K-theory and Characteristic Classes: A homotopical perspective

Inna Zakharevich

April 3, 2024

Contents

Introduction

Mathematics exposition struggles with two opposing forces. The first is the desire to explain things historically: A was understood before B, which motivated the definition for C, and thus we should tell it in this order. This often makes explaining motivations easier, as the discoveries of certain phenomena create new questions, which in turn spur the search for novel techniques to respond to them. Unfortunately, this is often not the best way to understand material—in the same way that a new arrival to town may not know the best ways of getting around, but only know a couple of streets and landmarks. Groups are no longer presented as subsets of permutations closed under compositions, even though their study originally arose in this way, because the axiomatic framework is clearer and easier to absorb.

I started writing this book by accident. I was teaching a course on characteristic classes and K-theory for the first time, and I was frustrated by the classic definition of the Steifel–Whitney classes in Milnor–Stasheff. All of the mathematics was correct, but it felt wrong to me, like someone was using a hammer to turn a screw. It's definitely possible, but it's not the way that the tools want to be used. I started to rewrite that section of the book to align with my intuition about how the material should fit together: with the cohomology of Grassmannians first, and the naming of the classes after, with all of their properties arising from the Grassmannian structure. But that one change snowballed into another, and another, and eventually this book was born.

I have written this book in the same way that I teach my classes: not a self-contained tour, but an introduction to the many branches that one's interest can take. I leave out technical details that I find less interesting, I skip computations that I don't think are illuminating, and I try to explain how one might come up with definitions from one's goals. The hope is to spark a student's interest—possibly not in something that I'm saying, but in something that I didn't say and merely alluded to. I also want students to understand that math does not arise cleanly from people coming up with definitions and then exploring their boundaries. It's messy and complicated, with one's goals often running up against what the mathematics knows it wants to do. It's often impossible to achieve one's goals at all, and we must instead follow the mathematics instead of leading it.

The mathematical goal of this book is to introduce topological K-theory through a more homotopy-theoretic viewpoint than the classic approach. The classic texts on these topics, including Milnor and Stasheff's *Characteristic Classes*, Atiyah's *K-theory* and Hatcher's unpublished *Vector Bundles & K-theory*, approach the topic from a geometric viewpoint: the goal is to construct geometric invariants of vector bundles and, as a side effect of this goal, homotopy-theoretic invariants emerge. The largest effect of this is that characteristic classes are introduced through their *use*: they are defined axiomatically, and shown to exist only so that the axioms may be comfortably used. Milnor–Stasheff's construction of characteristic classes uses Steenrod operations, and is therefore far more difficult to understand than the rest of the book. (In fact, the class that I took on characteristic classes skipped the construction entirely!)

In this work the perspective is flipped: the goal is to show how a homotopy-theoretic perspective can produce interesting geometric invariants. The book begins by classifying vector bundles in terms of homotopy classes of maps into Grassmannians. This result implies that if the set of homotopy classes of maps into a Grassmannian were understood, then all vector bundles could be easily classified. Cohomology is introduced (using spectra) as a computationally-feasible method for classifying homotopy classes of maps into a space, and is used to motivate constructing characteristic classes via pullbacks of classes in the cohomology of Grassmannians. The cohomology of Grassmannians is computed by induction, and shown to be a polynomial ring; characteristic classes are then defined as pullbacks of the generators. (This is done for real Grassmannians, rathan than complex, for two reasons: because the geometric applications we show use the real case, and because this is the more difficult computation. The complex case can be done analogously, without the additional difficulty of showing that the Euler class is nontrivial, while proving the real case knowing the proof for the complex one is still nontrivial. Surprisingly, I could not find this proof written up anywhere in the literature.) As a consequence of this definition, the four axioms of characteristic classes follow simply, showing its equivalence to the usual approach.

The next section of the book is applications of characteristic classes: non-parallelizability of $\mathbf{R}P^n$ for $n \neq 2^k - 1$ and the classification of cobordism groups as homotopy groups. In this way, Thom spaces and stable homotopy groups naturally arise. This motivates a homotopic proof of Bott periodicity: a construction, due to Dyer and Lashof, showing directly that $\Omega^8 O \cong \mathbb{Z} \times O$ and $\Omega^2 U \cong \mathbb{Z} \times U$. This proof has the advantage that it does not require much background other than basic algebraic topology (modulo a couple of computational results), while motivating topological K-theory as a cohomology theory: Bott periodicity implies that $\mathbb{Z} \times BU, U, \mathbb{Z} \times BU, U, \ldots$ is a spectrum, and the question becomes: what is the cohomology theory it represents? The book then moves into topological K-theory, doing the standard constructions and computation. Adams operations are introduced, and used to show that $\mathbf{R}P^n$ is paralellizable only for small n (thus harkening back to the geometric implications of characteristic classses).

For those who are interested in a more classic approach, each chapter will end with a "Further Reading" section which will give suggestions for students to read the same material from an alternate viewpoint. In addition, in chapters which cite theorems without proof, the "Further Reading" section will contain references to the proofs and background necessary for those results.

This book is targeted towards graduate students who have had one semester of algebraic topology: they are expected to have seen homology, a little bit of cohomology, and to be comfortable with the definition of homotopic maps. No other background is assumed. One of the goals of the exposition is to introduce not just homotopy-theoretic *results*, but the homotopy-theoretic *viewpoint*: the idea of classifying spaces by mapping into understood spaces, the idea of higher homotopy groups and spectra, and the notion that the homotopy type of a space can encode important invariants about all spaces. Therefore, although the geometric problems that we solve are treated as important, the *end* of the work and the *motivation* in its development is all from the point of view of homotopy.

Category theory is becoming a more and more foundational part of mathematics. As homotopy theory and category theory are inextricably intertwined, this book also expects the student to be familiar with certain definitions and results of category theory. The student comfortable with the basic notions (categories, functors, natural transformations, limits and colimits) can simply read the text; nothing beyond the basics is assumed. For the student who would like a refresher, or the student who enjoys a challege, we include a short crash course on category theory as Appendix A. This includes the definitions in use in the text as well as some basic exercises.

This book is not intended as a replacement for the classical works on topological K-theory and characteristic classes. Rather, it is intended as a complement: a different perspective on the material which may appeal to those students more inclined towards homotopy than geometry (or else as motivation for the geometrically-inclined to learn homotopy).

