 \to S^\infty \to \mathbb{C}\mathrm{P}^\infty$.

(6) Generalizing in a different direction, the field $\mathbb{C}$ can be replaced by the quaternions $\mathbb{H}$, and the preceding formulas yield fiber bundles $S^3 \to S^{4n+3} \to \mathbb{H}\mathrm{P}^n$ over quaternionic projective space $\mathbb{H}\mathrm{P}^n$. Here the fiber S^3 is the unit quaternions, and S^{4n+3} is the unit sphere in $\mathbb{H}^{n+1}$. Taking $n = 1$ gives a second Hopf bundle $S^3 \to S^7 \to S^4 = \mathbb{H}\mathrm{P}^1$.

(7) There is another Hopf bundle $S^7 \to S^{15} \to S^8$, whose definition uses the nonassociative 8-dimensional algebra $\mathbb{O}$ of Cayley octonions. Elements of $\mathbb{O}$ are pairs of quaternions (a_1, a_2) with multiplication defined by $(a_1, a_2)(b_1, b_2) = (a_1 b_1 - \overline{b}_2 a_2, a_1 \overline{b}_2 + b_2 a_1)$. Regarding S^{15} as the unit sphere in the 16-dimensional vector space $\mathbb{O}^2$, the projection map $p\colon S^{15} \to S^8 = \mathbb{O} \cup \{\infty\}$ is $(z_0, z_1) \mapsto z_0 z_1^{-1}$, just as for the other Hopf bundles, but a little care is needed to show this is a fiber bundle with fiber S^7, the unit octonions. Let U_0 and U_1 be the complements of ∞ and 0 in the base space $\mathbb{O} \cup \{\infty\}$. Define

$h_i : p^{-1}(U_i) \to U_i \times S^7$ and $g_i : U_i \times S^7 \to p^{-1}(U_i)$ by

$$\begin{aligned} h_0(z_0, z_1) &= (z_0 z_1^{-1}, z_1/|z_1|), & g_0(z, w) &= (zw, w)/|(zw, w)| \\ h_1(z_0, z_1) &= (z_0 z_1^{-1}, z_0/|z_0|), & g_1(z, w) &= (w, z^{-1}w)/|(w, z^{-1}w)| \end{aligned}$$

If one assumes the known fact that any subalgebra of $\mathbb{O}$ generated by two elements is associative, then it is a simple matter to check that g_i and h_i are inverse homeomorphisms, so we have a fiber bundle $S^7 \to S^{15} \to S^8$. Actually, the calculation that g_i and h_i are inverses needs only the following more elementary facts about octonions z, w, where the conjugate $\overline{z}$ of $z = (a_1, a_2)$ is defined by the expected formula $\overline{z} = (\overline{a}_1, -a_2)$:

(1) $\mathbb{R} \subset \mathbb{O}$ as the pairs $(r, 0)$ with $r \in \mathbb{R}$, and $rz = zr$ for all $z \in \mathbb{O}$.
(2) $|z|^2 = z\overline{z} = \overline{z}z$, hence $z^{-1} = \overline{z}/|z|^2$.
(3) $|zw| = |z||w|$.
(4) $\overline{zw} = \overline{w}\,\overline{z}$, hence $(zw)^{-1} = w^{-1}z^{-1}$.
(5) $z(\overline{z}w) = (z\overline{z})w$ and $(zw)\overline{w} = z(w\overline{w})$, hence $z(z^{-1}w) = w$ and $(zw)w^{-1} = z$.

These facts can be checked by somewhat tedious direct calculation, though more elegant derivations can be found in Chapter 8 of [Ebbinghaus].

There is an octonion 'projective plane' $\mathbb{O}\mathrm{P}^2$, obtained by attaching a cell e^{16} to S^8 via the Hopf map $S^{15} \to S^8$, just as $\mathbb{C}\mathrm{P}^2$ and $\mathbb{H}\mathrm{P}^2$ are obtained from the other Hopf maps. However, there is no octonion analog of $\mathbb{R}\mathrm{P}^n$, $\mathbb{C}\mathrm{P}^n$, and $\mathbb{H}\mathrm{P}^n$ for $n > 2$ since associativity of multiplication is needed for the relation $(z_0, \cdots, z_n) \sim \lambda(z_0, \cdots, z_n)$ to be an equivalence relation.

There are no fiber bundles with fiber, total space, and base space spheres of other dimensions than in these Hopf bundle examples. This is discussed in an exercise for §4.4d, which reduces the question to the famous 'Hopf invariant one' problem.

Since fiber bundles are a generalization of covering spaces, it is natural to ask whether the covering homotopy property holds for fiber bundles. As one can readily see by looking at simple examples like the product bundle $I \times I \to I$, lifted homotopies are generally not unique, as they are for covering spaces. But the existence of lifted homotopies does hold in fair generality. For example, a theorem of Huebsch and Hurewicz, proved in [Spanier, Chapter 2.7], says that fiber bundles over paracompact base spaces are fibrations, i.e., have the homotopy lifting property for all spaces. An easier result which suffices for many purposes in homotopy theory is:

Proposition 4.34. *A fiber bundle $p : E \to B$ has the homotopy lifting property with respect to all CW pairs (X, A).*

Proof: As noted earlier, the homotopy lifting property for CW pairs is equivalent to the homotopy lifting property for disks, or equivalently, cubes. Let $G: I^n \times I \to B$, $G(x,t) = g_t(x)$, be a homotopy we wish to lift, starting with a given lift $\widetilde{g}_0$ of g_0. Let $\{U_\alpha\}$ be an open cover of B for which there are local trivializations $h_\alpha: p^{-1}(U_\alpha) \to U_\alpha \times F$. Using compactness of $I^n \times I$, we may choose points $0 = t_1 < t_2 < \cdots < t_{k+1} = 1$ partitioning I into subintervals $I_j = [t_j, t_{j+1}]$ so that each of the resulting subcubes $I_{j_1} \times \cdots \times I_{j_{n+1}}$ of $I^n \times I$ is mapped by G into a single U_α. The subcubes $I_{j_1} \times \cdots \times I_{j_n}$ of I^n give a CW structure on I^n, and we may assume by induction on n that $\widetilde{g}_t$ has already been constructed over the $(n-1)$-skeleton of this CW structure. To extend this $\widetilde{g}_t$ over a cube $I_{j_1} \times \cdots \times I_{j_n}$ we may proceed in stages, first constructing $\widetilde{g}_t$ for t in the interval I_1, then for t in I_2, etc. This in effect reduces us to the case that no subdivision of $I^n \times I$ is necessary, so G maps all of $I^n \times I$ to a single U_α. Then we have $\widetilde{G}(I^n \times \{0\} \cup \partial I^n \times I) \subset p^{-1}(U_\alpha)$, and composing $\widetilde{G}$ with the local trivialization h_α reduces us to the case of a product bundle $U_\alpha \times F$. In this case the first coordinate of a lift $\widetilde{g}_t$ is just the given g_t, so only the second coordinate needs to be constructed. This can be obtained as a composition $I^n \times I \to I^n \times \{0\} \cup \partial I^n \times I \to F$ where the first map is a retraction and the second map is what we are given. □

Example 1. Applying this theorem to a covering space $p: E \to B$ with E and B path-connected, and discrete fiber F, the resulting long exact sequence of homotopy groups yields Proposition 4.1 that $p_*: \pi_n(E) \to \pi_n(B)$ is an isomorphism for $n \geq 2$. We also obtain a short exact sequence $0 \to \pi_1(E) \to \pi_1(B) \to \pi_0(F) \to 0$, consistent with the covering space theory facts that $p_*: \pi_1(E) \to \pi_1(B)$ is injective and that the fiber F can be identified, via path-lifting, with the set of cosets of $p_*\pi_1(E)$ in $\pi_1(B)$.

Example 2. From the bundle $S^1 \to S^\infty \to \mathbb{C}\mathrm{P}^\infty$ we obtain $\pi_i(\mathbb{C}\mathrm{P}^\infty) \approx \pi_{i-1}(S^1)$ for all i since S^∞ is contractible. Thus $\mathbb{C}\mathrm{P}^\infty$ is a $K(\mathbb{Z}, 2)$. In similar fashion the bundle $S^3 \to S^\infty \to \mathbb{H}\mathrm{P}^\infty$ gives $\pi_i(\mathbb{H}\mathrm{P}^\infty) \approx \pi_{i-1}(S^3)$ for all i, but these homotopy groups are far more complicated than for $\mathbb{C}\mathrm{P}^\infty$ and S^1. In particular, $\mathbb{H}\mathrm{P}^\infty$ is not a $K(\mathbb{Z}, 3)$.

Example 3. The long exact sequence for the Hopf bundle $S^1 \to S^3 \to S^2$ gives isomorphisms $\pi_2(S^2) \approx \pi_1(S^1)$ and $\pi_n(S^3) \approx \pi_n(S^2)$ for all $n \geq 3$, as we noted near the beginning of this section.

From the last two examples we see that S^2 and $S^3 \times \mathbb{C}\mathrm{P}^\infty$ are simply-connected CW complexes with isomorphic homotopy groups, though they are not homotopy equivalent since they have quite different homology groups.

Example 4: Whitehead Products. We compute $\pi_3(\bigvee_\alpha S^2_\alpha)$, showing that it is free abelian

with basis consisting of the Hopf maps $S^3 \to S^2_\alpha \subset \bigvee_\alpha S^2_\alpha$ together with the attaching maps $S^3 \to S^2_\alpha \vee S^2_\beta \subset \bigvee_\alpha S^2_\alpha$ of the cells $e^2_\alpha \times e^2_\beta$ in the products $S^2_\alpha \times S^2_\beta$ for all unordered pairs $\alpha \neq \beta$. Suppose first that there are only finitely many summands S^2_α. For a finite product $\prod_\alpha X_\alpha$ of path-connected spaces, the map $\pi_n(\bigvee_\alpha X_\alpha) \to \pi_n(\prod_\alpha X_\alpha)$ induced by inclusion is surjective since the group $\pi_n(\prod_\alpha X_\alpha) \approx \bigoplus_\alpha \pi_n(X_\alpha)$ is generated by the subgroups $\pi_n(X_\alpha)$. Thus the long exact sequence of homotopy groups for the pair $(\prod_\alpha X_\alpha, \bigvee_\alpha X_\alpha)$ breaks up into short exact sequences

$$0 \longrightarrow \pi_{n+1}(\textstyle\prod_\alpha X_\alpha, \bigvee_\alpha X_\alpha) \longrightarrow \pi_n(\bigvee_\alpha X_\alpha) \longrightarrow \pi_n(\prod_\alpha X_\alpha) \longrightarrow 0$$

For $n > 1$ these short exact sequences split since the inclusions $X_\alpha \hookrightarrow \bigvee_\alpha X_\alpha$ induce maps $\pi_n(X_\alpha) \to \pi_n(\bigvee_\alpha X_\alpha)$ and hence a splitting homomorphism $\bigoplus_\alpha \pi_n(X_\alpha) \to \pi_n(\bigvee_\alpha X_\alpha)$ since the latter group is abelian if $n > 1$. Taking $X_\alpha = S^2_\alpha$ and $n = 3$, we get an isomorphism

$$\pi_3(\textstyle\bigvee_\alpha S^2_\alpha) \approx \pi_4(\prod_\alpha S^2_\alpha, \bigvee_\alpha S^2_\alpha) \,\oplus\, \big(\bigoplus_\alpha \pi_3(S^2_\alpha)\big)$$

The factor $\bigoplus_\alpha \pi_3(S^2_\alpha)$ is free with basis the Hopf maps $S^3 \to S^2_\alpha$ by the preceding example. For the other factor, Proposition 4.27 gives an isomorphism $\pi_4(\prod_\alpha S^2_\alpha, \bigvee_\alpha S^2_\alpha) \approx \pi_4(\prod_\alpha S^2_\alpha / \bigvee_\alpha S^2_\alpha)$. The quotient $\prod_\alpha S^2_\alpha / \bigvee_\alpha S^2_\alpha$ has 5-skeleton a wedge of spheres $S^4_{\alpha\beta}$ for $\alpha \neq \beta$, so $\pi_4(\prod_\alpha S^2_\alpha / \bigvee_\alpha S^2_\alpha) \approx \pi_4(\bigvee_{\alpha\beta} S^4_{\alpha\beta})$ is free with basis the inclusions $S^4_{\alpha\beta} \hookrightarrow \bigvee_{\alpha\beta} S^4_{\alpha\beta}$. Hence $\pi_4(\prod_\alpha S^2_\alpha, \bigvee_\alpha S^2_\alpha)$ is free with basis the characteristic maps of the 4-cells $e^2_\alpha \times e^2_\beta$. Via the injection $\partial : \pi_4(\prod_\alpha S^2_\alpha, \bigvee_\alpha S^2_\alpha) \to \pi_3(\bigvee_\alpha S^2_\alpha)$ this means that the attaching maps of the cells $e^2_\alpha \times e^2_\beta$ form a basis for the summand $Im(\partial)$ of $\pi_3(\bigvee_\alpha S^2_\alpha)$. This finishes the proof for the case of finitely many summands S^2_α. The case of infinitely many S^2_α's follows immediately since any map $S^3 \to \bigvee_\alpha S^2_\alpha$ has compact image, lying in a finite union of summands, and similarly for any homotopy between such maps.

The maps $S^3 \to S^2_\alpha \vee S^2_\beta$ in this example are expressible in terms of a product in homotopy groups called the **Whitehead product**, defined as follows. Given basepoint-preserving maps $f : S^k \to X$ and $g : S^\ell \to X$, let $[f, g] : S^{k+\ell-1} \to X$ be the composition $S^{k+\ell-1} \longrightarrow S^k \vee S^\ell \xrightarrow{f \vee g} X$ where the first map is the attaching map of the $(k+\ell)$-cell of $S^k \times S^\ell$ with its usual CW structure. Since homotopies of f or g give rise to homotopies of $[f, g]$, we have a well-defined product $\pi_k(X) \times \pi_\ell(X) \to \pi_{k+\ell-1}(X)$. The notation $[f, g]$ is used since for $k = \ell = 1$ this is just the commutator product in $\pi_1(X)$. Similarly, it is an exercise to show that when $k = 1$ and $\ell > 1$, $[f, g]$ is the difference between g and its image under the π_1-action of f.

In these terms the map $S^3 \to S^2_\alpha \vee S^2_\beta$ in Example 4 is the Whitehead product $[\iota_\alpha, \iota_\beta]$ of the two inclusions of S^2 into the wedge sum as the summands S^2_α and S^2_β. Another

example of a Whitehead product which we have encountered previously is $[\mathbb{1},\mathbb{1}]\colon S^{2n-1}\to S^n$, which is the attaching map of the $2n$-cell of the space $J(S^n)$ considered in §3.2.

The calculation of $\pi_3(\bigvee_\alpha S^2_\alpha)$ is the first nontrivial case of a more general theorem of Hilton calculating all the homotopy groups of any wedge sum of spheres in terms of homotopy groups of spheres, using Whitehead products. A further generalization by Milnor extends this to wedge sums of suspensions of arbitrary connected CW complexes. See [G.Whitehead] for an exposition of these results and further information on Whitehead products.

Example 5: Stiefel and Grassmann Manifolds. The fiber bundles with total space a sphere and base space a projective space considered above are in fact just the cases $n=1$ of a broader family of fiber bundles in each of the real, complex, and quaternionic cases:

$$\begin{array}{ll} O(n)\longrightarrow V_n(\mathbb{R}^k)\longrightarrow G_n(\mathbb{R}^k) & \qquad O(n)\longrightarrow V_n(\mathbb{R}^\infty)\longrightarrow G_n(\mathbb{R}^\infty)\\ U(n)\longrightarrow V_n(\mathbb{C}^k)\longrightarrow G_n(\mathbb{C}^k) & \qquad U(n)\longrightarrow V_n(\mathbb{C}^\infty)\longrightarrow G_n(\mathbb{C}^\infty)\\ Sp(n)\longrightarrow V_n(\mathbb{H}^k)\longrightarrow G_n(\mathbb{H}^k) & \qquad Sp(n)\longrightarrow V_n(\mathbb{H}^\infty)\longrightarrow G_n(\mathbb{H}^\infty)\end{array}$$

Taking the real case first, the Stiefel manifold $V_n(\mathbb{R}^k)$ is the space of n-frames in $\mathbb{R}^k$, i.e., n-tuples of orthonormal vectors in $\mathbb{R}^k$. This is topologized as a subspace of the product of n copies of the unit sphere in $\mathbb{R}^k$. The Grassmann manifold $G_n(\mathbb{R}^k)$ is the space of n-dimensional vector subspaces of $\mathbb{R}^k$. There is a natural surjection $p\colon V_n(\mathbb{R}^k)\to G_n(\mathbb{R}^k)$ sending an n-frame to the subspace it spans, and $G_n(\mathbb{R}^k)$ is topologized as a quotient space of $V_n(\mathbb{R}^k)$ via this projection. The fibers of the map p are the spaces of n-frames in a fixed n-plane in $\mathbb{R}^k$ and so are homeomorphic to $V_n(\mathbb{R}^n)$. An n-frame in $\mathbb{R}^n$ is the same as an orthogonal $n\times n$ matrix, regarding the columns of the matrix as an n-frame, so the fiber can also be described as the orthogonal group $O(n)$. There is no difficulty in allowing $k=\infty$ in these definitions, and in fact $V_n(\mathbb{R}^\infty)=\bigcup_k V_n(\mathbb{R}^k)$ and $G_n(\mathbb{R}^\infty)=\bigcup_k G_n(\mathbb{R}^k)$.

The complex and quaternionic Stiefel manifolds and Grassmann manifolds are defined in the same way using the usual Hermitian inner products in $\mathbb{C}^k$ and $\mathbb{H}^k$. The unitary group $U(n)$ consists of $n\times n$ matrices whoses columns form orthonormal bases for $\mathbb{C}^n$, and the symplectic group is the quaternionic analog of this.

We should explain why the various projection maps $V_n\to G_n$ are fiber bundles. Let us take the real case for concreteness, though the argument is the same in all cases. If we fix an n-plane $P\in G_n(\mathbb{R}^k)$ and choose an orthonormal basis for P, then we obtain continuously-varying orthonormal bases for all n-planes P' in a neighborhood U of P by projecting the basis for P orthogonally onto P' to obtain a non-orthonormal basis for P', then applying the Gram-Schmidt process to this basis to make it orthonormal. The

formulas for the Gram-Schmidt process show that it is continuous. Having orthonormal bases for all n-planes in U, we can use these to identify these n-planes with $\mathbb{R}^n$, hence n-frames in these n-planes are identified with n-frames in $\mathbb{R}^n$, and so $p^{-1}(U)$ is identified with $U \times V_n(\mathbb{R}^n)$. This argument works for $k = \infty$ as well as for finite k.

In the case $n = 1$ the total spaces V_1 are spheres, which are highly connected, and the same is true in general:

— $V_n(\mathbb{R}^k)$ is $(k-n-1)$-connected.
— $V_n(\mathbb{C}^k)$ is $(2k-2n)$-connected.
— $V_n(\mathbb{H}^k)$ is $(4k-4n+2)$-connected.
— $V_n(\mathbb{R}^\infty)$, $V_n(\mathbb{C}^\infty)$, and $V_n(\mathbb{H}^\infty)$ are contractible.

The first three statements will be proved in the next example. For the last statement the argument is the same in the three cases, so let us consider the real case. Define a homotopy $h_t : \mathbb{R}^\infty \to \mathbb{R}^\infty$ by $h_t(x_1, x_2, \cdots) = (1-t)(x_1, x_2, \cdots) + t(0, x_1, x_2, \cdots)$. This is linear for each t, and its kernel is easily checked to be trivial. So if we apply h_t to an n-frame we get an n-tuple of independent vectors which can be made orthonormal by the Gram-Schmidt formulas. Thus we have a deformation retraction, in the weak sense, of $V_n(\mathbb{R}^\infty)$ onto the subspace of n-frames with first coordinate zero. Iterating this n times, we deform into the subspace of n-frames with first n coordinates zero. For such an n-frame $(v_1, \cdots, v_n)$ define a homotopy $(1-t)(v_1, \cdots, v_n) + t(e_1, \cdots, e_n)$ where e_i is the i^{th} standard basis vector in $\mathbb{R}^\infty$. This homotopy preserves linear independence, so after applying Gram-Schmidt again, we have a deformation through n-frames, which finishes the construction of a contraction of $V_n(\mathbb{R}^\infty)$.

Since $V_n(\mathbb{R}^\infty)$ is contractible, we obtain isomorphisms $\pi_i O(n) \approx \pi_{i+1} G_n(\mathbb{R}^\infty)$ for all i and n, and similarly in the complex and quaternionic cases.

Example 6. For $m < n \leq k$ there are fiber bundles

$$V_{n-m}(\mathbb{R}^{k-m}) \longrightarrow V_n(\mathbb{R}^k) \xrightarrow{p} V_m(\mathbb{R}^k)$$

where the projection p sends an n-frame onto the m-frame formed by its first m vectors, so the fiber consists of $(n-m)$-frames in the $(k-m)$-plane orthogonal to a given m-frame. Local trivializations can be constructed as follows. For an m-frame F, choose an orthonormal basis for the $(k-m)$-plane orthogonal to F. This determines orthonormal bases for the $(k-m)$-planes orthogonal to all nearby m-frames by orthogonal projection and Gram-Schmidt, as in Example 5. This allows us to identify these $(k-m)$-planes with $\mathbb{R}^{k-m}$, and in particular the fibers near $p^{-1}(F)$ are identified with $V_{n-m}(\mathbb{R}^{k-m})$, giving a local trivialization.

There are analogous bundles in the complex and quaternionic cases as well, with local triviality shown in the same way.

Taking $m = 1$, we have bundles $V_{n-1}(\mathbb{R}^{k-1}) \to V_n(\mathbb{R}^k) \to S^{k-1}$ whose associated long exact sequence of homotopy groups allows us deduce that $V_n(\mathbb{R}^k)$ is $(k-n-1)$-connected by induction on n. In the complex and quaternionic cases the same argument yields the other connectivity statements in Example 5.

Taking $k = n$ we obtain fiber bundles $O(k-m) \to O(k) \to V_m(\mathbb{R}^k)$. The fibers are in fact just the cosets $\alpha O(k-m)$ for $\alpha \in O(k)$, where $O(k-m)$ is regarded as the subgroup of $O(k)$ fixing the first m standard basis vectors. So we see that $V_m(\mathbb{R}^k)$ is identifiable with the coset space $O(k)/O(k-m)$, or in other words the orbit space for the free action of $O(k-m)$ on $O(k)$ by right-multiplication. In similar fashion one can see that $G_m(\mathbb{R}^k)$ is the coset space $O(k)/\big(O(m)\times O(k-m)\big)$ where the subgroup $O(m)\times O(k-m) \subset O(k)$ consists of the orthogonal transformations taking the m-plane spanned by the first m standard basis vectors to itself. The corresponding observations apply also in the complex and quaternionic cases, with the unitary and symplectic groups.

Example 7: Bott Periodicity. Specializing the preceding example further by taking $m = 1$ and $k = n$ we obtain bundles

$$\begin{aligned} O(n-1) &\longrightarrow O(n) \xrightarrow{p} S^{n-1} \\ U(n-1) &\longrightarrow U(n) \xrightarrow{p} S^{2n-1} \\ Sp(n-1) &\longrightarrow Sp(n) \xrightarrow{p} S^{4n-1} \end{aligned}$$

The map p can be described as evaluation of an orthogonal, unitary, or symplectic transformation on a fixed unit vector. These bundles show that computing homotopy groups of $O(n)$, $U(n)$, and $Sp(n)$ should be at least as difficult as computing homotopy groups of spheres. For example, if one knew the homotopy groups of $O(n)$ and $O(n-1)$, one could very nearly compute the homotopy groups of S^{n-1} from the long exact sequence for the first bundle.

However, the bundles above do imply a very interesting stability property. In the real case, the inclusion $O(n-1) \hookrightarrow O(n)$ induces an isomorphism on π_i for $i < n-1$, from the long exact sequence of the first bundle. Hence the groups $\pi_i O(n)$ are independent of n if n is sufficiently large, and the same is true for the groups $\pi_i U(n)$ and $\pi_i Sp(n)$ via the other two bundles. One of the most surprising results in all of algebraic topology is the Bott Periodicity Theorem which asserts that these stable groups repeat periodically, with a period of eight for SO and Sp and a period of two for U. Their values are given in the following table:

i mod 8	0	1	2	3	4	5	6	7
$\pi_i O(n)$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	0	$\mathbb{Z}$	0	0	0	$\mathbb{Z}$
$\pi_i U(n)$	0	$\mathbb{Z}$	0	$\mathbb{Z}$	0	$\mathbb{Z}$	0	$\mathbb{Z}$
$\pi_i Sp(n)$	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	0	$\mathbb{Z}$

Path Fibrations and Loopspaces

There is a simple but extremely useful way to turn arbitrary mappings into fibrations. Given a map $f: A \to B$, let E_f be the space of pairs (a, γ) where $a \in A$ and $\gamma: I \to B$ is a path in B with $\gamma(0) = f(a)$. We topologize E_f as a subspace of $A \times B^I$, where B^I is the space of mappings $I \to B$ with the compact-open topology; see the Appendix for the definition and basic properties of this topology, in particular Proposition A.13 which we will be using shortly.

Proposition 4.35. *For $E_f = \{\, (a, \gamma) \in A \times B^I \mid \gamma(0) = f(a) \,\}$ associated to $f: A \to B$, the map $p: E_f \to B$, $p(a, \gamma) = \gamma(1)$, is a fibration.*

Proof: Continuity of p follows from (a) of Proposition A.13 which says that the evaluation map $B^I \times I \to B$, $(\gamma, s) \mapsto \gamma(s)$, is continuous, hence also its restriction to $B^I \times \{1\}$.

Let a homotopy $g_t: X \to B$ and a lift $\widetilde{g}_0: X \to E_f$ of g_0 be given. Write $\widetilde{g}_0(x) = (h(x), \gamma_x)$ for $h: X \to A$ and $\gamma_x: I \to B$. Define a lift $\widetilde{g}_t: X \to E_f$ by $\widetilde{g}_t(x) = (h(x), \gamma_x \cdot g_{[0,t]}(x))$, the second coordinate being the path γ_x followed by the path traced out by $g_s(x)$ for $0 \le s \le t$. This composition of paths is defined since $g_0(x) = \gamma_x(1)$. Continuity of $\widetilde{g}_t$ as a map $X \times I \to E_f \subset A \times B^I$ follows from (b) of Proposition A.13 which implies that continuity of a map $X \times I \to A \times B^I$ is equivalent to continuity of the associated map $X \times I \times I \to A \times B$. □

For a map $f: A \to B$, we can regard A as the subspace of E_f consisting of pairs (a, γ) with γ the constant path at $f(a)$, and E_f deformation retracts onto this subspace by restricting all the paths γ to shorter and shorter initial segments. Thus an arbitrary map $f: A \to B$ factors as $A \hookrightarrow E_f \to B$, the composition of a homotopy equivalence and a fibration. The fiber F_f of $E_f \to B$ is called the **homotopy fiber** of f. It consists of all pairs (a, γ) with $a \in A$ and γ a path in B from $f(a)$ to a basepoint $b_0 \in B$.

If $f: A \to B$ is the inclusion of a subspace, then E_f is the space of paths in B starting at points of A. In this case a map $(I^{i+1}, \partial I^{i+1}, J^i) \to (B, A, x_0)$ is the same as a map $(I^i, \partial I^i) \to (F_f, \gamma_0)$ where γ_0 is the constant path at x_0 and F_f is the fiber of E_f over x_0. This means that $\pi_{i+1}(B, A, x_0)$ can be identified with $\pi_i(F_f, \gamma_0)$, hence the long exact sequences of homotopy groups of the pair (B, A) and of the fibration $E_f \to B$ can be identified.

An important special case is when f is the inclusion of the basepoint b_0 into B. Then E_f is the space PB of paths in B starting at b_0, and $p: PB \to B$ sends each path to its endpoint. The fiber $p^{-1}(b_0)$ is the **loopspace** ΩB consisting of all loops in B based at b_0. Since PB is contractible by progressively truncating paths, the long exact sequence of homotopy groups for the path fibration $PB \to B$ yields:

Proposition 4.36. *$\pi_n(X, x_0) \approx \pi_{n-1}(\Omega X, x_0)$ for all n.* □

This can also be seen directly. A map $I^n \times I \to X$ is equivalent to a map $I^n \to X^I$, so a map $(I^{n+1}, \partial I^{n+1}) \to (X, x_0)$ is equivalent to a map $(I^n, \partial I^n) \to (\Omega X, x_0)$ where the latter 'x_0' means the constant loop at the basepoint $x_0 \in X$. Similarly, a homotopy of maps $(I^{n+1}, \partial I^{n+1}) \to (X, x_0)$ is equivalent to a homotopy of maps $(I^n, \partial I^n) \to (\Omega X, x_0)$.

The operation of forming loopspaces can be iterated. For example, $\Omega^2 X = \Omega(\Omega X)$ is the space of maps $(I^2, \partial I^2) \to (X, x_0)$, and similarly $\Omega^n X$ is the space of maps $(I^n, \partial I^n) \to (X, x_0)$. Hence $\pi_n(X, x_0) = \pi_0(\Omega^n X)$.

Even if X is a rather small space like S^1, ΩX is usually quite large, generally infinite-dimensional. Nevertheless, if X is homotopy equivalent to a CW complex then so is ΩX, according to a theorem in [Milnor 59] which says more generally that if X is a CW complex and Y is a finite CW complex, then the subspace of X^Y consisting of maps taking subcomplexes $Y_1, \cdots, Y_n \subset Y$ to subcomplexes $X_1, \cdots, X_n \subset X$, respectively, is homotopy equivalent to a CW complex. This paper also contains the theorem that the homotopy fiber of a map between CW complexes has the homotopy type of a CW complex. These results can usually be circumvented by using CW approximations, though it is reassuring to know they are available if needed, or if one doesn't want to bother with CW approximations.

We have already encountered a few spaces which are homotopy equivalent to loopspaces. For example, an Eilenberg-MacLane space $K(\pi, n)$ has the homotopy type of $\Omega K(\pi, n+1)$ since the only nontrivial homotopy group of $\Omega K(\pi, n+1)$ is $\pi_n \approx \pi$. Here we need to assume π is abelian when $n = 1$ in order for $K(\pi, n+1)$ to exist. As a special case of the equivalence $K(\pi, n) \simeq \Omega K(\pi, n+1)$ we see that $\Omega \mathbb{C}\mathrm{P}^\infty \simeq S^1$, which shows that the loopspace of a space can have a much simpler homotopy type than the space itself. In the other direction, the loopspace of a simple space like S^n has a more complicated homotopy type, namely, we show in §4.4i that ΩS^n is homotopy equivalent to the James reduced product $J(S^{n-1})$ which was defined in §3.2.

Another way to find spaces homotopy equivalent to loopspaces is provided by the following:

Proposition 4.37. *If $F \to E \to B$ is a fibration or fiber bundle with E contractible, then there is a weak homotopy equivalence $F \to \Omega B$.*

Proof: If we compose a contraction of E with the projection $p : E \to B$ then we have for each point $x \in E$ a path in B from $p(x)$ to a basepoint. This yields a map $E \to E_p$ whose composition with the usual map $E_p \to B$ is p. By restriction this gives a map $F \to \Omega B$, and the long exact sequence of homotopy groups for $F \to E \to B$ maps to the long exact sequence for $\Omega B \to E_p \to B$. Since E and E_p are contractible, the five-lemma implies that the map $F \to \Omega B$ is a weak homotopy equivalence. □

Examples arising from fiber bundles constructed earlier in this section are $O(n) \simeq \Omega G_n(\mathbb{R}^\infty)$, $U(n) \simeq \Omega G_n(\mathbb{C}^\infty)$, and $Sp(n) \simeq \Omega G_n(\mathbb{H}^\infty)$. In particular, taking $n = 1$ in the latter two examples, we have $S^1 \simeq \Omega\mathbb{C}\mathrm{P}^\infty$ again and also $S^3 \simeq \Omega\mathbb{H}\mathrm{P}^\infty$. Note that in each of these examples it is a topological group which is homotopy equivalent to a loopspace. In [Milnor 56] this is shown to hold in general: For each topological group G there is a fiber bundle $G \to EG \to BG$ with EG contractible, hence by the proposition there is a weak equivalence $G \simeq \Omega BG$. There is also a converse statement: The loopspace of a CW complex is homotopy equivalent to a topological group.

The relationship between X and ΩX has been much studied, particularly the case that ΩX has the homotopy type of a finite CW complex, which is of special interest because of the examples of the classical Lie groups such as $O(n)$, $U(n)$, and $Sp(n)$. See [Kane] for an introduction to this subject.

Another reason loopspaces are of interest is their close connection with suspension. This starts with the observation that basepoint-preserving maps $X \to \Omega Y$ are the same as basepoint-preserving maps $\Sigma X \to Y$. Taking $X = S^n$, this amounts to the earlier observation that $\pi_n(\Omega Y) \approx \pi_{n+1}(Y)$. The more general case that X is an arbitrary CW complex plays an important role in §4.4e.

Let us return now to the general discussion of pathspace fibrations. If the fibration construction $f \mapsto E_f$ is applied to a map $p : E \to B$ which is already a fibration, one might expect the resulting fibration $E_p \to B$ to be in some way equivalent to the original $E \to B$. The natural notion of equivalence for fibrations is defined as follows. Given fibrations $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$, a map $f : E_1 \to E_2$ is called **fiber-preserving** if $p_1 = p_2 f$, or in other words, $f(p_1^{-1}(b)) \subset p_2^{-1}(b)$ for all $b \in B$. A fiber-preserving map $f : E_1 \to E_2$ is a **fiber homotopy equivalence** if there is a fiber-preserving map $g : E_2 \to E_1$ such that both compositions fg and gf are homotopic to the identity through fiber-preserving maps. Thus a fiber homotopy equivalence can be thought of as a family of homotopy equivalences between corresponding fibers of E_1 and E_2.