Assumed background

This text is aimed at a student who has taken a single semester of algebraic topology, and no more. Knowledge of the fundamental group and homology, and especially of homotopy classes of maps between spaces, is assumed. So is a basic level of comfort with exact sequences, both long and short. I assume that the student knows that higher homotopy groups exist, and their definition, but not anything beyond that; I also assume that they have heard of cohomology, Poincare duality, and the cup product, although I do not assume an expertise in these topics.

I also assume some basic category theory. A student who is able to do all of the problems in Appendix A knows everything they need to know for this book. A student who is not able to do the problems, but is brave and adventurous and willing to catch up as needed, will be fine.

Acknowledgements

I wrote this book over several years of teaching Math 6340 at Cornell, supported by NSF grants DMS-2052977 and CAREER-1846767. I would like to especially thank the early readers who bravely used it as their textbook, including especially Nicole Magill, Varinderjit Mann, Nikhil Sahoo, and Kimball Strong, who found errors in the early drafts. I would like to thank Ilya Zakharevich for reading some of the early chapters and tearing them apart (to my great benefit). I owe a great debt to the Saltonstall Foundation, whose retreat space benefitted this book massively. Lastly and especially, I would like to thank Thomas Barnet-Lamb, who in addition to being supportive and amazing in general, helped me go to Saltonstall in order to get this book into a human-readable state.

Chapter 1

A geometric beginning

In the lexicon of amusingly-named theorems, one of the most famous is the Hairy Ball Theorem:

Theorem 1.1 (Hairy Ball Theorem)**.** *You can't comb a hairy billiard ball.*

Apart from the immediate practical questions, the obvious question of what this means mathematically arises. Let us try to pull this apart. Think of the billiard ball as sitting as the unit sphere in $\mathbf{R}^3$. A "hair" growing from a point on the sphere can then be thought of as a unit vector originating at that point. If the ball is "combed" that means that the hairs lie smoothly next to one another along the ball: in other words, that at every point on the ball the vector is *tangent* to the sphere, changing continously on the sphere. We can thus rephrase the hairy ball theorem as follows:

Theorem 1.2 (Hairy Ball Theorem)**.** *Consider S^2 to be sitting inside $\mathbf{R}^3$ as the unit sphere. There does not exist a map $S^2 \longrightarrow S^2$ such that for all $x \in S^2$, $f(x) \cdot x = 0$.*

This statement of the theorem is unambiguous, but it does hide some of the structure present in the problem. The fact of the matter is that, morally speaking, we are assigning to each point on the sphere a *tangent vector* to that point. In the above statement we pretend that all tangent vectors live in the same place. Morally speaking, it should not be possible, for two distinct points $x, y \in S^2$, for their tangent vectors to be *equal*: they live in two different tangent planes. However, with the phrasing above, it is unexceptional to consider a non-injective map $S^2 \longrightarrow S^2$. This indicates that there ought to exist a better structure, with a better statement for the theorem, that will keep track of this information.

We define a space, called $S(TS^2)$, that will encode both the space of tangent vectors, as well as the points they are tangent to:

$$S(TS^2) \stackrel{\text{def}}{=} \{(x, v) \in S^2 \times S^2 \mid x \cdot v = 0\}.$$

Here, the dot product is defined by considering both x and v as points in $\mathbf{R}^3$. Each point in $S(TS^2)$ now knows both that it is a tangent vector, and also *which* point it is tangent to; this also gives the tangent space a natural topology (and in fact the structure of a manifold). There is a map $p\colon S(TS^2) \longrightarrow S^2$, given by $(x, v) \longmapsto x$. We can now rephrase the Hairy Ball Theorem as follows:

Theorem 1.3 (Hairy Ball Theorem)**.** *There does not exist a map $s\colon S^2 \longrightarrow S(TS^2)$ such that $p \circ s$ is the identity map.*

We are now in a position to generalize this definition, and the Hairy Ball Theorem, to more general manifolds.

Definition 1.4. Let M be a smooth n-manifold and pick an embedding $f\colon M \hookrightarrow \mathbf{R}^N$. Define the *tangent sphere bundle* of M to be

$$S(TM) \stackrel{\text{def}}{=} \{(x, v) \in M \times \mathbf{R}^N \mid v \text{ tangent to } M \text{ at } x \text{ and } |v| = 1\}.$$

The space $S(TM)$ is a manifol of dimension $2n-1$, with a natural (smooth) map $p\colon S(TM) \longrightarrow M$ given by forgetting the v-coordinate.

Although it is not immediately obvious that this is independent of the choice of embedding, it turns out to be the case. (See Example 2.5 and Exercise 2.3.)

The Hairy Ball Question can then be asked as follows:

> Which manifolds is it possible to comb? In other words, for which M does there exist a map $s\colon M \longrightarrow S(TM)$ such that $p \circ s$ is the identity map?

In this chapter we will formally introduce the main players and begin to develop results about them. The answer to this question, inasmuch as this book will provide it, is in Theorem 5.10. To begin we state explicitly what we mean by the words "space" and "map."

Definition 1.5. A *space* is a Hausdorff compactly-generated topological space.[a] A *map* $f\colon X \longrightarrow Y$ is a continuous function.

As with many questions in mathematics, analyzing this is easier if we work with a more general definition. This is one of the fundamental notions of this book: that of a *fiber bundle*.

First, a few motivational observations about the definition of the tangent sphere bundle:

- The tangent sphere bundle was encoded by a map $p\colon S(TM) \longrightarrow M$, where the preimage of every point was the same (a sphere of the appropriate dimension).

- In a neighborhood of every point, M is homeomorphic to an open subset $U \subseteq \mathbf{R}^n$, and $p^{-1}(U)$ is homeomorphic to $\mathbf{R}^n \times S^{n-1}$.

These are the properties emulated by the general definition of a fiber bundle.

Definition 1.6. Let $p\colon E \longrightarrow B$ be a map; for any $x \in B$, the set $p^{-1}(x)$ is the *fiber over* x. Let $U \subseteq B$ be a neighborhood. We say that p is *trivial on* U *with fiber* F if there exists a homeomorphism $\varphi\colon p^{-1}(U) \longrightarrow U \times F$ making the diagram

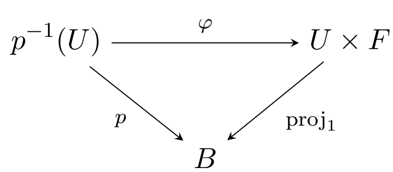

commute. The homeomorphism φ is called the *trivialization homeomorphism*, or simply as the *trivialization of* p *over* U.

A *fiber bundle with fiber* F is a map $p\colon E \longrightarrow B$ such that the set

$$\{U \overset{\text{open}}{\subseteq} B \mid p \text{ is locally trivial on } U \text{ with fiber } F\}$$

is an open cover. This is the *trivialization cover* of p. When the base B is clear from context, the fiber bundle is often referred to simply by naming

[a] The category of topological spaces is not good for doing many of the category-theoretic constructions we desire, as it is not cartesian-closed. The usual solution (and the one we take here) is to restrict to a "good" notion of space, which will be cartesian-closed. For an in-depth discussion of why this is necessary, as well as other types of good categories, see for example `https://ncatlab.org/nlab/show/convenient+category+of+topological+spaces`.

E and omitting p from the notation. The space B is the *base space*, the space E is the *total space*, and the map p is called the *structure map*. By an abuse of notation, when p is clear from context we sometimes say that "E is a fiber bundle over B."

We sometimes say F*-bundle* instead of "fiber bundle with fiber F." When $F = S^n$ and n is clear from context, it is also called a *sphere bundle*.

Definition 1.7. Let $p: E \longrightarrow B$ and $p': E' \longrightarrow B$ be two fiber bundles over B. A morphism of fiber bundles $f: p \longrightarrow p'$ is given by a map $E \longrightarrow E'$ (also by abuse of notation generally named f) such that

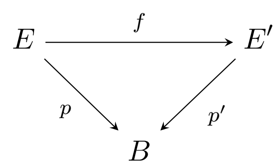

commutes. Compositions of morphisms of fiber bundles is done via compositions of the maps on total spaces.

For all $b \in B$, a morphism f of fiber bundles restricts to a map $f_b: p^{-1}(b) \longrightarrow (p')^{-1}(b)$.

By an abuse of notation, we sometimes denote a morphism of fiber bundles f by $f: E \longrightarrow E'$, leaving the structure over B implied.

Two fiber bundles p and p' are *isomorphic* if there exist morphisms of fiber bundles $f: p \longrightarrow p'$ and $g: p' \longrightarrow p$ such that $g \circ f$ and $f \circ g$ are the identities on p and p', respectively.

Some important examples of fiber bundles:

Example 1.8. The *trivial bundle* over B with fiber F is the bundle $B \times F \longrightarrow B$ where the map is just projection onto the first coordinate. When B is clear from context, this bundle is denoted ϵ^F.

A bundle which is isomorphic to a trivial bundle is called *trivializable*.

Example 1.9. Let M be a smooth manifold smoothly embedded into $\mathbf{R}^N$ for some N. The *tangent bundle* $p: TM \longrightarrow M$ is the space

$$TM \stackrel{\text{def}}{=} \left\{(x, v) \in M \times \mathbf{R}^N \;\middle|\; \text{the tranlation of } v \text{ to } x \text{ is tangent to } M\right\}$$

together with the projection forgetting the second coordinate. The *sphere bundle of* TM is the space

$$S(TM) \stackrel{\text{def}}{=} \{(x, v) \in TM \mid |v| = 1\}.$$

The *disk bundle of TM* is the space

$$D(TM) \overset{\text{def}}{=} \{(x, v) \in TM \mid |v| \leq 1\}.$$

$S(TM)$ is a closed manifold, while $D(TM)$ is a manifold with boundary (whose boundary is actually $S(TM)$). The space TM is also a manifold, but is clearly not compact.

For another example of a sphere bundle on a manifold, see Exercise 3.5.

Example 1.10. Let M be a smooth manifold smoothly embedded in $\mathbf{R}^N$. The *normal bundle* of M is the set of points (x, v) with $x \in M$ and v orthogonal to M at x, defined analogously to the tangent bundle. As above, this is naturally a bundle by projecting onto the M-coordinate.

With this new definition in mind, we can rephrase the Hairy Ball Question in terms of fiber bundles.

Definition 1.11. A *section* of a fiber bundle $p: E \longrightarrow B$ is a map $s: B \longrightarrow E$ such that $p \circ s = 1_B$.

The Hairy Ball Theorem states that $S(TS^2)$ does not have a section. On the other hand, we can see by inspection that $S(TS^1) \cong S^1 \amalg S^1$, with the projection down being the fold map, and thus has a section. It turns out that the answer is surprisingly simple:

Theorem 1.12 (Poincare–Hopf)**.** *For a smooth closed manifold M, the bundle $S(TM)$ has a section only if the Euler characteristic of M is 0.*

See Theorem 5.10.

The section in $S(TS^1)$ provides an isomorphism of TS^1 with the trivial line bundle on S^1. Using a similar construction we can give a section of $S(TS^3)$, and by working carefully we can extend it to show that TS^3 is also trivial. The construction can easily be generalized to produce a section of $S(TS^{2n-1})$ for all n. We might be tempted to conjecture, motivated by these observations, that the tangent bundles of odd spheres are all trivial.

Definition 1.13. A smooth manifold whose tangent bundle is trivial is called *parallelizable*.

The actual answer turns out to be rather surprising:

Theorem 1.14 (Hopf invariant 1)**.** *The only parallelizable spheres are S^0, S^1, S^3 and S^7.*

See Section 10.2.

The goal of this book is to develop most of the material needed to prove these two theorems. The discussion of the first one is completed at the end of Chapter 5,[b] but it will take the rest of the book to prove the second.