Proposition 4.38. *If $p : E \to B$ is a fibration, then the inclusion $E \hookrightarrow E_p$ is a fiber homotopy equivalence. In particular, the homotopy fiber of p is homotopy equivalent to the actual fiber $p^{-1}(b_0)$.*

Proof: We apply the homotopy lifting property to the homotopy $g_t : E_p \to B$, $g_t(e, \gamma) = \gamma(t)$, with initial lift $\widetilde{g}_0 : E_p \to E$, $\widetilde{g}_0(e, \gamma) = e$. The lifting $\widetilde{g}_t : E_p \to E$ is then the first coordinate of a homotopy $h_t : E_p \to E_p$ whose second coordinate is the restriction of the paths γ to the interval $[t, 1]$. Since the endpoints of the paths γ are unchanged, h_t is fiber-preserving. We have $h_0 = \mathbb{1}$, $h_1(E_p) \subset E$, and $h_t(E) \subset E$ for all t. If we let i denote the inclusion $E \hookrightarrow E_p$, then $ih_1 \simeq \mathbb{1}$ via h_t and $h_1 i \simeq \mathbb{1}$ via $h_t \big| E$, so i is a fiber homotopy equivalence. □

The process of turning maps into fibrations can be iterated in the following way. Given a fibration $p : E \to B$ with fiber $F = p^{-1}(b_0)$, we construct a diagram of the following form:

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & F_j & \longrightarrow & F_i & \xrightarrow{j} & F_p & \xrightarrow{i} & E & \xrightarrow{p} & B \\
 & & \Big\Vert \simeq & \nearrow & \Big\Vert \simeq & \nearrow & \uparrow \simeq & \nearrow & & & \\
\cdots & \longrightarrow & \Omega E & \xrightarrow[\Omega p]{} & \Omega B & & F & & & &
\end{array}$$

The homotopy fiber F_p of p consists of pairs (e, γ) with $e \in E$ and γ a path in B from $p(e)$ to b_0, and the map i is the projection $(e, \gamma) \mapsto e$. The inclusion $F \hookrightarrow F_p$ is a homotopy equivalence by the preceding proposition. The map $i : F_p \to E$ is evidently a fibration, and its fiber over a basepoint $e_0 \in F$ can be identified with ΩB. Its homotopy fiber F_i consists of pairs (γ, η) where η is a path in E ending at e_0 and γ is a path in B from $p(\eta(0))$ to b_0. The inclusion of the actual fiber ΩB into the homotopy fiber F_i is again a homotopy equivalence, and now there is a retraction $F_i \to \Omega B$ sending (γ, η) to the loop obtained by composing the inverse path of $p\gamma$ with η. Repeating these constructions for the fibration $j : F_i \to F_p$, produces the rest of the diagram.

Thus we produce a sequence

$$\cdots \to \Omega^2 B \to \Omega F \to \Omega E \to \Omega B \to F \to E \to B$$

where any two consecutive maps form a fibration, up to homotopy equivalence, and all the maps to the left of ΩB are obtained by applying the functor Ω to the later maps. The long exact sequence of homotopy groups for any fibration in the sequence coincides with the long exact sequence for $F \to E \to B$, as the reader can check.

Given a fibration $p : E \to B$ and a map $f : A \to B$, there is a **pullback** or **induced** fibration $f^*(E) \to A$ obtained by setting $f^*(E) = \{(a, e) \in A \times E \,\big|\, f(a) = p(e)\}$, with the

projections of $f^*(E)$ onto A and E giving a commutative diagram

$$\begin{array}{ccc} f^*(E) & \longrightarrow & E \\ \downarrow & & \downarrow p \\ A & \xrightarrow{f} & B \end{array}$$

The homotopy lifting property holds for $f^*(E) \to A$ since a homotopy $g_t : X \to A$ gives the first coordinate of a lift $\widetilde{g}_t : X \to f^*(E)$, the second coordinate being a lifting to E of the composed homotopy fg_t.

As a special case, if $E \to B$ is the path fibration $PB \to B$, then $f^*(PB)$ consists of the pairs $(a, \gamma) \in A \times PB$ with γ a path in B from the basepoint to $f(a)$, so $f^*(PB)$ is just the homotopy fiber of f.

Proposition 4.39. *Given a fibration $p : E \to B$ and a homotopy $f_t : A \to B$, the pullback fibrations $f_0^*(E) \to A$ and $f_1^*(E) \to A$ are fiber homotopy equivalent.*

Proof: Let $F : A \times I \to B$ be the homotopy f_t. The fibration $F^*(E) \to A \times I$ contains $f_0^*(E)$ and $f_1^*(E)$ over $A \times \{0\}$ and $A \times \{1\}$. So it suffices to prove the following: For a fibration $p : E \to B \times I$, the restricted fibrations $E_s = p^{-1}(B \times \{s\}) \to B$ are all fiber homotopy equivalent for $s \in [0, 1]$.

To prove this assertion the idea is to imitate the construction of the homotopy equivalences L_γ in the proof of Proposition 4.33. A path $\gamma : [0, 1] \to I$ gives rise to a fiber-preserving map $L_\gamma : E_{\gamma(0)} \to E_{\gamma(1)}$ by lifting the homotopy $g_t : E_{\gamma(0)} \to B \times I$, $g_t(x) = (p(x), \gamma(t))$, starting with the inclusion $E_{\gamma(0)} \hookrightarrow E$. As before, one shows the two basic properties (a) and (b), noting that in (a) the homotopy $L_\gamma \simeq L_{\gamma'}$ is fiber-preserving since it is obtained by lifting a homotopy $h_t : E_{\gamma(0)} \times [0, 1] \to B \times I$ of the form $h_t(x, u) = (p(x), -)$. From (a) and (b) it follows that L_γ is a fiber homotopy equivalence with inverse $L_{\overline{\gamma}}$. □

A consequence of the proposition is that a fibration $E \to B$ over a contractible base is fiber homotopy equivalent to a product fibration $B \times F \to B$, since the identity map $B \to B$ pulls back E to itself while a constant map $B \to B$ pulls back E to the product $B \times F$. More generally, the proposition implies that the set of fiber-homotopy equivalence classes of fibrations over a fixed base B depends only on the homotopy type of B.

Postnikov Towers

Using fibrations we shall now describe a general procedure for 'factoring' an arbitrary CW complex X into Eilenberg-MacLane spaces $K(\pi_n(X), n)$. Conversely, with the assumption that $\pi_1(X)$ acts trivially on the higher homotopy groups $\pi_n(X)$, we describe

the data specifying how to reconstruct X, up to homotopy equivalence, from its factors $K(\pi_n(X), n)$, i.e., from its homotopy groups.

A **Postnikov tower** for a path-connected space X is a commutative diagram

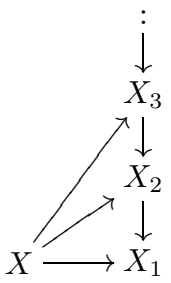

where:

(1) The map $X \to X_n$ induces an isomorphism on π_i for $i \leq n$.
(2) $\pi_i(X_n) = 0$ for $i > n$.
(3) The map $X_n \to X_{n-1}$ is a fibration with fiber a $K(\pi_n X, n)$.

In particular, X_1 is a $K(\pi_1 X, 1)$. The condition (3) is somewhat redundant since (1) and (2), together with commutativity of the diagram, imply that the fiber of $X_n \to X_{n-1}$ is a $K(\pi_n X, n)$, as is apparent from a glance at the long exact sequence of homotopy groups for this fibration. A map $X \to X_n$ satisfying (1) and (2) is called an n-approximation to X, according to the definition in §4.2.

Theorem 4.40. *Every connected CW complex X has a Postnikov tower in which the n-approximations X_n are homotopy equivalent to CW complexes. Such a Postnikov tower for X is unique up to homotopy equivalence.*

Proof: First choose CW n-approximations $X \to X_n$ for all n. These exist by Proposition 4.19 specialized to the map from X to a point. Applying Proposition 4.22 yields maps $X_n \to X_{n-1}$, unique up to homotopy, forming homotopy-commutative triangles as in the diagram of a Postnikov tower. We may convert each map $X_n \to X_{n-1}$ into a homotopy equivalent fibration, starting with $X_2 \to X_1$ and working upward. Then by successive applications of the homotopy lifting property we can arrange that all the triangles are strictly commutative, not just commutative up to homotopy. As already mentioned, (1) and (2) imply that the fiber of $X_n \to X_{n-1}$ is a $K(\pi_n X, n)$. □

For many purposes, a CW complex X can be replaced by an n-approximation, for example if one is interested in homotopy or homology groups in only a finite range of dimensions. However, to determine the full homotopy type of X from its Postnikov tower,

some sort of limit process is needed. Let us investigate this question is somewhat greater generality.

Given a sequence of maps $\cdots \to X_2 \to X_1$, define their **inverse limit** $\varprojlim X_n$ to be the subspace of the product $\prod_n X_n$ consisting of sequences of points $x_n \in X_n$ with x_n mapping to x_{n-1} under the map $X_n \to X_{n-1}$. The corresponding algebraic notion is the inverse limit $\varprojlim X_n$ of a sequence of group homomorphisms $\cdots \to G_2 \to G_1$, which is the subgroup of $\prod_n G_n$ consisting of sequences of elements $g_n \in G_n$ with g_n mapping to g_{n-1} under the homomorphism $G_n \to G_{n-1}$.

Proposition 4.41. *For a sequence of fibrations $\cdots \to X_2 \to X_1$ the natural map $\lambda : \pi_i(\varprojlim X_n) \to \varprojlim \pi_i(X_n)$ is surjective, and is injective if the maps $\pi_{i+1}(X_n) \to \pi_{i+1}(X_{n-1})$ are surjective for sufficiently large n.*

Proof: For surjectivity of λ, represent an element of $\varprojlim \pi_i(X_n)$ by maps $f_n : (S^i, s_0) \to (X_n, x_n)$. Since the projection $p_n : X_n \to X_{n-1}$ takes $[f_n]$ to $[f_{n-1}]$, by applying the homotopy lifting property for the pair (S^i, s_0) we can homotope f_n, fixing s_0, so that $p_n f_n = f_{n-1}$. Doing this inductively for $n = 2, 3, \cdots$, we get $p_n f_n = f_{n-1}$ for all n simultaneously, which gives surjectivity of λ.

For injectivity, note first that inverse limits are unaffected by throwing away a finite number of terms at the end of the sequence of spaces or groups, so we may assume the maps $\pi_{i+1}(X_n) \to \pi_{i+1}(X_{n-1})$ are surjective for all n. Given a map $f : S^i \to \varprojlim X_n$, suppose we have nullhomotopies $F_n : D^{i+1} \to X_n$ of the coordinate functions $f_n : S^i \to X_n$ of f. We have $p_n F_n = F_{n-1}$ on S^i, so $p_n F_n$ and F_{n-1} are the restrictions to the two hemispheres of S^{i+1} of a map $g_{n-1} : S^{i+1} \to X_{n-1}$. If $\pi_{i+1}(X_n) \to \pi_{i+1}(X_{n-1})$ is surjective, we can rechoose F_n so that the new g_{n-1} is nullhomotopic, i.e., so that $p_n F_n \simeq F_{n-1} \text{rel } S^i$. Applying the homotopy lifting property for (D^{i+1}, S^i), we can make $p_n F_n = F_{n-1}$. Doing this inductively for $n = 2, 3, \cdots$ we see that $f : S^i \to \varprojlim X_n$ is nullhomotopic and λ is injective. □

One might wish to have a description of the kernel of λ in the case of an arbitrary sequence of fibrations $\cdots \to X_2 \to X_1$, though for our present purposes this question is not relevant. It is not hard to produce a natural isomorphism of $Ker\lambda$ with $\varprojlim{}^1 \pi_{i+1}(X_n)$, where $\varprojlim{}^1$ is the functor defined in §3.4f. Namely, given $f : S^i \to \varprojlim X_n$ determining an element of $Ker\lambda$, then the sequence of maps $g_n : S^{i+1} \to X_n$ constructed above gives an element of $\prod_n \pi_{i+1}(X_n)$, well-defined up to the choice of the nullhomotopies F_n. Any new choice of F_n is obtained by adding a map $G_n : S^{i+1} \to X_n$ to F_n. The effect of this is to change g_n to $g_n + G_n$ and g_{n-1} to $g_{n-1} - p_n G_n$. Since $\varprojlim{}^1 \pi_{i+1}(X_n)$ is the quotient

of $\prod_n \pi_{i+1}(X_n)$ under exactly these identifications, we get $Ker\lambda \approx \varprojlim{}^1 \pi_{i+1}(X_n)$. Thus for each $i > 0$ there is a natural exact sequence

$$0 \longrightarrow \varprojlim{}^1 \pi_{i+1}(X_n) \longrightarrow \pi_i(\varprojlim X_n) \longrightarrow \varprojlim \pi_i(X_n) \longrightarrow 0$$

The proposition says that the $\varprojlim{}^1$ term vanishes if the maps $\pi_{i+1}(X_n) \to \pi_{i+1}(X_{n-1})$ are surjective for sufficiently large n.

Returning to the situation of a Postnikov tower, there is a natural map $X \to \varprojlim X_n$.

Corollary 4.42. *In a Postnikov tower for a connected CW complex X, the map $X \to \varprojlim X_n$ is a weak homotopy equivalence, so X is a CW approximation to $\varprojlim X_n$.*

Proof: The composition $\pi_i(X) \longrightarrow \pi_i(\varprojlim X_n) \xrightarrow{\lambda} \varprojlim \pi_i(X_n)$ is an isomorphism since $\pi_i(X) \to \pi_i(X_n)$ is an isomorphism for large n. □

Having seen how to decompose a space X into the terms in its Postnikov tower, we consider now the inverse process of building a Postnikov tower, starting with X_1 as a $K(\pi, 1)$ and inductively constructing X_n from X_{n-1}. It would be very nice if the fibration $K(\pi, n) \to X_n \to X_{n-1}$ could be extended another term to the right, to form a fibration sequence

$$K(\pi, n) \to X_n \to X_{n-1} \to K(\pi, n+1)$$

for this would say that X_n is the homotopy fiber of a map $X_{n-1} \to K(\pi, n+1)$, and homotopy classes of such maps are in one-to-one correspondence with elements of $H^{n+1}(X_{n-1}; \pi)$ by Theorem 4.17. Note that the last term $K(\pi, n+1)$ in the fibration sequence above cannot be anything else but a $K(\pi, n+1)$ since its loopspace must be homotopy equivalent to the first term in the sequence, a $K(\pi, n)$.

In general, a fibration $F \to E \to B$ is called **principal** if there is a commutative diagram

$$\begin{array}{ccccc} F & \longrightarrow & E & \longrightarrow & B \\ \downarrow & & \downarrow & & \downarrow \\ \Omega B' & \longrightarrow & F' & \longrightarrow & E' \longrightarrow B' \end{array}$$

where the second row is a fibration sequence and the vertical maps are weak homotopy equivalences.

Thus if all the fibrations in a Postnikov tower for X happen to be principal, we have

a diagram

$$\begin{array}{ccccc} & & \vdots & & \\ & & \downarrow & & \\ K(\pi_3X,3) & \longrightarrow & X_3 & \xrightarrow{k_3} & K(\pi_4X,5) \\ & & \downarrow & & \\ K(\pi_2X,2) & \longrightarrow & X_2 & \xrightarrow{k_2} & K(\pi_3X,4) \\ & & \downarrow & & \\ & & K(\pi_1X,1) = X_1 & \xrightarrow{k_1} & K(\pi_2X,3) \end{array}$$

where each X_{n+1} is, up to weak homotopy equivalence, the homotopy fiber of $k_n : X_n \to K(\pi_{n+1}, n+2)$. The map k_n is equivalent to a class in $H^{n+2}(X_n; \pi_{n+1}(K))$ called the n^{th} $\boldsymbol{k}$**-invariant** of X. These k-invariants specify how to construct X inductively from Eilenberg-MacLane spaces. For example, if all the k_n's are zero, X is just the product of the spaces $K(\pi_nX, n)$, and in the general case X can be viewed as a sort of 'twisted product' of $K(\pi_nX, n)$'s. To actually build a space X from its k-invariants is usually too unwieldy a procedure to be carried out in practice, but as a theoretical tool this procedure can be quite useful. The next result tells us when this tool is available:

Theorem 4.43. *A connected CW complex X has a Postnikov tower of principal fibrations iff $\pi_1(X)$ acts trivially on $\pi_n(X)$ for all $n > 1$.*

Notice that in the definition of a principal fibration, the map $F \to \Omega B'$ automatically exists and is a homotopy weak equivalence once one has the right-hand square of the commutative diagram with its vertical maps weak homotopy equivalences. Thus the question of whether a fibration is principal can be rephrased in the following way: Given a map $A \to X$, which one can always replace by an equivalent fibration if one likes, does there exist a fibration $F \to E \to B$ and a commutative square

$$\begin{array}{ccc} A & \longrightarrow & X \\ \downarrow & & \downarrow \\ F & \longrightarrow & E \end{array}$$

where the vertical maps are weak homotopy equivalences? By replacing A and X with CW approximations and converting the resulting map $A \to X$ into an inclusion via the usual mapping cylinder trick, the question becomes whether a CW pair (X, A) is equivalent to a fibration pair (E, F), i.e., whether there is a fibration $F \to E \to B$ and a map $(X, A) \to (E, F)$ for which both $X \to E$ and $A \to F$ are weak homotopy equivalences. In general the answer will rarely be yes, since the homotopy fiber of $A \hookrightarrow X$ would have to have the weak homotopy type of a loopspace, which is a rather severe restriction. However,

in the situation of Postnikov towers, the homotopy fiber is a $K(\pi, n)$ with π abelian since $n \geq 2$, so it is a loopspace. But there is another requirement: The action of $\pi_1(A)$ on $\pi_n(X, A)$ must be trivial for all $n \geq 1$. For this is equivalent to the action of $\pi_1(F)$ on $\pi_n(E, F)$ being trivial, which is always the case in a fibration since under the isomorphism $\pi_n(E, F) \approx \pi_n(B, x_0)$ this action corresponds to the action of $\pi_1(x_0)$ on $\pi_n(B, x_0)$.