An observation about the parallelizable spheres: these are exactly the unit spheres in $\mathbf{R}$, $\mathbf{C}$, $\mathbf{H}$, and $\mathbf{O}$, the real, complex, Hermitian, and octonion numbers. The multiplicative structure in these examples can be used to construct the trivialization of the tangent bundle, so it may not be surprising that these are the examples which are easy to construct. In fact, the Hopf invariant 1 Theorem implies the following:

Theorem 1.15. *Suppose that $\mathbf{R}^n$ is equipped with a bilinear multiplication $\mu\colon \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}^n$ with no zero divisors. Then $n = 1, 2, 4$ or 8.*

Recall that an *H-space* is a topological space X equipped with a function $\mu\colon X \times X \longrightarrow X$. Every topological group is an H-space. The spheres S^0, S^1, S^3 and S^7 are H-spaces (and the first three are groups), and one may ask which spheres have H-space structures (and thus which can have group structures). The answer turns out to be that it is only these examples—and that this is the same as the parallelizability question.

Adams has a famous chart of these types of properties, which we reproduce in Figure 1.1. We will mostly be taking this diagram for granted. In Theorem 6.25 we show that a bilinear product without zero divisors implies that $\mathbf{R}P^n$ (and therefore, S^n) is parallelizable. H-spaces and some of their properties are discussed in Section 8.3. Elements of Hopf invariant 1, and their nonexistence above dimension 8, are discussed in Section 10.2.

[b]A key component of the proof is unfortunately omitted, as it requires results from differential geometry that are beyond the scope of the course.

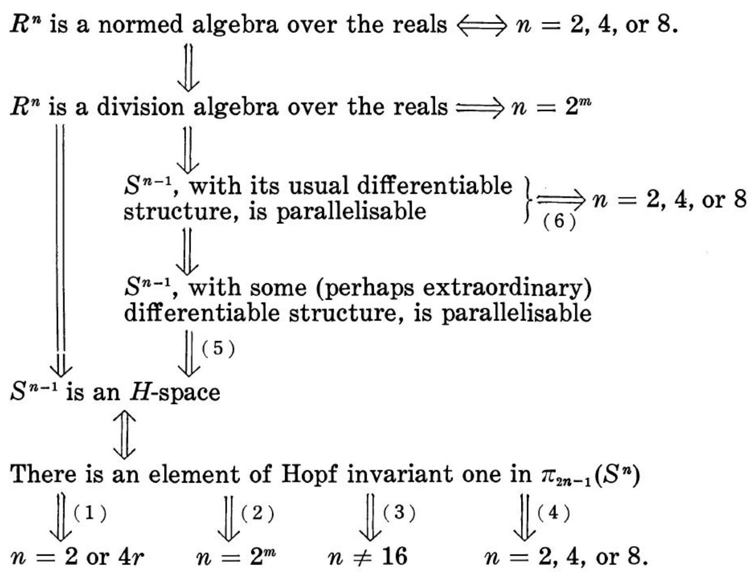

Figure 1.1: Adams's diagram on the Hopf Invariant 1 problem; [Ada60]

Chapter 2

Formal constructions and fundamental examples

This chapter introduces two of our most important players: the formal definition of a vector bundle, and the example of the Grassmannian. Vector bundles are our basic objects of study in the book, and this chapter contains the basic ideas about their construction and analysis.

We also introduce several important examples of spaces and vector bundles, most notably the Stiefel manifolds and the Grassmannians of planes. These are examples of *moduli spaces*: spaces in which each point represents a geometric object. In our case, each point of a Grassmannian represents a subspace of a fixed vector space; each point of a Stiefel manifold represents a choice of orthogonal basis for a subspace. These inherit a topology from that of the vector space, and their geometry will form the basis of most of the analysis in this book. Each Grassmannian is equipped with a *canonical bundle*: since each point represents a vector space, it is possible to construct a new space by "replacing each point with the space it represents."

This chapter is organized as follows. Section 2.1 defines fiber bundles and introduces some of their theory. Section 2.2 defines Stiefel manifolds and their properties. Section 2.3 defines a special class of fiber bundles called vector bundles, and explains how the general results in Section 2.1 specialize to this case. Section 2.4 introduces Grassmannians and some of their properties.

2.1 Constructions using local data

In Example 1.9 we defined the tangent bundle by looking at individual tangent planes for a specific embedding of the manifold. This is not wholly satisfying, as we would like the tangent bundle to exist independently of any embedding, the way that the tangent space at a point of a manifold exists independently of an embedding. (In contrast, the normal bundle cannot exist without an embedding, as even the dimension of each fiber depends fundamentally on the choice of embedding.) In order to give a better construction of the tangent bundle, we will need the notion of a mapping space.

Definition 2.1. Let X and Y be spaces. The *mapping space* $\mathrm{Map}(X,Y)$ has as its underlying set the set of all maps $X \longrightarrow Y$, with the topology given by the compact-open topology. A subbase for this topology is given by the sets of maps

$$V(K,U) = \{f \in \mathrm{Hom}(X,Y) \mid f(K) \subseteq U\}.$$

Here, K is a compact subset of X and U is an open subset of Y.

Lemma 2.2. *There is a natural bijection*

$$\left\{ \begin{array}{c} \textit{maps} \\ f\colon U \times F \longrightarrow U \times F' \\ \textit{s.t. } \mathrm{proj}_1 \circ f = \mathrm{proj}_1 \end{array} \right\} \rightsquigarrow \left\{ \begin{array}{c} \textit{maps} \\ f'\colon U \longrightarrow \mathrm{Map}(F,F') \end{array} \right\}$$

$$f \longmapsto (u \longmapsto \mathrm{proj}_2 \circ f|_{u \times F})$$

In other words, this lemma states that we can think of a map $U \times F \longrightarrow U \times F'$ as on the left-hand side as a "continuous family of maps" parametrized by U.

Proof. Let $f\colon U \times F \longrightarrow U \times F'$ satisfy $\mathrm{proj}_1 \circ f = \mathrm{proj}_1$; this condition implies that

$$f(u,x) = (u, \widetilde{f}_2(u,x)), \qquad \text{where } f_2\colon U \times F \longrightarrow F'.$$

The map f_2 determines and is uniquely determined by a map

$$\widetilde{f}\colon U \longrightarrow \mathrm{Map}(F,F').$$

This map $\widetilde{f}$ is given by precisely the formula in the statement of the lemma. A fundamental property of spaces is that this function will be continuous if and only if f was continuous.[a] This proves the bijection in the first part of the lemma. $\square$

Using this lemma we can build bundles by patching together trivial bundles.

Example 2.3. Let $p\colon E \longrightarrow B$ be a fiber bundle with fiber F. Write $\{U_\alpha\}_{\alpha\in A}$ for the trivialization cover and $\varphi_\alpha\colon p^{-1}(U_\alpha) \longrightarrow U_\alpha \times F$ for the trivialization homeomorphism. For every pair $(\alpha, \beta) \in A^2$ the homeomorphisms φ_α and φ_β induce a composite homeomorphism

$$(U_\alpha \cap U_\beta) \times F \xrightarrow{\varphi_\alpha^{-1}} p^{-1}(U_\alpha \cap U_\beta) \xrightarrow{\varphi_\beta} (U_\alpha \cap U_\beta) \times F$$

which is the identity after projection to the first coordinate and a homeomorphism on each fiber. By Lemma 2.2, this determines a map

$$g_{\alpha\beta}\colon U_\alpha \cap U_\beta \longrightarrow \mathrm{Map}(F, F).$$

This satisfies:

- $g_{\alpha\alpha}$ is uniformly the identity.
- $g_{\alpha\beta}(x) = g_{\beta\alpha}(x)^{-1}$ for all $x \in U_\alpha \cap U_\beta$.
- $g_{\alpha\beta}(x) \circ g_{\beta\gamma}(x) \circ g_{\gamma\alpha}(x) = 1_F$ for all $x \in U_\alpha \cap U_\beta \cap U_\gamma$.

On the other hand, consider a collection of g's which satisfy the above conditions. Then we can assemble a bundle $p\colon E \longrightarrow B$ by taking

$$E = \coprod_\alpha U_\alpha \times F/\sim,$$

where for any $x \in U_\alpha \cap U_\beta$ we say that $(x, v) \sim (x, g_{\alpha\beta}(x) \cdot v)$. The conditions above exactly state that $\sim$ is an equivalence relation, and the continuity conditions on the φ_α are enforced because each $g_{\alpha\beta}$ is continuous. We leave the proof that this is well-defined to Exercise 3.4.

This last example is a procedure that is often seen in mathematics. We take an object that we understand ($\mathbf{R}^n$, $\mathbf{C}^n$, trivial bundle, ring) "glue" a whole bunch of them together in a nice way, and produce a new object ((real/complex) manifold, vector bundle, scheme) which is more general and interesting, while still retaining many of the properties of the simpler object.

[a] It is *not* true for general topological spaces.

Example 2.4 (Möbius bundle). The Möbius bundle is a bundle on S^1 using the cover $U_1 = S^1 \setminus \{\text{north}\}$ and $U_2 = S^1 \setminus \{\text{south}\}$. Define the function

$$g_{12}\colon S^1 \setminus \{\text{poles}\} = U_1 \cap U_2 \longrightarrow \mathrm{GL}_1(\mathbf{R})$$

by letting it be -1 on the part of S^1 with negative x-coordinate and 1 on the part of S^1 with positive x-coordinate.

We can visualize this by taking two strips of paper and drawing a line down the middle of each. This line is the original circle S^1 visualized as $U_1 \times \{0\}$ and $U_2 \times \{0\}$; for any point $x \in S^1$, the line in the paper perpendicular to S^1 is the copy of $\mathbf{R}^1$ sitting above x. The gluing function above says to glue two ends of the paper together "correctly", but to flip the other end over before gluing them together—exactly the construction of a Möbius strip.

Example 2.5. The approach in Example 2.3 can also be used to give an intrinsic construction of the tangent bundle. Let M be a smooth n-manifold. Fix any smooth atlas $\{(U_\alpha, \phi_\alpha\colon U_\alpha \longrightarrow \mathbf{R}^n)\}$ on M. For any α, β, define

$$f_{\alpha\beta}\colon \phi_\alpha(U_\alpha \cap U_\beta) \xrightarrow{\phi_\beta\phi_\alpha^{-1}} \phi_\beta(U_\alpha \cap U_\beta)$$

to be the transition map from α to β. As both the domain and codomain are subsets of $\mathbf{R}^n$ we can define

$$g_{\alpha\beta}\colon U_\alpha \cap U_\beta \longrightarrow GL_n(\mathbf{R}) \subseteq \mathrm{Map}(\mathbf{R}^n, \mathbf{R}^n) \qquad \text{by} \qquad g_{\alpha\beta}(x) \overset{\text{def}}{=} df_{\phi_\alpha(x)}.$$

This is continuous by definition. To check that this produces a vector bundle it suffices to check conditions (1)-(3) in Example 2.3. Condition (1) holds automatically, because $f_{\alpha\alpha}$ is the identity map. Condition (2) holds by the inverse function theorem, since $d(f^{-1})_{f(x)} = (df_x)^{-1}$. By the definition of the transition maps,

$$f_{\alpha\beta} \circ f_{\beta\gamma} \circ f_{\gamma\alpha} = 1_{\mathbf{R}^n} \qquad \text{for all } x \in U_\alpha \cap U_\beta \cap U_\gamma.$$

Fix $x \in U_\alpha \cap U_\beta \cap U_\gamma$ and write $x_\alpha \overset{\text{def}}{=} \phi_\alpha(x)$ (and analogously for β and γ). The chain rule tells us that, in $GL_n(\mathbf{R})$,

$$\begin{aligned} 1_{\mathbf{R}^n} &= d1_{x_\gamma} = d(f_{\alpha\beta} \circ f_{\beta\gamma} \circ f_{\gamma\alpha})_{x_\gamma} \\ &= d(f_{\alpha\beta})_{x_\alpha} d(f_{\beta\gamma})_{x_\beta} d(f_{\gamma\alpha})_{x_\gamma} = g_{\alpha\beta}(x) g_{\beta\gamma}(x) g_{\gamma\alpha}(x), \end{aligned}$$

which is exactly condition (3).

Our next goal is to classify fiber bundles up to isomorphism. A useful lemma for constructing isomorphisms of fiber bundles is the following:

Lemma 2.6. *Let* $p\colon E \longrightarrow B$ *and* $p'\colon E' \longrightarrow B$ *be two fiber bundles with fibers* F *and* F'. *If* $f\colon E \longrightarrow E'$ *is an isomorphism on bundles then for all* $b \in B$ *the restriction* f_b *is a homeomorphism. Conversely, if* F *and* F' *are compact Hausdorff spaces, then if for all* $b \in B$ *the restriction* f_b *is a homeomorphism then* f *is an isomorphism of fiber bundles.*

The condition on fibers in the second half of the lemma is necessary because in general inversion of homeomorphisms is not continuous with respect to the compact-open topology. For more general fibers a more complicated condition is necessary, as the proof fundamentally uses this continuity.