The relative group $\pi_n(X, A)$ is always isomorphic to π_{n-1} of the homotopy fiber of the inclusion $A \hookrightarrow X$, so in the case at hand when the homotopy fiber is a $K(\pi, n)$, the only nontrivial relative homotopy group is $\pi_{n+1}(X, A) \approx \pi$. In this case the necessary condition of trivial action is also sufficient:

Lemma 4.44. *Let (X, A) be a CW pair with both X and A connected, such that the homotopy fiber of the inclusion $A \hookrightarrow X$ is a $K(\pi, n)$, $n \geq 1$. Then there exists a fibration $F \to E \to B$ and a map $(X, A) \to (E, F)$ inducing weak homotopy equivalences $X \to E$ and $A \to F$ iff the action of $\pi_1(A)$ on $\pi_{n+1}(X, A)$ is trivial.*

Proof: It remains only to prove the 'if' implication. As we noted just before the statement of the lemma, the groups $\pi_i(X, A)$ are zero except for $\pi_{n+1}(X, A) \approx \pi$. If the action of $\pi_1(A)$ on $\pi_{n+1}(X, A)$ is trivial, the relative Hurewicz theorem gives an isomorphism $\pi_{n+1}(X, A) \approx H_{n+1}(X, A)$. Since (X, A) is n-connected, we may assume A contains the n-skeleton of X, so X/A is n-connected and the absolute Hurewicz theorem gives $\pi_{n+1}(X/A) \approx H_{n+1}(X/A)$. Hence the quotient map $X \to X/A$ induces an isomorphism $\pi_{n+1}(X, A) \approx \pi_{n+1}(X/A)$ since the analogous statement for homology is certainly true.

Since $\pi_{n+1}(X/A) \approx \pi$, we can build a $K(\pi, n+1)$ from X/A by attaching cells of dimension $n+3$ and greater. This leads to a commutative diagram

$$\begin{array}{ccccc} A & \longrightarrow & X & \longrightarrow & X/A \\ \downarrow & & \downarrow \simeq & \overset{k}{\searrow} & \downarrow \\ F_k & \longrightarrow & E_k & \longrightarrow & K(\pi, n+1) \end{array}$$

where the vertical maps are inclusions and the lower row is obtained by converting the map k into a fibration. The map $A \to F_k$ is a weak homotopy equivalence by the five-lemma applied to the map between the long exact sequences of homotopy groups for the pairs (X, A) and (E_k, F_k), since the only nontrivial relative groups are π_{n+1}, both of which map isomorphically to $\pi_{n+1}(K(\pi, n+1))$. □

Proof of 4.43: In view of the lemma, all that needs to be done is identify the action of $\pi_1(X)$ on $\pi_n(X)$ with the action of $\pi_1(X_n)$ on $\pi_{n+1}(X_{n-1}, X_n)$ for $n \geq 2$, thinking of

the map $X_n \to X_{n-1}$ as an inclusion. From the exact sequence

$$0 = \pi_{n+1}(X_{n-1}) \longrightarrow \pi_{n+1}(X_{n-1}, X_n) \xrightarrow{\partial} \pi_n(X_n) \longrightarrow \pi_n(X_{n-1}) = 0$$

we have an isomorphism $\pi_{n+1}(X_{n-1}, X_n) \approx \pi_n(X_n)$ respecting the action of $\pi_1(X_n)$. And the map $X \to X_n$ induces isomorphisms on π_1 and π_n, so we are done. □

Let us describe now a natural generalization of Postnikov towers, in which one starts with a map $f : X \to Y$ between path-connected spaces rather than just a space X. A **Moore-Postnikov tower** for f is a commutative diagram

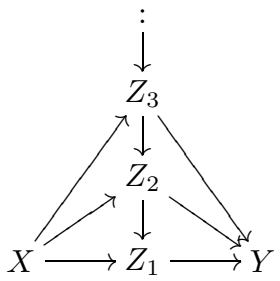

with each composition $X \to Z_n \to Y$ homotopic to f, and such that:

(1) The map $X \to Z_n$ induces an isomorphism on π_i for $i < n$ and a surjection for $i = n$.
(2) The map $Z_n \to Y$ induces an isomorphism on π_i for $i > n$ and an injection for $i = n$.
(3) The map $Z_{n+1} \to Z_n$ is a fibration with fiber a $K(\pi_n F, n)$ where F is the homotopy fiber of f.

A composition $X \to Z_n \to Y$ homotopic to f and satisfying (1) and (2) is what we called an n-factorization of f in §4.2. A Moore-Postnikov tower specializes to a Postnikov tower by taking Y to be a point and then setting $X_n = Z_{n+1}$, discarding the space Z_1 which has trivial homotopy groups.

Theorem 4.45. *Every map $f : X \to Y$ between connected CW complexes has a Moore-Postnikov tower, which is unique up to homotopy equivalence. A Moore-Postnikov tower of principal fibrations exists iff $\pi_1(X)$ acts trivially on $\pi_n(M_f, X)$ for all $n > 1$, where M_f is the mapping cylinder of f.*

Proof: The existence and uniqueness of a diagram satisfying (1) and (2) follows from Propositions 4.19 and 4.22, just as in the Postnikov case. Having this diagram, one again replaces each map $Z_n \to Z_{n-1}$ by a homotopy equivalent fibration, starting with $Z_2 \to Z_1$ and working upward, then one applies the homotopy lifting property to make all the triangles in the left half of the tower strictly commutative. To make the triangles in the

right half of the tower commute, simply replace each map to Y by the composition through Z_1.

It remains only to see that the fibers of the maps $Z_{n+1} \to Z_n$ are Eilenberg-MacLane spaces as in condition (3). Consider two successive levels of the tower:

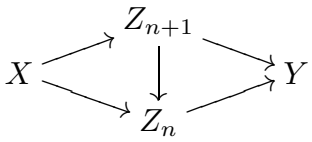

We may arrange that the maps $X \to Z_{n+1} \to Z_n \to Y$ are inclusions by taking mapping cylinders, first of $X \to Z_{n+1}$, then of the new $Z_{n+1} \to Z_n$, and then of the new $Z_n \to Y$. From the left-hand triangle we see that $Z_{n+1} \to Z_n$ induces an isomorphism on π_i for $i < n$ and a surjection for $i = n$, hence $\pi_i(Z_n, Z_{n+1}) = 0$ for $i < n+1$. Similarly, the other triangle gives $\pi_i(Z_n, Z_{n+1}) = 0$ for $i > n+1$. To show that $\pi_{n+1}(Z_n, Z_{n+1}) \approx \pi_{n+1}(Y, X)$ we use the following diagram:

$$\begin{array}{ccccccccc}
\pi_{n+1}(Z_{n+1}) & \longrightarrow & \pi_{n+1}(Z_n) & \longrightarrow & \pi_{n+1}(Z_n, Z_{n+1}) & \longrightarrow & \pi_n(Z_{n+1}) & \longrightarrow & \pi_n(Z_n) \\
\downarrow{=} & & \downarrow{\approx} & & \downarrow & & \downarrow{=} & & \downarrow \\
\pi_{n+1}(Z_{n+1}) & \longrightarrow & \pi_{n+1}(Y) & \longrightarrow & \pi_{n+1}(Y, Z_{n+1}) & \longrightarrow & \pi_n(Z_{n+1}) & \longrightarrow & \pi_n(Y) \\
\uparrow & & \uparrow{=} & & \uparrow & & \uparrow{\approx} & & \uparrow{=} \\
\pi_{n+1}(X) & \longrightarrow & \pi_{n+1}(Y) & \longrightarrow & \pi_{n+1}(Y, X) & \longrightarrow & \pi_n(X) & \longrightarrow & \pi_n(Y)
\end{array}$$

The upper-right vertical map is injective and the lower-left vertical map is surjective, so the five-lemma implies the two middle vertical maps are isomorphisms. Since the homotopy fiber of an inclusion $A \hookrightarrow B$ has π_i equal to $\pi_{i+1}(B, A)$, we see that condition (3) is satisfied.

To obtain the statement about a tower of principal fibrations we apply Lemma 4.44. As we saw in the previous paragraph, there are isomorphisms $\pi_{n+1}(Y, X) \approx \pi_{n+1}(Z_n, Z_{n+1})$, and these respect the action of $\pi_1(X) \approx \pi_1(Z_{n+1})$, so Lemma 4.44 gives the result. □

Just as in the Postnikov case, X is a CW approximation to $\varprojlim Z_n$.

Besides the case that Y is a point, which yields Postnikov towers, another interesting special case of Moore-Postnikov towers is when X is a point. In this case the space Z_n is an n-connected covering of Y, as in §4.2. As an example, the Hopf map $S^3 \to S^2$ exhibits S^3 as a 2-connected covering of S^2. This is one of the rare naturally-occurring n-connected coverings with $n > 1$.

The n-connected covering of Y can also be obtained as the homotopy fiber of an n-approximation $Y \to Y_n$. The tower of n-connected coverings of Y can be built by

iterating this construction:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & Z_2 & \longrightarrow & Z_1 & \longrightarrow & Y \\ & & \downarrow & & \downarrow & & \downarrow \\ & & K(\pi_3 Y, 3) & & K(\pi_2 Y, 2) & & K(\pi_1 Y, 1) \end{array}$$

Here Z_n is the homotopy fiber of the n-approximation $Z_{n-1} \to K(\pi_n Y, n)$. In particular this shows one has a tower of principal fibrations.

Obstruction Theory

Many questions in algebraic topology can be phrased in one of the following two forms:

The Extension Problem. Given a CW pair (W, A) and a map $A \to X$, does this extend to a map $W \to X$?

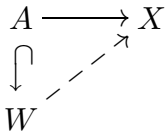

The Lifting Problem. Given a fibration $X \to Y$ and a map $W \to Y$, is there a lift $W \to X$?

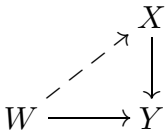

These two problems can be combined into a single problem:

The Relative Lifting Problem. Given a CW pair (W, A), a fibration $X \to Y$, and a map $W \to Y$, does there exist a lift $W \to X$ extending a given lift on A?

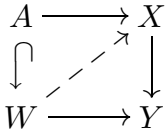

Besides reducing to the absolute lifting problem when $A = \varnothing$, this includes the extension problem by taking Y to be a point.

For example, the familiar homotopy extension property is a special case of the extension problem, and the homotopy lifting property is a special case of the relative lifting problem.

The term 'obstruction theory' refers to a procedure for defining a sequence of cohomology classes which are the obstructions to finding a solution to the extension, lifting, or relative lifting problem. In the most favorable cases these obstructions lie in cohomology groups which are all zero, so the problem has a solution. But even when the obstructions are nonzero it can be very useful to have the problem expressed in cohomological terms.

There are two ways of developing obstruction theory, which produce essentially the same result in the end. In the more elementary approach one tries to construct the extension or lifting one cell of W at a time, proceeding inductively over skeleta of W. Our

proof of Theorem 4.17 gives some flavor of this technique. This skeletal approach has an appealing directness, but the technical details of working at the level of cochains are perhaps a little tedious. Instead, we shall use a slightly more sophisticated approach in which the theory becomes an almost trivial application of Postnikov towers for the extension problem, and Moore-Postnikov towers for the lifting problems. (The skeletal viewpoint is explained in [VBKT], where it is useful in the study of characteristic classes of vector bundles.)

Let us consider the extension problem first, where we wish to extend a map $A \to X$ to the larger complex W. Suppose that X has a Postnikov tower of principal fibrations. Then we have a commutative diagram

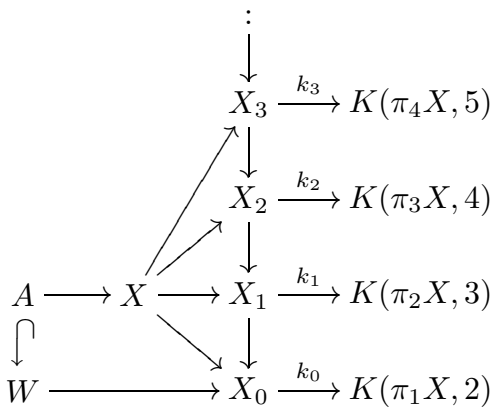

where we have enlarged the tower by adjoining the space X_0, which is just a point, at the bottom. The map $X_1 \to X_0$ is then a fibration, and to say it is principal says that X_1, which in any case is a $K(\pi_1 X, 1)$, is the loopspace of $K(\pi_1 X, 2)$, hence $\pi_1(X)$ must be abelian. Conversely, if $\pi_1(X)$ is abelian and acts trivially on all the higher homotopy groups of X, then there is an extended Postnikov tower of principal fibrations as shown.

Our strategy will be to try to lift the constant map $W \to X_0$ to maps $W \to X_n$ for $n = 1, 2, \cdots$ in succession, extending the given maps $A \to X_n$. If we succeed in finding all these lifts $W \to X_n$ it will then be easy to construct the desired extension $W \to X$.

For the inductive step we have a commutative diagram

$$\begin{array}{ccccc} A & \longrightarrow & X_n & \longrightarrow & PK \\ \downarrow & & \downarrow & & \downarrow \\ W & \longrightarrow & X_{n-1} & \longrightarrow & K = K(\pi_n X, n+1) \end{array}$$

Since X_n is the pullback, its points are pairs consisting of a point in X_{n-1} and a path from its image in K to the basepoint. A lift $W \to X_n$ therefore amounts to a nullhomotopy

of the composition $W \to X_{n-1} \to K$. We already have such a lift defined on A, hence a nullhomotopy of $A \to K$, and we want a nullhomotopy of $W \to K$ extending this nullhomotopy on A.

The map $W \to K$ together with the nullhomotopy on A gives a map $W \cup CA \to K$, where CA is the cone on A. Since K is a $K(\pi_n X, n+1)$, the map $W \cup CA \to K$ determines an **obstruction** class $\omega_n \in H^{n+1}(W \cup CA; \pi_n X) \approx H^{n+1}(W, A; \pi_n X)$.

Proposition 4.46. *A lift $W \to X_n$ extending the given $A \to X_n$ exists iff $\omega_n = 0$.*

Proof: We need to show that the map $W \cup CA \to K$ extends to a map $CW \to K$ iff $\omega_n = 0$, i.e., $W \cup CA \to K$ is homotopic to a constant map.

Suppose that $g_t : W \cup CA \to K$ is such a homotopy. The constant map g_1 then extends to the constant map $g_1 : CW \to K$, so by the homotopy extension property for the pair $(CW, W \cup CA)$, applied to the reversed homotopy g_{1-t}, we have a homotopy $g_t : CW \to K$ extending the previous homotopy $g_t : W \cup CA \to K$. The map $g_0 : CW \to K$ then extends the given map $W \cup CA \to K$.

Conversely, if we have an extension $CW \to K$, then this is nullhomotopic since the cone CW is contractible, and we may restrict such a nullhomotopy to $W \cup CA$. □

If we succeed in extending the lifts $A \to X_n$ to lifts $W \to X_n$ for all n, then we obtain a map $W \to \varprojlim X_n$ extending the given $A \to X \to \varprojlim X_n$. Let M be the mapping cylinder of $X \to \varprojlim X_n$. Since the restriction of $W \to \varprojlim X_n \subset M$ to A factors through X, this gives a homotopy of this restriction to the map $A \to X \subset M$. Extend this to a homotopy of $W \to M$, producing a map $(W, A) \to (M, X)$. Since the map $X \to \varprojlim X_n$ is a weak homotopy equivalence, $\pi_i(M, X) = 0$ for all i, so by Lemma 4.5 the map $(W, A) \to (M, X)$ can be homotoped to a map $W \to X$ extending the given $A \to X$, and we have solved the extension problem.

Thus if it happens that at each stage of the inductive process of constructing lifts $W \to X_n$ the obstruction $\omega_n \in H^{n+1}(W, A; \pi_n X)$ vanishes, then the extension problem has a solution. In particular, this yields:

Corollary 4.47. *If X is a connected abelian CW complex and (W, A) is a CW pair such that $H^{n+1}(W, A; \pi_n X) = 0$ for all n, then every map $A \to X$ can be extended to a map $W \to X$.* □

This is a considerable improvement on the more elementary result that extensions exist if $\pi_n(X) = 0$ for all n such that $W - A$ has cells of dimension $n+1$, which is an exercise for §4.1.

We can apply the Hurewicz theorem and obstruction theory to extend the homology version of Whitehead's theorem to CW complexes with trivial action of π_1 on all homotopy groups:

Proposition 4.48. *If X and Y are connected abelian CW complexes, then a map $f: X \to Y$ which induces isomorphisms on all homology groups is a homotopy equivalence.*

Proof: Taking the mapping cylinder of f reduces us to the case of an inclusion $X \hookrightarrow Y$ of a subcomplex. The key point then is to see that $\pi_1(X)$ acts trivially on $\pi_n(Y, X)$ for all n, since if this is true, then the relative Hurewicz theorem will imply that $\pi_n(Y, X) = 0$ for all n, so $X \to Y$ will be a weak homotopy equivalence. The assumptions guarantee that $\pi_1(X) \to \pi_1(Y)$ is an isomorphism, so we know at least that $\pi_1(Y, X) = 0$.