Proof. If a morphism of fiber bundles is an isomorphism then it has a fiberwise inverse because it has an inverse which is a morphism of fiber bundles.

Now suppose that f is a homeomorphism on each fiber and the fibers are compact Hausdorff spaces. Let $g\colon E' \longrightarrow E$ be given on each fiber by the inverse of the restriction of f. In other words, $g|_{(p')^{-1}(x)} = (f|_{p^{-1}(x)})^{-1}$. This function is well-defined and is the inverse of f pointwise, so it suffices to check that it is gives a well-defined morphism of fiber bundles. It is compatible with the projection maps by definition, so the only thing left to check is continuity.

First, consider the case when E and E' are trivial. By Lemma 2.2, g corresponds to a function $B \longrightarrow \mathrm{Homeo}(F', F) \subseteq \mathrm{Map}(F', F)$ and is continuous if and only if this function is. But by definition, this function is exactly the composition

$$B \xrightarrow{\widetilde{f}} \mathrm{Homeo}(F, F') \xrightarrow{\cdot^{-1}} \mathrm{Homeo}(F', F),$$

and is therefore continuous (since for compact Hausdorff spaces the homeomorphism group is a topological group [Are46, Theorem 3]). Thus g is continuous, as desired.

We now consider the general case. Let U be any open in the trivialization cover, and consider the following diagram:

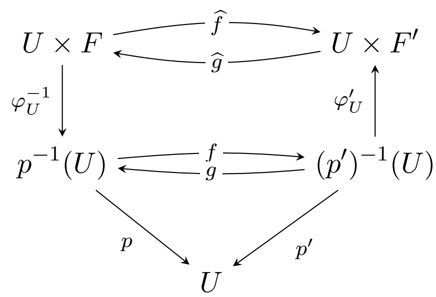

By the special case, $\widehat{g}$ is continuous. Thus g induces a map

$$(p')^{-1}(U) \longrightarrow p^{-1}(U).$$

Since continuity is a local property and U was arbitrary, it follows that g is continuous everywhere, as desired. □

Example 2.7. Let $B = S^1$ and consider TS^1. A point in TS^1 is a point in S^1 together with a vector tangent to S^1 at that point. In other words, we can write a point of TS^1 as a point $(\cos\theta, \sin\theta)$ and a vector $(-\lambda\sin\theta, \lambda\cos\theta)$. There is a map $TS^1 \longrightarrow \mathbf{R}^3$ given by

$$((\cos\theta, \sin\theta), (-\lambda\sin\theta, \lambda\cos\theta)) \longmapsto (\cos\theta, \sin\theta, \lambda).$$

The image of this map is exactly $S^1 \times \mathbf{R}$, and this gives an isomorphism between TS^1 and a trivial bundle on S^1.

2.2 Stiefel manifolds

One of the useful features of fiber bundles is that the homotopy groups of the fiber, the base, and the total space are closely related. If sequences of the form

$$A \hookrightarrow X \longrightarrow X/A$$

are the fundamental building blocks for spaces for the purposes of homology and cohomology, then fiber bundles are such fundamental building blocks from the point of view of homotopy groups. In particular, there is an important theorem about the relationship between these groups:

Theorem 2.8 ([Hat02, Theorem 4.41, Proposition 4.48])**.** *Let $E \longrightarrow B$ be a fiber bundle with fiber F. There is a long exact sequence of homotopy groups*

$$\cdots \longrightarrow \pi_n F \longrightarrow \pi_n E \longrightarrow \pi_n B \longrightarrow \pi_{n-1} F \longrightarrow \cdots \longrightarrow \pi_1 F \longrightarrow \pi_1 E \longrightarrow \pi_1 B.$$

For a suitable definition of "exact" this can be continued to the sets of connected components; in particular, if F is path-connected then a 0 can be appended on the right.

In the rest of this section we introduce *Stiefel manifolds*, which will be fundamental to our construction of universal bundles, and use this theorem to compute some of their homotopy groups.

Definition 2.9. We write

$$\mathbf{R}^\infty \stackrel{\text{def}}{=} \bigcup_{n\geq 0} \mathbf{R}^n.$$

In other words, $\mathbf{R}^\infty$ is the vector space of infinite tuples in which all but finitely many entries are 0.

Definition 2.10. Let n be a nonnegative integer and k be any nonnegative integer or ∞. The *Stiefel manifold* $V_n(\mathbf{R}^k)$ is the set of orthogonal n-frames of $\mathbf{R}^k$: the points of it are ordered n-tuples of orthonormal vectors in $\mathbf{R}^k$. We can think of $V_n(\mathbf{R}^k)$ as a subset of $(S^{k-1})^n$; it inherits its topology from this space. As this is a closed subspace of a compact space, it is compact. When $k < \infty$, $V_n(\mathbf{R}^k)$ is a manifold.

Example 2.11. If $n > k$ then $V_n(\mathbf{R}^k)$ is empty. We also have

$$V_n(\mathbf{R}^n) \cong O(n) \qquad \text{and} \qquad V_1(\mathbf{R}^k) \cong S^{k-1}.$$

There is a map $V_n(\mathbf{R}^k) \longrightarrow S^{k-1}$ given by projecting an n-tuple $(v_1, \ldots, v_n)$ to v_n.

Lemma 2.12. *$p\colon V_n(\mathbf{R}^k) \longrightarrow S^{k-1}$ is a fiber bundle with fiber $V_{n-1}(\mathbf{R}^{k-1})$.*

Proof. Fix a point $x \in S^{k-1}$. The preimage of this point is orthonormal n-tuples $(v_1, \ldots, v_n)$ with $v_n = x$. This means that $v_1, \ldots, v_{n-1} \in v_n^\perp \cong \mathbf{R}^{k-1}$ and form an orthonormal tuple. Conversely, any such orthonormal tuple gives a point in the preimage of x. This identifies the fiber with $V_{n-1}(\mathbf{R}^{k-1})$.