We use obstruction theory to construct a retraction $Y \to X$, i.e., to extend the identity map $X \to X$ to a map $Y \to X$. To apply the theory we need $\pi_1(X)$ acting trivially on $\pi_n(X)$, which holds by hypothesis. Since the inclusion $X \hookrightarrow Y$ induces isomorphisms on homology, we have $H_*(Y, X) = 0$, hence $H^{n+1}(Y, X; \pi_n(X)) = 0$ for all n by the universal coefficient theorem. So there are no obstructions, and a retraction $Y \to X$ exists. This implies that the maps $\pi_n(Y) \to \pi_n(Y, X)$ are onto, so trivial action of $\pi_1(X)$ on $\pi_n(Y)$ implies trivial action on $\pi_n(Y, X)$ by naturality of the action. □

The generalization of the preceding analysis of the extension problem to the relative lifting problem is straightforward. Assuming the fibration $p: X \to Y$ in the statement of the lifting problem has a Moore-Postnikov tower of principal fibrations, we have a diagram

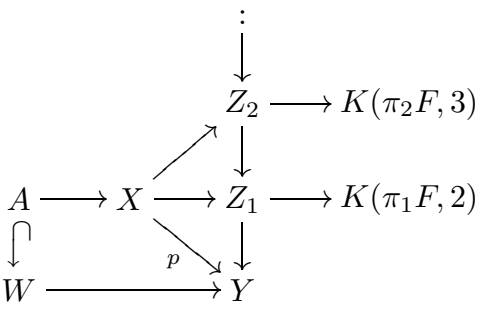

The first step is to lift the map $W \to Y$ to Z_1, extending the given lift on A. We may take Z_1 to be the covering space of Y corresponding to the subgroup $p_*(\pi_1(X))$ of $\pi_1(Y)$, so covering space theory tells us when we can lift $W \to Y$ to Z_1, and the unique lifting property for covering spaces can be used to see whether a lift can be chosen to agree with the lift on A given by the diagram; this could only be a problem when A has more than one component.

Having a lift to Z_1, the analysis proceeds exactly as before. One encounters a sequence of obstructions $\omega_n \in H^{n+1}(W, A; \pi_n F)$, assuming $\pi_1 F$ is abelian in the case $n = 1$, and a lift to X exists, extending the given lift on A, if each successive ω_n is zero.

One can ask the converse question: If a lift exists, must the obstructions ω_n all be zero? Since Proposition 4.46 is an 'if and only if' statement, one might expect the answer to be yes, but upon closer inspection the matter becomes less clear. The difficulty is that, even if at some stage the obstruction ω_n is zero, so a lift to Z_{n+1} exists, there may be many choices of such a lift, and different choices could lead to different ω_{n+1}'s, some zero and others nonzero. Examples of such ambiguities are not hard to produce, both for the lifting and the extension problem, and the ambiguities only become worse with each subsequent choice of a lift. So it is only in rather special circumstances that one can say that there are well-defined obstructions. A simple case is when $\pi_i(F) = 0$ for $i < n$, so the Moore-Postnikov factorization begins with Z_n:

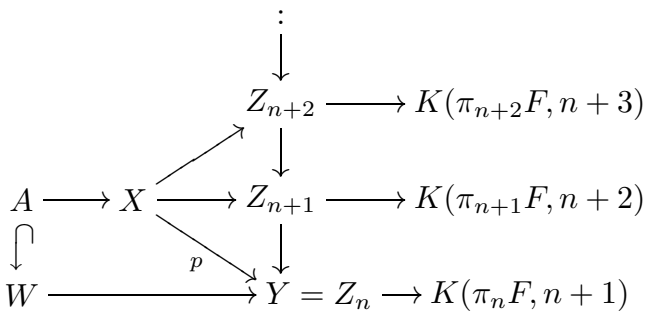

Hence there is a well-defined **primary obstruction** $\omega_n \in H^{n+1}(W, A; \pi_n F)$.

Exercises

1. For a fiber bundle $F \to E \to B$ such that the inclusion $F \hookrightarrow E$ is homotopic to a constant map, show that the long exact sequence of homotopy groups breaks up into split short exact sequences, giving isomorphisms $\pi_n(B) \approx \pi_n(E) \oplus \pi_{n-1}(F)$. In particular, for the Hopf bundles $S^3 \to S^7 \to S^4$ and $S^7 \to S^{15} \to S^8$ this yields isomorphisms

$$\pi_n(S^4) \approx \pi_n(S^7) \oplus \pi_{n-1}(S^3) \qquad \pi_n(S^8) \approx \pi_n(S^{15}) \oplus \pi_{n-1}(S^7)$$

for all n.

2. Show that if $F \to E \to B$ is a fiber bundle with F and B manifolds of dimensions m and n, then E is a manifold of dimension $m + n$.

3. Show that if $S^k \to S^m \to S^n$ is a fiber bundle, then $k = n - 1$ and $m = 2n - 1$. [Use the previous problem and the long exact sequence of homotopy groups.]

4. Show that if there were fiber bundles $S^{n-1} \to S^{2n-1} \to S^n$ for all n, then the groups $\pi_i(S^n)$ would be finitely generated free abelian groups computable by induction, and nonzero for $i \geq n \geq 2$.

5. Let $p : S^3 \to S^2$ be the Hopf bundle and let $q : T^3 \to S^3$ be the quotient map which collapses the complement of a ball in the 3-dimensional torus $T^3 = S^1 \times S^1 \times S^1$ to a point. Show that $pq : T^3 \to S^2$ induces the trivial map on π_* and $\widetilde{H}_*$, yet is not homotopic to a constant map.

6. Show there is a map $\mathbb{R}\mathrm{P}^\infty \to \mathbb{C}\mathrm{P}^\infty = K(\mathbb{Z}, 2)$ which induces the trivial map on $\widetilde{H}_*(-;\mathbb{Z})$ but a nontrivial map on $\widetilde{H}^*(-;\mathbb{Z})$. How is this consistent with the universal coefficient theorem?

7. For basepoint-preserving maps $f : S^1 \to X$ and $g : S^n \to X$ with $n > 1$, show that the Whitehead product $[f, g]$ is $\pm(g - fg)$, where fg denotes the action of f on g.

8. Show that a map $p : E \to B$ is a fibration iff the map $\pi : E^I \to E_p$, $\pi(\gamma) = (\gamma(0), p\gamma)$, has a section, i.e., a map $s : E_p \to E^I$ such that $ps = \mathbb{1}$.

9. Show that a linear projection of a 2-simplex onto one of its edges is a fibration but not a fiber bundle.

10. Given a fibration $F \to E \to B$, use the homotopy lifting property to define an action of $\pi_1(E)$ on $\pi_n(F)$, i.e., a homomorphism $\pi_1(E) \to Aut\big(\pi_n(F)\big)$, such that the composition $\pi_1(F) \to \pi_1(E) \to Aut\big(\pi_n(F)\big)$ is the usual action of $\pi_1(F)$ on $\pi_n(F)$. Deduce that if $\pi_1(E) = 0$, then the action of $\pi_1(F)$ on $\pi_n(F)$ is trivial.

11. Show that if $p : E \to B$ is a fibration, then its fiber F is homotopy equivalent to the pullback $p^*(PB)$ of the path space PB in the diagram

$$\begin{array}{ccc} p^*(PB) & \longrightarrow & PB \\ \downarrow & & \downarrow \\ E & \xrightarrow{\ p\ } & B \end{array}$$

12. Show that for homotopic maps $f, g : A \to B$ the fibrations $E_f \to B$ and $E_g \to B$ are fiber homotopy equivalent.

13. Given a map $f : A \to B$ and a deformation retraction in the weak sense of A onto a subspace A', with $i : A' \to A$ the inclusion, show that the fibrations $E_f \to B$ and $E_{fi} \to B$ are fiber homotopy equivalent.

14. If the fibration $p : E \to B$ is a homotopic equivalence, show that p is a fiber homotopy equivalence of E with the trivial fibration $\mathbb{1} : B \to B$.

15. Show that a map $f: X \to Y$ of connected CW complexes is a homotopy equivalence if it induces an isomorphism on π_1 and its homotopy fiber F_f has $\widetilde{H}_*(F_f; \mathbb{Z}) = 0$.

16. Show that ΩX is an H-space with the multiplication given by composition of loops.

17. Show that a fibration sequence $\cdots \to \Omega B \to F \to E \to B$ induces a long exact sequence $\cdots \to \langle X, \Omega B\rangle \to \langle X, F\rangle \to \langle X, E\rangle \to \langle X, B\rangle$, with groups and group homomorphisms except for the last three terms, and abelian groups except for the last six terms.

18. Given a fibration $F \longrightarrow E \xrightarrow{p} B$, define a natural action of ΩB on the homotopy fiber F_p and use this to show that exactness at $\langle X, F\rangle$ in the long exact sequence in the preceding problem can be improved to the statement that two elements of $\langle X, F\rangle$ have the same image in $\langle X, E\rangle$ iff they are in the same orbit of the induced action of $\langle X, \Omega B\rangle$ on $\langle X, F\rangle$.

19. For a space B, let $\mathcal{F}(B)$ denote the set of fiber homotopy equivalence classes of fibrations $E \to B$. Let $\mathcal{S}(B)$ denote the set of equivalence classes of maps $f: A \to B$ where $f_1: A_1 \to B$ is equivalent to $f_2: A_2 \to B$ if there is a homotopy equivalence $g: A_1 \to A_2$ such that $f_1 \simeq f_2 g$. Show the natural map $\mathcal{F}(B) \to \mathcal{S}(B)$ is a bijection, and show that a homotopy equivalence $h: B_1 \to B_2$ induces bijections $\mathcal{S}(B_1) \to \mathcal{S}(B_2)$ and $\mathcal{F}(B_1) \to \mathcal{F}(B_2)$.

20. Show that a map $f: X \to Y$ between connected CW complexes factors as a composition $X \to Z_n \to Y$ where the first map induces isomorphisms on π_i for $i \le n$ and the second map induces isomorphisms on π_i for $i \ge n+1$. [Does this data determine the homotopy type of Z_n?]

21. Show that by applying the loopspace functor to a Postnikov tower for X one obtains a Postnikov tower of principal fibrations for ΩX.

22. Show that in the Postnikov tower of an H-space, all the spaces are H-spaces and the maps are H-maps, commuting with the multiplication, up to homotopy.

23. Show that a principal fibration $\Omega C \longrightarrow E \xrightarrow{p} B$ has a section (i.e., a map $s: B \to E$ with $ps = \mathbb{1}$) iff E is fiber homotopy equivalent to the product fibration $\Omega C \times B$.

24. Prove the following uniqueness result for the Quillen plus construction: Given a connected CW complex X, if there is an abelian CW complex Y and a map $X \to Y$ inducing an isomorphism $H_*(X; \mathbb{Z}) \approx H_*(Y; \mathbb{Z})$, then such a Y is unique up to homotopy equivalence. [Apply Corollary 4.47 with W the mapping cylinder of one such map $X \to Y$.]

25. In the situation of the relative lifting problem, suppose one has two different lifts $W \to X$ which agree on the subspace $A \subset W$. Show that the obstructions to finding a homotopy rel A between these two lifts lie in the groups $H^n(W, A; \pi_n F)$.

Appendix

1. Topology of Cell Complexes

Here we collect a number of basic topological facts about CW complexes for convenient reference.

Let us first recall from Chapter 0 that a CW complex is a space X constructed in the following way:

(1) Start with a discrete set X^0, the 0-cells of X.
(2) Inductively, form the n-skeleton X^n from X^{n-1} by attaching n-cells e^n_α via maps $\varphi_\alpha : S^{n-1} \to X^{n-1}$. This means that X^n is the quotient space of $X^{n-1} \coprod_\alpha D^n_\alpha$ under the identifications $x \sim \varphi_\alpha(x)$ for $x \in \partial D^n_\alpha$.
(3) $X = \bigcup_n X^n$ with the weak topology: A set $A \subset X$ is open (or closed) iff $A \cap X^n$ is open (or closed) in X^n for each n.

Note that condition (3) is superfluous when X is finite-dimensional, say $X = X^n$. For if A is open in $X = X^n$, the definition of the quotient topology on X^n implies that $A \cap X^{n-1}$ is open in X^{n-1}, and then by the same reasoning $A \cap X^{n-2}$ is open in X^{n-2}, etc.

Each cell e^n_α has its **characteristic map** Φ_α, the composition $D^n_\alpha \hookrightarrow X^{n-1} \coprod_\alpha D^n_\alpha \to X^n \hookrightarrow X$. This is continuous since it is a composition of continuous maps, the inclusion $X^n \hookrightarrow X$ being continuous by (3). The restriction of Φ_α to the interior of D^n_α is a homeomorphism onto e^n_α.

An alternative way to describe the topology on X is to say that a set $A \subset X$ is open (or closed) iff $\Phi_\alpha^{-1}(A)$ is open (or closed) in D^n_α for each characteristic map Φ_α. In one direction this follows from continuity of the Φ_α's, and in the other direction, suppose $\Phi_\alpha^{-1}(A)$ is open in D^n_α for each Φ_α, and suppose by induction on n that $A \cap X^{n-1}$ is open in X^{n-1}. Then since $\Phi_\alpha^{-1}(A)$ is open in D^n_α for all α, $A \cap X^n$ is open in X^n by the definition of the quotient topology on X^n. Hence by (3), A is open in X.

A consequence of this characterization of the topology on X is that X is a quotient space of $\coprod_{n,\alpha} D^n_\alpha$.

A **subcomplex** of a CW complex X is a subspace $A \subset X$ which is a union of cells of X, such that the closure of each cell in A is also contained in A. Thus for each cell in A, the image of its attaching map is contained in A, so A is itself a CW complex. Its CW complex topology is the same as the topology induced from X, as one sees by noting inductively that the two topologies agree on $A^n = A \cap X^n$. It is easy to see by induction over skeleta that a subcomplex is a closed subspace. Conversely, a subcomplex could be defined as a closed subspace which is a union of cells.

A finite CW complex, i.e., one with only finitely many cells, is compact since attaching a single cell preserves compactness. A sort of converse to this is:

Proposition A.1. *A compact subspace of a CW complex is contained in a finite subcomplex.*

Proof: First we show that a compact set $C \subset X$ can meet only finitely many cells e^n_α of X. Suppose on the contrary that there is an infinite sequence of points $x_i \in C$ all lying in distinct cells e^n_α. Then the set $S = \{x_1, x_2, \cdots\}$ is closed in X. Namely, assuming $S \cap X^{n-1}$ is closed in X^{n-1} by induction on n, then for each cell e^n_α of X, $\varphi_\alpha^{-1}(S)$ is closed in ∂D^n_α, and $\Phi_\alpha^{-1}(S)$ consists of at most one more point in D^n_α, hence is closed in D^n_α. Therefore $S \cap X^n$ is closed in X^n for each n, hence S is closed in X. The same argument shows that any subset of S is closed, so S has the discrete topology. But it is compact, being a closed subset of the compact set C. Therefore S must be finite, a contradiction.

Since C is contained in a finite union of cells, it suffices to show that a finite union of cells is contained in a finite subcomplex of X. A finite union of finite subcomplexes is again a finite subcomplex, so this reduces to showing that a single cell e^n_α is contained in a finite subcomplex. The image of the attaching map φ_α for e^n_α is compact, hence by induction on dimension is contained in a finite subcomplex $A \subset X^{n-1}$. So e^n_α is contained in the finite subcomplex $A \cup e^n_\alpha$. □

The letters 'CW' refer to two properties satisfied by CW complexes:

(1) Closure-finiteness: the closure of each cell meets only finitely many other cells. This follows from the preceding proposition since the closure of a cell is compact, being the image of a characteristic map.

(2) Weak topology: a set is closed iff it meets the closure of each cell in a closed set. For if a set meets the closure of each cell in a closed set, it pulls back to a closed set under each characteristic map, hence is closed by an earlier remark.

In J.H.C.Whitehead's original definition of CW complexes these two properties played a more central role. However, the inductive construction we have given has subsequently become the standard definition.

Next we describe a convenient way of constructing open neighborhoods $N_\varepsilon(A)$ of subsets A of a CW complex X. The construction is inductive over the skeleta X^n, so suppose we have already constructed $N^n_\varepsilon(A)$, a neighborhood of $A \cap X^n$ in X^n, starting the process with $N^0_\varepsilon(A) = A \cap X^0$. Then we define $N^{n+1}_\varepsilon(A)$ by specifying its preimage under the characteristic map $\Phi_\alpha : D^{n+1} \to X$ of each cell e^{n+1}_α, namely, $\Phi_\alpha^{-1}\big(N^{n+1}_\varepsilon(A)\big)$

is the union of two parts: an open ε-neighborhood of $\Phi_\alpha^{-1}(A) - \partial D^{n+1}$ in $D^{n+1} - \partial D^{n+1}$, and a product $(1-\varepsilon, 1] \times \Phi_\alpha^{-1}\big(N_\varepsilon^n(A)\big)$ with respect to 'spherical' coordinates (r, θ) in D^{n+1}, where $r \in [0,1]$ is the radial coordinate and $\theta \in \partial D^{n+1} = S^n$. The number $\varepsilon = \varepsilon_\alpha > 0$ we allow to depend on the cell e_α^{n+1}. Then we define $N_\varepsilon(A) = \bigcup_n N_\varepsilon^n(A)$. This is an open set in X since it pulls back to an open set under each characteristic map.

Proposition A.2. *CW complexes are normal, and in particular, Hausdorff.*

Proof: Points are closed in a CW complex X since they pull back to closed sets under all characteristic maps Φ_α. For disjoint closed sets A and B in X, we show that $N_\varepsilon(A)$ and $N_\varepsilon(B)$ are disjoint for small enough ε_α's. In the inductive process for building these open sets, assume $N_\varepsilon^n(A)$ and $N_\varepsilon^n(B)$ have been chosen to be disjoint. For a characteristic map $\Phi_\alpha : D^{n+1} \to X$, observe that $\Phi_\alpha^{-1}\big(N_\varepsilon^n(A)\big)$ and $\Phi_\alpha^{-1}(B)$ are a positive distance apart, otherwise using compactness we would have a sequence in $\Phi_\alpha^{-1}(B)$ converging to a point of $\Phi_\alpha^{-1}(B)$ in ∂D^{n+1} of distance zero from $\Phi_\alpha^{-1}\big(N_\varepsilon^n(A)\big)$, but this is impossible since $\Phi_\alpha^{-1}\big(N_\varepsilon^n(B)\big)$ is a neighborhood of $\Phi_\alpha^{-1}(B)$ in ∂D^{n+1} disjoint from $\Phi_\alpha^{-1}\big(N_\varepsilon^n(A)\big)$. Similarly, $\Phi_\alpha^{-1}\big(N_\varepsilon^n(B)\big)$ and $\Phi_\alpha^{-1}(A)$ are a positive distance apart. Also, $\Phi_\alpha^{-1}(A)$ and $\Phi_\alpha^{-1}(B)$ are a positive distance apart. So a small enough ε_α will make $\Phi_\alpha^{-1}\big(N_\varepsilon^{n+1}(A)\big)$ disjoint from $\Phi_\alpha^{-1}\big(N_\varepsilon^{n+1}(B)\big)$ in D^{n+1}. □

Proposition A.3. *A CW complexes is locally contractible, in the strong sense that each point has arbitrarily small contractible open neighborhoods.*

Proof: Given $x \in X$ and a neighborhood U of x in X, we can choose the ε_α's small enough so that $N_\varepsilon(x) \subset U$ by requiring that the closure of $N_\varepsilon^n(x)$ be contained in U for each n. It remains to see that $N_\varepsilon(x)$ is contractible. If $x \in X^m - X^{m-1}$ and $n > m$ we can construct a deformation retraction of $N_\varepsilon^n(x)$ onto $N_\varepsilon^{n-1}(x)$ by sliding out 'radial segments' in cells e_β^n, i.e., the images under the characteristic maps Φ_β of radial segments in D^n. A deformation retraction of $N_\varepsilon(x)$ onto $N_\varepsilon^m(x)$ is then obtained by performing the deformation retraction of $N_\varepsilon^n(x)$ onto $N_\varepsilon^{n-1}(x)$ during the t-interval $[1/2^n, 1/2^{n-1}]$, points of $N_\varepsilon^n(x) - N_\varepsilon^{n-1}(x)$ being stationary outside this t-interval. Finally, $N_\varepsilon^m(x)$ is an open ball about x, and so deformation retracts onto x. □

In particular, CW complexes are locally path-connected. So a CW complex is path-connected iff it is connected.

Proposition A.4. *For a subcomplex A of a CW complex X, the open neighborhoods $N_\varepsilon(A)$ deformation retract onto A if $\varepsilon_\alpha < 1$ for all α.*

Proof: In each cell of $X - A$, $N_\varepsilon(A)$ is a product neighborhood of the boundary of this cell, so a deformation retraction of $N_\varepsilon(A)$ onto A can be constructed just as in the previous proof. □

Note that for subcomplexes A and B of X, we have $N_\varepsilon(A) \cap N_\varepsilon(B) = N_\varepsilon(A \cap B)$. This implies for example that the van Kampen theorem and Mayer-Vietoris sequences hold for decompositions $X = A \cup B$ into subcomplexes A and B as well as into open sets A and B.

A map $f : X \to Y$ with domain a CW complex is continuous iff its restrictions to the closures $\overline{e}^n_\alpha$ of all cells e^n_α are continuous, and it is useful to know that the same is true for homotopies $f_t : X \to Y$. With this objective in mind, let us introduce a little terminology. A topological space X is said to be **generated** by a collection of subspaces X_α if $X = \bigcup_\alpha X_\alpha$ and a set $A \subset X$ is closed iff $A \cap X_\alpha$ is closed in X_α for each α. Equivalently, we could say 'open' instead of 'closed' here, but 'closed' is more convenient for our present purposes. As noted earlier, though not in these words, a CW complex X is generated by the closures $\overline{e}^n_\alpha$ of its cells e^n_α. Since every finite subcomplex of X is a finite union of closures $\overline{e}^n_\alpha$, X is also generated by its finite subcomplexes. It follows that X is also generated by its compact subspaces, or more briefly, X is **compactly generated**.

Proposition A.14 later in the Appendix asserts that if X is a compactly generated Hausdorff space and Z is locally compact Hausdorff, then $X \times Z$, with the product topology, is compactly generated. In particular, $X \times I$ is compactly generated if X is a CW complex. Since every compact set in $X \times I$ is contained in the product of a compact subspace of X with I, hence in the product of a finite subcomplex of X with I, such product subspaces also generate $X \times I$. Since such a product subspace is a finite union of products $\overline{e}^n_\alpha \times I$, it is also true that $X \times I$ is generated by its subspaces $\overline{e}^n_\alpha \times I$. This implies that a homotopy $F : X \times I \to Y$ is continuous iff its restrictions to the subspaces $\overline{e}^n_\alpha \times I$ are continuous, which is the statement we were seeking.

Products of CW Complexes

There are some unexpected point-set-topological subtleties which arise with products of CW complexes. As we shall show, the product of two CW complexes does have a natural CW structure, but its topology is in general finer, with more open sets, than the product topology. However, the distinctions between the two topologies are rather small, and indeed nonexistent in most cases of interest, so there is no real problem for algebraic topology.

Given a space X and a collection of subspaces X_α whose union is X, these subspaces generate a finer topology on X by defining a set $A \subset X$ to be open iff $A \cap X_\alpha$ is open in X_α for all X_α's. The axioms for a topology are easily verified for this definition. In case $\{X_\alpha\}$ is the collection of compact subsets of X, we write X_c for this new compactly generated topology. It is easy to see that X and X_c have the same compact subsets, and the two induced topologies on these compact subsets coincide. If X is compact, or even locally compact, then $X = X_c$, i.e., X is compactly generated.

Theorem A.5. *For CW complexes X and Y with characteristic maps Φ_α and Ψ_β, the product maps $\Phi_\alpha \times \Psi_\beta$ are the characteristic maps for a CW complex structure on $(X \times Y)_c$. If either X or Y is compact or more generally locally compact, then $(X \times Y)_c = X \times Y$. Also, $(X \times Y)_c = X \times Y$ if X and Y have at most countably many cells.*

The simplest example where the two topologies $X \times Y$ and $(X \times Y)_c$ differ has X and Y graphs consisting of infinitely many edges emanating from a single vertex, with countably many edges for X and uncountably many for Y; see [Dowker].

Proof: When X and Y are finite complexes, hence $X \times Y$ is compact and $(X \times Y)_c = X \times Y$, the proof is easy: Let $(X \times Y)^n = \bigcup_{i+j=n} X^i \times Y^j$, the union of the cells $e^k_\alpha \times e^\ell_\beta$ with $k + \ell \leq n$. Assume inductively that in the topology induced from the product topology on $X \times Y$, $(X \times Y)^{n-1}$ is a CW complex with characteristic maps $\Phi_\alpha \times \Psi_\beta$. The characteristic maps $\Phi_\alpha \times \Psi_\beta$ for the products $e^i_\alpha \times e^j_\beta$ with $i + j = n$ give a surjective map $(X \times Y)^{n-1} \coprod_{i+j=n} D^i_\alpha \times D^j_\beta \to (X \times Y)^n$. The domain of this map is compact and the range is Hausdorff, so it is a quotient map. Hence $(X \times Y)^n$ is a CW complex in the topology induced from the product topology on $X \times Y$, completing the induction step.

In the general case we have the following situation: The space $Z = (X \times Y)_c$ is generated by a collection of finite CW complexes $Z_\gamma \subset Z$, the products of finite subcomplexes of X and Y, such that any two Z_γ's are subcomplexes of a third Z_γ. Then we assert that Z has a unique CW complex structure with all the Z_γ's as subcomplexes. This is because the condition on the Z_γ's guarantees that Z has a well-defined decomposition into cells with well-defined characteristic maps, and so when we build Z up skeleton by skeleton, the step from Z^{n-1} to Z^n is accomplished by attaching n-cells, since Z has compactly generated topology.

The second assertion of the theorem is a special case of Proposition A.14, having nothing to do with CW complexes, which says that a product $X \times Y$ is compactly generated if X is compactly generated and Y is locally compact.

For the final statement of the theorem, suppose X and Y each have at most countably

many cells. Then there are finite subcomplexes $X_1 \subset X_2 \subset \cdots$ of X with $X = \bigcup_i X_i$, and similarly for Y. For an open set $W \subset (X \times Y)_c$ and a point $(a, b) \in W$ we need to find a product $U \times V \subset W$ with U an open neighborhood of a in X and V an open neighborhood of b in Y. Suppose inductively that we have a compact product $K_i \times L_i \subset W$ which is a neighborhood of (a, b) in $X_i \times Y_i$. We know such a product $K_i \times L_i$ exists since the two topologies on $X_i \times Y_i$ are the same. We would like to find a compact product neighborhood $K_{i+1} \times L_{i+1} \subset W$ of $K_i \times L_i$ in $X_{i+1} \times Y_{i+1}$. To do this, we first choose for each $x \in K_i$ compact neighborhoods K_x of x in X_{i+1} and L_x of L_i in Y_{i+1} such that $K_x \times L_x \subset W$, using the compactness of L_i. By compactness of K_i, a finite number of the K_x's cover K_i. Let K_{i+1} be the union of these K_x's and let L_{i+1} be the intersection of the corresponding L_x's. This defines the desired $K_{i+1} \times L_{i+1}$. Let U_i be the interior of K_i in X_i, hence $U_i \subset U_{i+1}$ for each i. The union $U = \bigcup_i U_i$ is then open in X since it intersects each X_i in a union of open sets and the X_i's generate X. Similarly, the L_i's yield an open set V in Y, and then $U \times V \subset W$ is a product neighborhood of (a, b) in $X \times Y$. □

Euclidean Neighborhood Retracts

At certain places in this book it is desirable to know that a given compact space is a retract of a finite simplicial complex, or equivalently (as we shall see) a retract of a neighborhood in some Euclidean space. For example, this condition occurs in the Lefschetz fixed point theorem, and it was used in the proof of Alexander duality. So let us study this situation in more detail.

Theorem A.6. *A compact subspace K of $\mathbb{R}^n$ is a retract of some neighborhood iff K is locally contractible in the weak sense that for each $x \in K$ and each neighborhood U of x in K there exists a neighborhood $V \subset U$ of x such that the inclusion $V \hookrightarrow U$ is nullhomotopic.*

Note that if K is a retract of some neighborhood, then it is a retract of every smaller neighborhood, just by restriction of the retraction. So it does not matter if we require the neighborhoods to be open. Similarly it does not matter if the neighborhoods U and V in the statement of the theorem are required to be open.

Proof: Let us do the harder half first, constructing a retraction of a neighborhood of K onto K under the local contractibility assumption. The first step is to put a CW structure on the open set $X = \mathbb{R}^n - K$, with the size of the cells approaching zero near K. Consider the subdivision of $\mathbb{R}^n$ into unit cubes of dimension n with vertices at the points with integer coordinates. Call this collection of cubes C_0. For an integer $k > 0$,

we can subdivide the cubes of C_0 by taking n-dimensional cubes of edgelength $1/2^k$ with vertices having coordinates of the form $i/2^k$ for $i \in \mathbb{Z}$. Denote this collection of cubes by C_k. Let $A_0 \subset C_0$ be the cubes disjoint from K, and inductively, let $A_k \subset C_k$ be the cubes disjoint from K and not contained in cubes of A_j for $j < k$. The open set X is then the union of all the cubes in the combined collection $A = \bigcup_k A_k$. Note that the collection A is locally finite: each point of X has a neighborhood meeting only finitely many cubes in A, since the point has a positive distance from the closed set K.

If two cubes of A intersect, their intersection is an i-dimensional face of one of them for some $i < n$. Likewise, when two faces of cubes of A intersect, their intersection is a face of one of them. This implies that the open faces of cubes of A which are minimal with respect to inclusion among such faces form the cells of a CW structure on X, since the boundary of such a face is a union of such faces. The vertices of this CW structure are thus the vertices of all the cubes of A, and the n-cells are the interiors of the cubes of A.

Next we define inductively a subcomplex Z of this CW structure on X and a map $r : Z \to K$. The 0-cells of Z are exactly the 0-cells of X, and we let r send each 0-cell to the closest point of K, or if this is not unique, any one of the closest points of K. Assume inductively that Z^k and $r : Z^k \to K$ have been defined. For a cell e^{k+1} of X with boundary in Z^k, if the restriction of r to this boundary extends over e^{k+1} then we include e^{k+1} in Z^{k+1} and we let r on e^{k+1} be such an extension which is not too large, say an extension for which the diameter of its image $r(e^{k+1})$ is less than twice the infimum of the diameters for all possible extensions. This defines Z^{k+1} and $r : Z^{k+1} \to K$. At the end of the induction we set $Z = Z^n$.

It remains to verify that by letting r equal the identity on K we obtain a continuous retraction $Z \cup K \to K$, and that $Z \cup K$ contains a neighborhood of K. Given a point $x \in K$, let U be a ball in the metric space K centered at x. Choose a sequence of balls in K centered at x, $U = U_n \supset V_n \supset U_{n-1} \supset V_{n-1} \supset \cdots \supset U_0 \supset V_0$, each having radius equal to some small fraction, e.g., a quarter, of the radius of the preceding one, and with V_i contractible in U_i. Let $B \subset \mathbb{R}^n$ be a ball centered at x with radius less than half the radius of V_0, and let Y be the subcomplex of X formed by the cells whose closures are contained in B. Thus $Y \cup K$ contains a neighborhood of x in $\mathbb{R}^n$. By the choice of B and the definition of r on 0-cells we have $r(Y^0) \subset V_0$. Since V_0 is contractible in U_0, r is defined on the 1-cells of Y. Also, $r(Y^1) \subset V_1$ by the definition of r on 1-cells and the fact that U_0 is much smaller than V_1. Similarly, by induction we have r defined on Y^i with $r(Y^i) \subset V_i$ for all i. In particular, r maps Y to U. Since U could be arbitrarily small, this shows that extending r by the identity map on K gives a continuous map $r : Z \cup K \to K$. And since $Y \subset Z$, we see that $Z \cup K$ contains a neighborhood of K by

the earlier observation that $Y \cup K$ contains a neighborhood of x. Thus $r : Z \cup K \to K$ retracts a neighborhood of K onto K.

Now for the converse. Since open sets in $\mathbb{R}^n$ are locally contractible, it suffices to show that a retract of a locally contractible space is locally contractible. Let $r : X \to A$ be a retraction and let $U \subset A$ be a neighborhood of a given point $x \in A$. If X is locally contractible, then inside the open set $r^{-1}(U)$ there is a neighborhood V of x which is contractible in $r^{-1}(U)$, say by a homotopy $f_t : V \to r^{-1}(U)$. Then $V \cap A$ is contractible in U via the restriction of the composition rf_t. □

A space X is called a **Euclidean neighborhood retract** or **ENR** if for some n there exists an embedding $i : X \hookrightarrow \mathbb{R}^n$ such that $i(X)$ is a retract of some neighborhood in $\mathbb{R}^n$. The preceding theorem implies that the existence of the retraction is independent of the choice of embedding, at least when X is compact.

Corollary A.7. *A compact space is an ENR iff it can be embedded as a retract of a finite simplicial complex. Hence the homology groups and the fundamental group of a compact ENR are finitely generated.*

Proof: A finite simplicial complex K with n vertices is a subcomplex of a simplex Δ^{n-1}, hence embeds in $\mathbb{R}^n$. The proposition then implies that K is a retract of some neighborhood in $\mathbb{R}^n$, hence any retract of K is also a retract of such a neighborhood, via the composition of the two retractions. Conversely, let K be a compact space which is a retract of some open neighborhood U in $\mathbb{R}^n$. Since K is compact it is bounded, hence lies in some large simplex $\Delta^n \subset \mathbb{R}^n$. Subdivide Δ^n, say by repeated barycentric subdivision, so that all simplices of the subdivision have diameter less than the distance from K to the complement of U. Then the union of all the simplices in this subdivision which intersect K is a finite simplicial complex which retracts onto K via the restriction of the retraction $U \to K$. □

Corollary A.8. *Every closed manifold, with or without boundary, is an ENR.*

Proof: Since manifolds are locally contractible, it suffices to show that a compact manifold M can be embedded in $\mathbb{R}^k$ for some k. If M is not closed, it embeds in a closed manifold, namely the manifold obtained from two copies of M by identifying their boundaries. So it suffices to consider the case that M is closed. By compactness there exist finitely many closed balls $B_i^n \subset M$ whose interiors cover M, where n is the dimension of M. Let $f_i : M \to S^n$ be the quotient map which collapse the complement of the interior of B_i^n to a point. These f_i's are the components of a map $f : M \to (S^n)^m$ which is

injective since if x and y are distinct points of M with x in the interior of B^n_i, say, then $f_i(x) \neq f_i(y)$. Composing f with an embedding $(S^n)^m \hookrightarrow \mathbb{R}^k$, for example the product of the standard embeddings $S^n \hookrightarrow \mathbb{R}^{n+1}$, we obtain a continuous injection $M \hookrightarrow \mathbb{R}^k$, and this is a homeomorphism onto its image since M is compact. □

Corollary A.9. *Every finite CW complex is an ENR.*

Proof: Since CW complexes are locally contractible, it suffices to show that a finite CW complex can be embedded in some $\mathbb{R}^n$. This is proved by induction on the number of cells. Suppose the CW complex X is obtained from a subcomplex A by attaching a cell e^k via a map $f : S^{k-1} \to A$, and suppose that we have an embedding $A \hookrightarrow \mathbb{R}^m$. Then we can embed X in $\mathbb{R}^k \times \mathbb{R}^m \times \mathbb{R}$ as the union of the subspaces $D^k \times \{0\} \times \{0\}$, $\{0\} \times A \times \{1\}$, and all line segments joining points $(x, 0, 0)$ and $(0, f(x), 1)$ for $x \in S^{k-1}$. □

Spaces Dominated by CW Complexes

We have been considering spaces which are retracts of finite simplicial complexes, and now we show that such spaces have the homotopy type of CW complexes. In fact, we can just as easily prove something a little more general than this. A space Y is said to be **dominated** by a space X if there are maps $Y \xrightarrow{i} X \xrightarrow{r} Y$ with $ri \simeq \mathbb{1}$. This makes the notion of a retract into something which depends only on the homotopy types of the spaces involved.