To prove that it is a fiber bundle it suffices to show that there exists an open cover of S^{k-1} by subsets on which the bundle is trivial. Let U_i^+ be the subset of S^{k-1} of those points with positive x_i-coordinate. The preimage $p^{-1}(U_i)$ contains those orthonormal tuples $(v_1, \ldots, v_n)$ where the i-th coordinate of v_n is positive. Let

$$c_i(v) \stackrel{\text{def}}{=} v - (v \cdot e_i)e_i.$$

For all points inside $p^{-1}(U_i)$, the only vector that can equal e_i is v_n, so for all other coordinates $c_i(v_j)$ will be linearly independent. Define

$$\varphi\colon p^{-1}(U_i) \longrightarrow U_i \times V_{n-1}(\mathbf{R}^{k-1})$$

by

$$\varphi(v_1,\ldots,v_n) = (v_n, GS(c_i(v_1),\ldots,c_i(v_{n-1}))) \in U_i^+ \times V_{n-1}(\mathbf{R}^{k-1}),$$

where GS is the Gram–Schmidt process. Since the U_i form a cover, this proves that p is a fiber bundle. □

Proposition 2.13. *For* $1 < m < k-2$*,*

$$\pi_m V_{n-1}(\mathbf{R}^{k-1}) \cong \pi_m V_n(\mathbf{R}^k).$$

In particular, when $k > m+n$*,*

$$\pi_m V_n(\mathbf{R}^k) = 0.$$

Proof. Since $V_n(\mathbf{R}^k) \longrightarrow S^{k-1}$ is a fiber bundle with fiber $V_{n-1}(\mathbf{R}^{k-1})$, there exists a long exact sequence in homotopy groups

$$\cdots \longrightarrow \pi_{m+1}S^{k-1} \longrightarrow \pi_m V_{n-1}(\mathbf{R}^{k-1}) \longrightarrow \pi_m V_n(\mathbf{R}^k) \longrightarrow \pi_m S^{k-1} \longrightarrow \cdots.$$

The homotopy groups of S^{k-1} are 0 for $m \le k-2$ and $\mathbb{Z}$ for $m = k-1$ (the homotopy groups $\pi_m S^{k-1}$ are not, in general, known for $m > k$). When $1 < m < k-2$ the above fragment of the long exact sequence is

$$0 \longrightarrow \pi_m V_{n-1}(\mathbf{R}^{k-1}) \longrightarrow \pi_m V_n(\mathbf{R}^k) \longrightarrow 0;$$

thus $\pi_m V_{n-1}(\mathbf{R}^{k-1}) \cong \pi_m V_n(\mathbf{R}^k)$. By iterating this and remembering the assumption that $k > n+m$, it follows that

$$\pi_m V_n(\mathbf{R}^k) \cong \pi_m V_1(\mathbf{R}^{k-n+1}) = \pi_m S^{k-n} = 0.$$

□

Thus, within a particular range, $\pi_m V_n(\mathbf{R}^k)$ depends only on $k-n$. It turns out that this range is good enough to give good behavior on the infinite-dimensional Stiefel manifold.

There is a sequence of inclusions induced by the standard inclusion $\mathbf{R}^n \subseteq \mathbf{R}^{n+1} \subseteq \cdots$

$$V_n(\mathbf{R}^n) \subseteq V_n(\mathbf{R}^{n+1}) \subseteq \cdots.$$

Write

$$V_n \overset{\text{def}}{=} V_n(\mathbf{R}^\infty) \cong \operatorname*{colim}_{k\to\infty} V_n(\mathbf{R}^k).$$

Theorem 2.14. *The map $V_n \longrightarrow *$ taking V_n to a point is an isomorphism on all homotopy groups. (I.e. V_n is* weakly contractible.*)*

Proof. To check connectedness it suffices to describe a path connecting any two orthonormal n-frames $(v_1, \ldots, v_n)$ and $(v'_1, \ldots, v'_n)$, in $\mathbf{R}^\infty$. As $\mathbf{R}^\infty$ is a colimit, we can assume that these n-frames both sit inside some $\mathbf{R}^k$, with k at least $n+1$. Let $g \in O(k)$ be such that $g \cdot v_i = v'_i$ and $\det g = 1$; this is possible by our assumption on k. The group $O(k)$ has two connected components, consisting of those elements with determinant 1, and those with determinant -1. (See Exercise 2.5.) Since g has determinant 1 there exists a path $\gamma\colon I \longrightarrow O(k)$ with $\gamma(0) = 1$ and $\gamma(1) = g$. The path $\gamma'\colon I \longrightarrow V_n(\mathbf{R}^k)$, defined by

$$\gamma'(t) \stackrel{\text{def}}{=} (\gamma(t) \cdot v_1, \ldots, \gamma(t) \cdot v_n)$$

has $\gamma'(0) = (v_1, \ldots, v_n)$ and $g'(1) = (v'_1, \ldots, v'_n)$. Thus in $V_n(\mathbf{R}^k)$ there is a path connecting the two n-frames, and this must therefore hold in V_n, as well, as desired.

We turn our attention to showing that $\pi_m V_n = 0$ for $m \geq 1$. An element in $\pi_m V_n$ is represented by a homotopy class of maps $S^m \longrightarrow V_n = \bigcup_{k=n}^\infty V_n(\mathbf{R}^k)$. Since S^m is compact this map must factor[b] through the inclusion $V_n(\mathbf{R}^k) \longrightarrow V_n$ for some k, which we assume to be at least $m+n$. Thus the map factors as

$$S^m \longrightarrow V_n(\mathbf{R}^k) \longrightarrow V_n.$$

Since $\pi_m V_n(\mathbf{R}^k) = 0$ for $k > m+n$, it follows that $\pi_m V_n = 0$, as desired. □

There is a left action of $O(n)$ on $V_n(\mathbf{R}^k)$ for all k, which can be defined in the following way. A point in $V_n(\mathbf{R}^k)$ is an n-tuple of vectors in $\mathbf{R}^k$. We can write these as an $n \times k$ matrix. $O(n)$ then acts on this by multiplication on the left.