Proposition A.10. *A space dominated by a CW complex is homotopy equivalent to a CW complex.*

Proof: We will use mapping telescopes. Recall from §3.4f that the telescope $T(f_1, f_2, \cdots)$ of a sequence of maps

$$X_1 \xrightarrow{f_1} X_2 \xrightarrow{f_2} X_3 \longrightarrow \cdots$$

is the quotient space of the disjoint union of the products $X_i \times [i, i+1]$, obtained by identifying $(x, i+1) \in X_i \times [i, i+1]$ with $(f(x), i+1) \in X_{i+1} \times [i+1, i+2]$. We shall need the following elementary facts:

(1) $T(f_1, f_2, \cdots) \simeq T(g_1, g_2, \cdots)$ if $f_i \simeq g_i$ for each i.
(2) $T(f_1, f_2, \cdots)$ deformation retracts onto $T(f_2, f_3, \cdots)$.
(3) $T(f_2 f_1, f_4 f_3, \cdots)$ is a deformation retract, in the weak sense, of $T(f_1, f_2, \cdots)$.

The first of these follows since $T(f_1, f_2, \cdots)$ is a deformation retract of $T(F_1, F_2, \cdots)$ where $F_i : X_i \times I \to X_{i+1} \times I$ is defined by $(x, t) \in (h_t(x), t)$ for h_t a homotopy of f_i. Namely, $T(F_1, F_2, \cdots)$ is a quotient of the disjoint union of products $X_i \times I \times [i, i+1]$,

and the desired deformation retraction of $T(F_1, F_2, \cdots)$ is the quotient of deformation retractions of $X_i \times I \times [i, i+1]$ first onto $X_i \times (\{0\} \times [i, i+1] \cup I \times \{i, i+1\})$, then onto $X_i \times \{0\} \times [i, i+1]$. The deformation retraction in (2) comes from shrinking $X_1 \times [1, 2]$ onto $X_1 \times \{2\}$. The proof of (3) is similar, $X_{2i} \times [2i, 2i+1]$ shrinking to $X_{2i} \times \{2i+1\}$ while $X_{2i-1} \times [2i-1, 2i]$ expands to $X_{2i-1} \times [2i-1, 2i+1]$.

Now to prove the proposition, suppose that the space Y is dominated by the CW complex X via maps $Y \xrightarrow{i} X \xrightarrow{r} Y$ with $ri \simeq \mathbb{1}$. By (2) and (3) we have $T(ir, ir, \cdots) \simeq T(r, i, r, i, \cdots) \simeq T(i, r, i, r, \cdots) \simeq T(ri, ri, \cdots)$. Since $ri \simeq \mathbb{1}$, $T(ri, ri, \cdots)$ is homotopy equivalent to the telescope of the identity maps $Y \to Y \to Y \to \cdots$, which is $Y \times [0, \infty) \simeq Y$. On the other hand, the map ir is homotopic to a cellular map $f : X \to X$, so $T(ir, ir, \cdots) \simeq T(f, f, \cdots)$, which is a CW complex. □

One might ask whether a space dominated by a finite CW complex is homotopy equivalent to a finite CW complex. In the simply-connected case this follows from Proposition 4.52 since such a space has finitely generated homology groups. But there are counterexamples in the general case; see [Wall].

In view of Corollary A.9 the preceding proposition implies:

Corollary A.11. *Every compact manifold is homotopy equivalent to a CW complex.* □

One could ask more refined questions. For example, do all compact manifolds have CW complex structures, or even simplicial complex structures? Answers here are considerably harder to come by. Restricting attention to closed manifolds for simplicity, the present status of these questions is the following. For manifolds of dimensions less than 4, simplicial complex structures always exist. In dimension 4 there are closed manifolds which do not have simplicial complex structures, while the existence of CW structures is an open question. In dimensions greater than 4, CW structures always exist, but whether simplicial structures always exist is unknown, though it is known that there are n-manifolds not having simplicial structures locally isomorphic to any linear simplicial subdivision of $\mathbb{R}^n$, for all $n \geq 4$. See [Kirby-Siebenmann] and [Freedman-Quinn] for more information on these questions.

Exercises

1. Show that a covering space of a CW complex is also a CW complex, with cells projecting homeomorphically onto cells.

2. Let X be a CW complex and x_0 any point of X. Construct a new CW complex structure on X having x_0 as a 0-cell, and having each of the original cells a union of

the new cells. The latter condition is expressed by saying the new CW structure is a **subdivision** of the old one.

3. Show that a CW complex X is path-connected iff its 1-skeleton X^1 is path-connected.

4. Show that a CW complex is locally compact iff each point has a neighborhood meeting only finitely many cells.

5. For a space X, show that the identity map $X_c \to X$ induces an isomorphism on π_1, where X_c denotes X with the compactly generated topology.

2. The Compact-Open Topology

By definition, the compact-open topology on the space X^Y of maps $f: Y \to X$ has a sub-basis consisting of the sets $M(K,U)$ of mappings taking a compact set $K \subset Y$ to an open set $U \subset X$. Thus a basis for X^Y consists of sets of maps taking a finite number of compact sets $K_i \subset Y$ to open sets $U_i \subset X$. If Y is compact, which is the only case we consider in this book, convergence to $f \in X^Y$ then means, loosely speaking, that finer and finer compact covers $\{K_i\}$ of Y are taken to smaller and smaller open covers $\{U_i\}$ of $f(Y)$. One of the main cases of interest in homotopy theory is when $Y = I$, so X^I is the space of paths in X. In this case one can check that a system of basic neighborhoods of a path $f: I \to X$ consists of the open sets $\bigcap_i M(K_i, U_i)$ where the K_i's are a partition of I into non-overlapping closed intervals and U_i is an open neighborhood of $f(K_i)$.

A fact which we will not actually need but which is reassuring is:

Proposition A.12. *If X is a metric space and Y is compact, the compact-open topology on X^Y is the same as the topology defined by the metric $d(f,g) = \sup_{y\in Y} d(f(y), g(y))$.*

Proof: First we show that every open ε-ball $B_\varepsilon(f)$ about $f \in X^Y$ contains a neighborhood $\bigcap_i M(K_i, U_i)$ of f in the compact-open topology. Since $f(Y)$ is compact, it is covered by finitely many balls $B_{\varepsilon/3}\big(f(y_i)\big)$. Let $K_i \subset Y$ be the closure of $f^{-1}\big(B_{\varepsilon/3}(f(y_i))\big)$, so K_i is compact, $Y = \cup_i K_i$, and $f(K_i) \subset B_{\varepsilon/2}\big(f(y_i)\big) = U_i$, hence $f \in \bigcap_i M(K_i, U_i)$. To show that $\bigcap_i M(K_i, U_i) \subset B_\varepsilon(f)$, suppose $g \in \bigcap_i M(K_i, U_i)$. For any $y \in Y$, say $y \in K_i$, we have $d\big(g(y), f(y_i)\big) < \varepsilon/2$ since $g(K_i) \subset U_i$. Likewise $d\big(f(y), f(y_i)\big) < \varepsilon/2$, so $d\big(f(y), g(y)\big) \le d\big(f(y), f(y_i)\big) + d\big(g(y), f(y_i)\big) < \varepsilon$. Since y was arbitrary, this shows $g \in B_\varepsilon(f)$.

Conversely, we show that for each open set $M(K,U)$ and each $f \in M(K,U)$ there is a ball $B_\varepsilon(f) \subset M(K,U)$. Since $f(K)$ is compact, it has a distance $\varepsilon > 0$ from the

complement of U. Then $d(f,g) < \varepsilon/2$ implies $g(K) \subset U$ since $g(K)$ is contained in an $\varepsilon/2$-neighborhood of $f(K)$. So $B_{\varepsilon/2}(f) \subset M(K,U)$. □

The next proposition contains the essential properties of the compact-open topology from the viewpoint of algebraic topology.

Proposition A.13. *If Y is locally compact Hausdorff, then:*

(a) The evaluation map $e : X^Y \times Y \to X$, $e(f,y) = f(y)$, is continuous.

(b) A map $f : Y \times Z \to X$ is continuous iff the associated map $\widehat{f} : Z \to X^Y$, $\widehat{f}(z)(y) = f(y,z)$, is continuous.

Proof: (a) For $(f,y) \in X^Y \times Y$ let $U \subset X$ be an open neighborhood of $f(y)$. By the hypothesis on Y, continuity of f implies there is a compact neighborhood $K \subset Y$ of y such that $f(K) \subset U$. Let $M(K,U)$ be the set of maps $g : Y \to X$ with $g(K) \subset U$. This is a neighborhood of f in X^Y, and e takes $M(K,U) \times K$ to U, so e is continuous at (f,y).

(b) Suppose $f : Y \times Z \to X$ is continuous. To show continuity of $\widehat{f}$ it suffices to show that for a sub-basic set $M(K,U) \subset X^Y$, the set $\widehat{f}^{-1}\big(M(K,U)\big) = \{\, z \in Z \mid f(K,z) \subset U \,\}$ is open in Z. Let $z \in \widehat{f}^{-1}\big(M(K,U)\big)$. Since $f^{-1}(U)$ is an open neighborhood of the compact set $K \times \{z\}$, there are open sets $V \subset Y$ and $W \subset Z$ with $K \times \{z\} \subset V \times W \subset f^{-1}(U)$, and so W is a neighborhood of z in $\widehat{f}^{-1}\big(M(K,U)\big)$. (The hypothesis that Y is locally compact Hausdorff is not needed here.)

For the converse half of (b) note that f is the composition $Y \times Z \to Y \times X^Y \to X$ of $\mathbb{1} \times \widehat{f}$ and the evaluation map, so part (a) gives the result. □

Proposition A.14. *If X is a compactly generated Hausdorff space and Y is locally compact Hausdorff, then the product topology on $X \times Y$ is compactly generated.*

Proof: First a preliminary observation: A function $f : X \times Y \to Z$ is continuous iff its restrictions $f : C \times Y \to Z$ are continuous for all compact $C \subset X$. For, using the previous proposition, the first statement is equivalent to $\widehat{f} : X \to Z^Y$ being continuous and the second statement is equivalent to $\widehat{f} : C \to Z^Y$ being continuous for all compact $C \subset X$. Since X is compactly generated, the latter two statements are equivalent.

To prove the proposition we just need to show that the identity map $X \times Y \to (X \times Y)_c$ is continuous. By the previous paragraph, this is equivalent to continuity of $C \times Y \to (X \times Y)_c$ for all compact $C \subset X$. Since Y is locally compact, it is compactly generated, so by the same reasoning, continuity of $C \times Y \to (X \times Y)_c$ is equivalent to continuity of $C \times C' \to (X \times Y)_c$ for all compact $C' \subset Y$. But on compact subsets such as $C \times C'$, the two topologies on $X \times Y$ agree, so we are done. (This proof is from [Dugundji].) □

Bibliography

Books

J. F. Adams, *Stable Homotopy and Generalised Homology*, Univ. of Chicago Press, 1974.

J. F. Adams, *Infinite Loop Spaces*, Ann. of Math. Studies 90, 1978.

A. Adem and R. J. Milgram, *Cohomology of Finite Groups*, Springer-Verlag, 1994.

M. A. Armstrong, *Basic Topology*, Springer-Verlag, 1983.

M. F. Atiyah, *K-Theory*, W.A.Benjamin, 1967.

H. J. Baues, *Homotopy Type and Homology*, Oxford Univ. Press, 1996.

D. J. Benson, *Polynomial Invariants of Finite Groups*, Cambridge Univ. Press, 1993.

D. J. Benson, *Representations and Cohomology, Volume II: Cohomology of Groups and Modules*, Cambridge Univ. Press, 1992.

R. Bott and L. Tu, *Differential Forms in Algebraic Topology*, Springer-Verlag GTM 82, 1982.

G. Bredon, *Topology and Geometry*, Springer-Verlag GTM 139, 1993.

K. Brown, *Cohomology of Groups*, Springer-Verlag GTM 87, 1982.

R. Brown, *The Lefschetz Fixed Point Theorem*, Scott Foresman, 1971.

M. Cohen, *A Course in Simple-Homotopy Theory*, Springer-Verlag GTM 10, 1973.

J. Dieudonné, *A History of Algebraic and Differential Topology 1900-1960*, Birkhäuser, 1989.

A. Dold, *Lectures on Algebraic Topology*, Springer-Verlag, 1980.

J. Dugundji, *Topology*, Allyn & Bacon, 1966.

H.-D. Ebbinghaus et. al., *Numbers*, Springer-Verlag GTM 123, 1991.

S. Eilenberg and N. Steenrod, *Foundations of Algebraic Topology*, Princeton Univ. Press, 1952.

R. Fenn, *Techniques of Geometric Topology*, Cambridge Univ. Press, 1983.

A. T. Fomenko and D. B. Fuks, *A Course in Homotopic Topology*, Izd. "Nauka," 1989. (In Russian; an English translation of an earlier version was published by Akadémiai Kiadó, Budapest, 1986.)

M. Freedman and F. Quinn, *Topology of 4-Manifolds*, Princeton Univ. Press, 1990.

W. Fulton, *Algebraic Topology: A First Course*, Springer-Verlag, 1995.

S. I. Gelfand and Yu. I. Manin, *Methods of Homological Algebra*, Springer-Verlag, 1996.

B. Gray, *Homotopy Theory*, Academic Press, 1975.

M. Greenberg and J. Harper, *Algebraic Topology: A First Course*, Addison-Wesley, 1981.

P. Griffiths and J. Morgan, *Rational Homotopy Theory and Differential Forms*, Birkhäuser, 1981.

P. J. Hilton, *An Introduction to Homotopy Theory*, Cambridge University Press, 1953.

P. J. Hilton, *Homotopy Theory and Duality*, Gordon and Breach, 1966.

P. J. Hilton and S. Wylie, *Homology Theory*, Cambridge Univ. Press, 1967.

P. J. Hilton and U. Stammbach, *A Course in Homological Algebra*, Springer-Verlag, 1970.

P. J. Hilton, G. Mislin, and J. Roitberg, *Localization of Nilpotent Groups and Spaces*, North-Holland, 1975.

C. Hog-Angeloni, W. Metzler, A. Sieradski eds., *Two-dimensional Homotopy and Combinatorial Group Theory*, Cambridge Univ. Press, 1993.

D. Husemoller, *Fibre Bundles*, McGraw-Hill, 1966 (later editions by Springer-Verlag).

I. M. James, ed., *Handbook of Algebraic Topology*, North-Holland, 1995.

R. M. Kane, *The Homology of Hopf Spaces*, North-Holland, 1988.

J. Kelley, *General Topology*, Van Nostrand Reinhold, 1955 (reprinted by Springer-Verlag).

R. Kirby and L. Siebenmann, *Foundational Essays on Topological Manifolds, Smoothings, and Triangulations*, Ann. of Math. Studies 88, 1977.

S. Kochman, *Bordism, Stable Homotopy and Adams Spectral Sequences*, Fields Institute Monographs 7, A. M. S., 1996.

S. Kochman, *Stable Homotopy Groups of Spheres*, Springer Lecture Notes 1423, 1990.

D. König, *Theory of Finite and Infinite Graphs*, Birkhäuser, 1990.

A. Lundell and S. Weingram, *The Topology of CW Complexes*, Van Nostrand Reinhold, 1969.

S. MacLane, *Homology*, Springer-Verlag, 1963.

S. MacLane, *Categories for the Working Mathematician*, Springer-Verlag GTM 5, 1971.

I. Madsen and R. Milgram, *The Classifying Spaces for Surgery and Cobordism of Manifolds*, Ann. of Math. Studies 92, 1979.

W. Massey, *Algebraic Topology: An Introduction*, Harcourt, Brace & World, 1967 (reprinted by Springer-Verlag).

W. Massey, *A Basic Course in Algebraic Topology*, Springer-Verlag, 1993.

C. R. F. Maunder, *Algebraic Topology*, Cambridge Univ. Press, 1980 (reprinted by Dover Publications).

J. P. May, *Simplicial Objects in Algebraic Topology*, Van Nostrand, 1967 (reprinted by Univ. Chicago Press).

J. Milnor, *Topology from the Differentiable Viewpoint*, Univ. Press of Virginia, 1965.

J. Milnor and J. Stasheff, *Characteristic Classes*, Ann. of Math. Studies 76, 1974.

R. Mosher and M. Tangora, *Cohomology Operations and Applications in Homotopy Theory*, Harper and Row, 1968.

D. Ravenel, *Complex Cobordism and Stable Homotopy Groups of Spheres*, Academic Press, 1986.

D. Ravenel, *Nilpotence and Periodicity in Stable Homotopy Theory*, Ann. of Math. Studies 128, 1992.

E. Rees and J.D.S. Jones, eds., *Homotopy Theory: Proceeding of the Durham Symposium 1985*, London Math. Soc. Lect. Note Ser. 117, Cambridge Univ. Press, 1987.

D. Rolfsen, *Knots and Links*, Publish or Perish, 1976.

P. Selick, *Introduction to Homotopy Theory*, Fields Institute Monographs 9, A. M. S., 1997.

J-P. Serre, *A Course in Arithmetic*, Springer-Verlag, 1973.

E. Spanier, *Algebraic Topology*, McGraw-Hill, 1966 (reprinted by Springer-Verlag).

N. Steenrod, *The Topology of Fibre Bundles*, Princeton Univ. Press, 1951.

N. Steenrod and D. Epstein, *Cohomology Operations*, Ann. of Math. Studies 50, 1962.

R. Stong, *Notes on Cobordism Theory*, Princeton Univ. Press, 1968.

D. Sullivan, *Geometric Topology*, xeroxed notes from MIT, 1970.

R. Switzer, *Algebraic Topology*, Springer-Verlag, 1975.

H. Toda, *Composition Methods in Homotopy Groups of Spheres*, Ann. of Math. Studies 49, 1962.

K. Varadarajan, *The Finiteness Obstruction of C.T.C.Wall*, Wiley, 1989.

G. Whitehead, *Elements of Homotopy Theory*, Springer-Verlag GTM 62, 1978.

J. Wolf, *Spaces of Constant Curvature*, Publish or Perish, 1984.

Papers

J. F. Adams, On the non-existence of elements of Hopf invariant one, *Ann. of Math.* 72 (1960), 20-104.

J. F. Adams, Vector fields on spheres, *Ann. of Math.* 75 (1962), 603-632.

J. F. Adams, On the groups J(X) IV, *Topology* 5 (1966), 21-71.

J. F. Adams and M. F. Atiyah, K-theory and the Hopf invariant, *Quart. J. Math.* 17 (1966), 31-38.

J. F. Adams, A variant of E. H. Brown's representability theorem, *Topology* 10 (1971), 185-198.

J. F. Adams and C. Wilkerson, Finite H-spaces and algebras over the Steenrod algebra, *Ann. of Math.* 111 (1980), 95-143.

M. Barratt and J. Milnor, An example of anomalous singular homology, *Proc. A.M.S.* 13 (1962), 293-297.

A. Borel, Sur la cohomologie des espaces fibrés principaux, *Ann. of Math.* 57 (1953), 115-207.

R. Bott and J. Milnor, On the parallelizability of spheres, *Bull. A.M.S.* 64 (1958), 87-89.

E. Brown and A. Copeland, An homology analog of Postnikov systems, *Michigan Math. J.* 6 (1959), 313-330.

M. Brown, A proof of the generalized Schoenflies theorem, *Bull. A.M.S.* 66 (1960), 74-76.

G. Carlsson and R. J. Milgram, Stable homotopy and iterated loopspaces, pp. 505-583 in *Handbook of Algebraic Topology*, ed. I.M. James, Elsevier 1995.

J. Davis and R. J. Milgram, A survey of the spherical space form problem, *Math. Reports* 2 (1985), 223-283, Harwood Acad. Pub.

A. Dold and R. Thom, Quasifaserungen und unendliche symmetrischen Produkte, *Ann. of Math.* 67 (1958), 239-281.

C. Dowker, Topology of metric complexes, *Am. J. Math.* 74 (1952), 555-577.

W. Dwyer, H. Miller, and C. Wilkerson, Homotopical uniqueness of classifying spaces, *Topology* 31 (1992), 29-45.

W. Dwyer and C. Wilkerson, A new finite loop space at the prime two, *J. Am. Math. Soc.* (6) 1993, 37-63.

E. Dyer and A. Vasquez, Some small aspherical spaces, *J. Austral. Math. Soc.* 16 (1973), 332-352.

S. Feder and S. Gitler, Mappings of quaternionic projective spaces, *Bol. Soc. Mat. Mexicana* 18 (1978), 33-37.

J. Harris and N. Kuhn, Stable decompositions of classifying spaces of finite abelian p-groups, *Math. Proc. Camb. Phil. Soc.* 103 (1988), 427-449.

P. Hoffman and G. Porter, Cohomology realizations of $\mathbb{Q}[x]$, *Ox. Q.* 24 (1973), 251-255.

R. Holzsager, Stable splitting of K(G,1), *Proc. A.M.S.* 31 (1972), 305-306.

I. M. James, Reduced product spaces, *Ann. of Math.* 62 (1955), 170-197.

S. Kochman and M. Mahowald, On the computation of stable stems, *Contemp. Math.* 181 (1995), 299-316.

I. Madsen, C. B. Thomas, and C. T. C. Wall, The topological spherical space form problem, II: existence of free actions, *Topology* 15 (1976), 375-382.

J. P. May, A general approach to Steenrod operations, *Springer Lecture Notes* 168 (1970), 153-231.

J. P. May, Weak equivalences and quasifibrations, *Springer Lecture Notes* 1425 (1990), 91-101.

J. Milnor, Construction of universal bundles I, II, *Ann. of Math.* 63 (1956), 272-284, 430-436.

J. Milnor, Groups which act on S^n without fixed points, *Am. J. Math* 79 (1957), 623-630.

J. Milnor, On spaces having the homotopy type of a CW complex, *Trans. A.M.S.* 90 (1959), 272-280.

J. Milnor, On axiomatic homology theory, *Pac. J. Math.* 12 (1962), 337-341.

J. Milnor and J. Moore, On the structure of Hopf algebras, *Ann. of Math.* 81 (1965), 211-264.

J. Milnor, On the Steenrod homology theory, pp. 79-96 in *Novikov Conjectures, Index Theorems, and Rigidity, vol. 1*, ed. S. Ferry, A. Ranicki, and J. Rosenberg, Cambridge Univ. Press, 1995.

J. Morgan and I. Morrison, A van Kampen theorem for weak joins, *Proc. London Math. Soc.* 53 (1986), 562-576.

J. Neisendorfer, Primary homotopy theory, *Mem. A.M.S.* 232 (1980).

D. Notbohm, Spaces with polynomial mod-p cohomology, *to appear.*

D. Quillen, Rational homotopy theory, *Ann. of Math.* 90 (1969), 205-295.

D. Rector, Loop structures on the homotopy type of S^3, *Springer Lecture Notes* 249 (1971), 99-105.

P. Scott and T. Wall, Topological methods in group theory, *London Math. Soc. Lecture Notes* 36 (1979), 137-203.

J. Stallings, A finitely presented group whose 3-dimensional integral homology is not finitely generated, *Am. J. Math.* 85 (1963), 541-543.

D. Sullivan, Genetics of homotopy theory and the Adams conjecture, *Ann. of Math.* 100 (1974), 1-79.

H. Toda, Note on the cohomology ring of certain spaces, *Proc. A.M.S.* 14 (1963), 89-95.

C. T. C. Wall, Finiteness conditions for CW complexes, *Ann. of Math.* 81 (1965), 56-69.

G. Whitehead, Generalized homology theories, *Trans. A.M.S.* 102 (1962), 227-283.

Index

abelian space 328, 407
action of π_1 on π_n 328, 332, 412
action of a group 66
acyclic space 136
Adams 415
Adém relations 489, 494
adjoint functors 453
admissible monomial 492
Alexander 125, 166
Alexander duality 234
Alexander horned sphere 161
aspherical space 329
attaching 8, 16
augmented chain complex 105

Barratt-Kahn-Priddy theorem 366
barycenter 114
barycentric coordinates 98
barycentric subdivision 115
base space 385
basepoint-preserving homotopy 37, 363, 412
basepoint 28, 29
Betti number 125
binomial coefficient 264, 483
Bockstein homomorphism 283, 480
Borel construction 307
Borel theorem 265
Borsuk-Ulam theorem 60, 74
Bott periodicity 392, 447
boundary homomorphism 100, 103, 112
boundary 101, 232
Brouwer 35, 109, 119, 128, 163, 166
Brown representability 442
BSO(n) 429
BSU(n) 429
bundle of groups 315
Burnside problem 75

cap product 218
Cartan formula 481, 482
category 154
Cayley graph, complex 71
Čech cohomology, homology 237
cell complex 9
cellular approximation theorem 343, 357
cellular chain complex 133
cellular cohomology 186
cellular homology 134, 145
cellular map 251, 343, 357
cell 8
chain complex 101
chain homotopy 108
chain map 106
chain 99, 103
change of basepoint 29, 328
characteristic map 10, 513
circle 32
classifying space 157
closed manifold 213
closure-finite 514
coboundary 183
coboundary map 174, 182
cochain 174, 182
cocycle 183
coefficients 145, 152, 182, 454
cofiber 453
cofibration 452
cofibration sequence 438
Cohen-Macauley ring 210

cohomology group 174, 182
cohomology operation 479
cohomology ring 193
cohomology theory 185, 296, 437, 440, 445, 450
cohomology with compact supports 224
cohomotopy groups 450
collar 232
commutative diagram 106
commutative graded ring 194
commutativity of cup product 193
compact supports 224, 321
compact-open topology 523
compactly generated topology 516, 524
cone 12
connected graded algebra 263
contractible 5, 148
contravariant 156, 184
coproduct 263
covariant 156
covering homotopy property 57
covering space action 67
covering space 53, 305, 329, 385
covering transformation 65
cross product 200, 250, 258
cup product 189, 228
CW approximation 359, 364
CW complex 9, 513
CW pair 12
cycle 101

deck transformation 65
decomposable operation 489
deformation retraction 2, 36, 333, 515
deformation retraction, weak 6
degree 128

Δ-complex 98
diagonal 263
dihedral group 71
direct limit 223, 292, 451
directed set 223
divided polynomial algebra 206, 266, 270
division algebra 163, 204, 415
dodecahedral group 136
Dold-Thom theorem 472
dominated 521
dual Hopf algebra 269

Eckmann-Hilton duality 452
edge 77
edgepath 79
EHP sequence 465
Eilenberg 125
Eilenberg-MacLane space 82, 344, 347, 394, 397, 437, 449, 467
ENR 520
Euclidean neighborhood retract 520
Euler characteristic 9, 140
Euler class 426, 433
exact sequence 108
excess 492
excision 117, 185, 366
excisive triad 467
extension problem 405
exterior algebra 197, 263
external cup product 196, 200
Ext 180, 298, 300

face 98
fiber bundle 384, 420
fiber homotopy equivalence 395
fiber-preserving map 395

fiber 382
fibration 382
fibration sequence 396
finitely generated 342
finitely generated homology 520
five-lemma 123
fixed point 35, 68, 110, 168, 210, 487
flag 424, 436
frame 281, 390
free action 68
free algebra 209
free group 42, 72, 80
free product 40
free product with amalgamation 87
free resolution 178, 243
Freudenthal suspension theorem 366, 374
function space 523
functor 155
fundamental class 219, 221, 347
fundamental group 28
fundamental theorem of algebra 34

general linear group GL_n 273
good pair 109
Gram-Schmidt orthogonalization 273, 391
graph of groups 87
graph product of groups 87
graph 9, 14, 77
Grassmann manifold 208, 390, 424, 428, 434
groups acting freely on spheres 70, 129, 355
Gysin sequence 426, 433

H-space 261, 355, 411, 413, 415
HNN extension 88
hocolim 453
holim 453
homologous 101
homology class 101
homology decomposition 456
homology group 101
homology theory 151, 296, 450
homotopy class 28
homotopy equivalence 4, 14, 37, 333
homotopy extension property 17
homotopy fiber 393, 473
homotopy group 325
homotopy group with coefficients 454
homotopy lifting property 57, 382, 387
homotopy limits 453
homotopy of attaching maps 16, 19
homotopy type 4
homotopy 3, 27
Hopf 128, 163, 204, 261, 265
Hopf algebra 263
Hopf bundle 367, 382, 385, 386, 409
Hopf invariant 415, 436, 481, 483
Hopf map 376, 387, 389, 415, 419, 465, 466, 491
Hurewicz homomorphism 477
Hurewicz map 335
Hurewicz theorem 337, 354, 369, 373

induced fibration 396
induced homomorphism 36, 105, 106, 113, 184, 195
infinite loopspace 447
invariance of dimension 119
invariance of domain 162
inverse limit 293, 399
inverse path 29
isomorphism of covering spaces 64

J-homomorphism 378
J(X), James reduced product 205, 261, 267, 269, 458, 461
join 12
Jordan curve theorem 161

k-invariant 467, 401
K(G,1) space 82
Klein bottle 49, 88, 97
Künneth formula 196, 250, 254, 256, 363, 420

Lefschetz 125, 168, 210
Lefschetz duality 233
Lefschetz number 168
lens space 70, 82, 137, 230, 261, 285, 290, 356
Leray-Hirsch theorem 420
lifting criterion 59
lifting problem 405
lift 32, 57
$\varprojlim^1$ 294, 399
linking 45
local coefficients: cohomology 320
local coefficients: homology 314
local degree 130
local homology 120, 236
local orientation 217
local trivialization 385
locally contractible 515, 518
locally finite homology 323
locally path-connected 59
long exact sequence: cohomology 184
long exact sequence: fibration 383
long exact sequence: homology 109, 110, 111, 112, 114
long exact sequence: homotopy 330
loop 28
loopspace 394, 437, 461

manifold 213, 520, 522
manifold with boundary 232
mapping cone 16, 170
mapping cylinder 3, 170, 334
mapping telescope 132, 293, 521
maximal tree 78
Mayer-Vietoris axiom 442
Mayer-Vietoris sequence 143, 149, 154, 187
Milnor 394, 395
minimal chain complex 283
Mittag-Leffler condition 299
monoid 155
Moore space 136, 257, 293, 303, 344, 355, 454, 455, 467
Moore-Postnikov tower 403
morphism 154

n-approximation 360, 398
n-connected cover 359, 404
n-connected space, pair 332
n-factorization 360
natural transformation 157
naturality 120
nerve 236
nonsingular pairing 229
normal covering space 66
nullhomotopic 5

object 154
obstruction theory 405
obstruction 407
octonion 163, 261, 274, 386, 491
Ω-spectrum 437
orbit, orbit space 67

orientable manifold 217
orientable sphere bundle 431
orientation class 219
orientation 100, 217
orthogonal group O(n) 273, 288, 423

p-adic integers 295
path lifting property 58
path 27
pathspace 393
plus construction 365, 411
Poincaré 125
Poincaré conjecture 354
Poincaré duality 219, 224, 233 322
Poincaré series 211, 425
Pontryagin product 267, 278
Postnikov tower 398
primary obstruction 409
primitive element 264, 278
principal fibration 400, 411
prism 107
product of Δ-complexes 259
product of CW complexes 11, 516
product of paths 28
product of simplices 259
product space 31, 250, 329, 524
projective plane 49, 97, 101, 193, 387
projective space: complex 10, 134, 202, 207, 211, 230, 261, 306, 388, 427, 484
projective space: quaternionic 203, 207, 211, 230, 306, 386, 388, 427, 484, 485
projective space: real 9, 70,82, 137, 146, 168, 202, 211, 230, 306, 427, 484
properly discontinuous 67
pullback 396, 422
Puppe sequence 438
pushout 453

quasi-circle 75
quasifibration 473
quaternion 70, 163, 261, 274, 436
Quillen 365
quotient 12

rank 140
reduced cohomology 183
reduced homology 105
relative cohomology 183
relative cycle, boundary 111
relative homology 111
relative homotopy group 329
rel 4
reparametrization 28
retract 36, 110, 518
retraction 4, 142

semilocally simply-connected 62
sheet 58
short exact sequence 109, 111
shrinking wedge 47
shuffle 259
simplex 13, 97
simplicial approximation theorem 166
simplicial cohomology 186
simplicial complex 102
simplicial homology 101, 122
simplicial map 166
simply-connected 30
simply-connected 4-manifold 418
singular complex 103, 362
singular homology 103
singular simplex 103

skeleton 8, 513
slant product 260
smash product 14, 200, 252
spectrum 449
sphere bundle 431, 433
split exact sequence 142
stable homotopy group 375, 448
stable splitting 458, 484
stable stem 375
star 166
Steenrod algebra 489
Steenrod homology 237
Steenrod squares, powers 479
Stiefel manifold 281, 390, 424, 436, 487
subcomplex 11, 513
subgraph 78
surface 48, 82, 88, 97, 135, 190, 220, 354
suspension spectrum 449
suspension 12, 131, 153, 200, 373, 458, 464
symmetric polynomials 424
symmetric product 261, 345, 470
symplectic group Sp(n) 208, 390, 422

tensor algebra 268, 462
tensor product of chain complexes 254
tensor product 241
Thom class 430, 503
Thom isomorphism 430
Thom space 431, 503
Toda bracket 379
topological group 261
torsion coefficient 125
torus knot 76
torus 35, 68, 97, 101, 209
Tor 244, 248
total space 385
transfer homomorphism 305
transitive action 66
tree 78
triple 114, 331
truncated polynomial algebra 264

unique lifting property 59
unitary group U(n) 208, 390, 422
universal coefficient theorem 180, 244, 454
universal cover 56, 65

van Kampen 42
vector field 129, 487
vertex 77, 98

weak homotopy equivalence 362
weak topology 9, 77, 514
wedge sum 13, 42, 119, 151, 185, 388, 458
Whitehead product 389, 419
Whitehead's theorem 333, 342, 408