Definition 2.15. The map

$$\tau_V\colon V_n(\mathbf{R}^k) \longrightarrow V_{n+1}(\mathbf{R}^{k+1})$$

is given by adding the vector $(0, \ldots, 0, 1)$ to each n-frame. This is $O(n)$-equivariant, if we think of having $O(n)$ act on $V_{n+1}(\mathbf{R}^{k+1})$ by including it into $O(n+1)$ as the transformations that fix the last coordinate.

[b] This fact is actually more complicated than we pretend above. [Hov99, Proposition 2.4.2] shows that this fact holds for compact spaces assuming that the sequential limit is given by closed inclusions of Hausdorff spaces, which is the case here.

2.3 Vector Bundles

Fiber bundles are useful, but they are quite difficult to classify because they do not have enough structure to work with. We showed in Lemma 2.6 that isomorphisms between two fiber bundles $E \longrightarrow B$ and $E' \longrightarrow B$ with fiber F correspond to maps $B \longrightarrow \mathrm{Homeo}(F, F)$. But $\mathrm{Homeo}(F, F)$ is quite a large space, and the compact-open topology is difficult to work with, so analyzing it can be extremely difficult even for simple F. (Consider, for example, how large the space $\mathrm{Homeo}(\mathbf{R}^k, \mathbf{R}^k)$ is.) In many situations allowing all homeomorphisms is overkill, and a restriction to a smaller subspace produces many examples of interest. For example, we can put some kind of algebraic structure on the fiber and demand that all morphisms between bundles respect this structure. The general case of this type of construction is rich and varied, and far beyond the scope of this course so we focus on one type of example.

Definition 2.16. Let $p\colon E \longrightarrow B$ be a fiber bundle with fiber $\mathbf{R}^k$, together with a choice of $\mathbf{R}$-vector space structure on each fiber. A local trivialization map $\varphi\colon p^{-1}(U) \longrightarrow U \times \mathbf{R}^k$ is *fiberwise-linear* if on each fiber it is a linear map. If a trivialization map compatible with the linear structure exists on an open set U, we say that p is *linear* on U.

A *vector bundle of rank k* is a fiber bundle $p\colon E \longrightarrow B$ with fiber $\mathbf{R}^k$, together with a choice of $\mathbf{R}$-vector space structure on each fiber, such that the set

$$\{U \overset{\text{open}}{\subseteq} B \mid p \text{ is locally linear on } U\}$$

is an open cover of B. We sometimes use the term *k-bundle* instead of "vector bundle of rank k."

A *morphism of vector bundles* is a morphism of fiber bundles which restricts to a linear map on each fiber.

Examples 2.4 and 2.5 were both examples of vector bundles. Many of the proofs in Section 2.1 generalize immediately to the case of vector bundles. We state the important ones here, without proof. Instead of mapping spaces appearing we instead get spaces of linear maps or general linear groups, making the corresponding analyses significantly simpler.

Lemma 2.17. *There is a natural bijection*

$$\left\{ \begin{matrix} \textit{fiberwise-linear maps} \\ f\colon U \times \mathbf{R}^k \longrightarrow U \times \mathbf{R}^{k'} \\ s.t.\ \mathrm{proj}_1 \circ f = \mathrm{proj}_1 \end{matrix} \right\} \quad \leftrightsquigarrow \quad \left\{ \begin{matrix} \textit{maps} \\ f'\colon U \longrightarrow \mathrm{Lin}(F, F') \end{matrix} \right\}$$

$$f \qquad \longmapsto \qquad (u \longmapsto \mathrm{proj}_2 \circ f|_{u \times F})$$

Corollary 2.18. *A morphism of vector bundles is an isomorphism if and only if it is a linear isomorphism on each fiber.*

Example 2.19. Given an open cover $\{U_\alpha\}_{\alpha \in A}$ of B, pick a collection of maps $g_{\alpha\beta}\colon U_\alpha \cap U_\beta \longrightarrow GL(\mathbf{R}^k)$ satisfying the three conditions

- $g_{\alpha\alpha}$ is uniformly the identity.
- $g_{\alpha\beta}(x) = g_{\beta\alpha}(x)^{-1}$ of all $x \in U_\alpha \cap U_\beta$.
- $g_{\alpha\beta}(x) g_{\beta\gamma}(x) g_{\gamma\alpha}(x) = 1$ for all $x \in U_\alpha \cap U_\beta \cap U_\gamma$.

Then there is a vector bundle $p\colon E \longrightarrow B$ of rank k constructed by taking

$$E = \coprod_{\alpha \in A} U_\alpha \times \mathbf{R}^k / (x, v) \sim (x, g_{\alpha\beta}(x), v) \qquad \text{for } x \in U_\alpha \cap U_\beta.$$

Example 2.20. The construction of the tangent bundle as a fiber bundle in Example 2.5 gives it the structure of a vector bundle.

Unlike general fiber bundles, every vector bundle has at least one section: the section which takes each point x to the origin in the vector space in the preimage of x. (This is called the *zero section.*) The zero section is preserved by any morphism of vector bundles, and can therefore be used to make some basic distinctions between bundles.

Lemma 2.21. *TS^1 is not isomorphic to the Mobius bundle.*

Proof. Consider $TS^1 \backslash s_0(S^1)$ and Mobius$\backslash s_0(S^1)$. Since TS^1 is trivializable, this is isomorphic to $S^1 \times (\mathbf{R}^1 \backslash \{0\})$, which is not connected. However, Mobius$\backslash s_0(S^1)$ is connected (as we all know from cutting a Mobius band down the middle). Thus these are not homeomorphic. □

Using the vector space structure on the fibers also allows us to determine, from a finite amount of information, when a vector bundle is trivializable.

Proposition 2.22. *Let $p\colon E \longrightarrow B$ be an n-dimensional vector bundle. There exist n sections $s_1, \ldots, s_n\colon B \longrightarrow E$ such that for all $x \in B$, $s_1(x), \ldots, s_n(x)$ are linearly independent if and only if E is trivializable.*

Proof. If E is isomorphic to $B \times \mathbf{R}^n$ then we can define $s_i\colon B \longrightarrow E$ by $s_i(b) = (b, e_i)$ for a fixed basis $e_1, \ldots, e_n$ of $\mathbf{R}^n$. Then $s_1, \ldots, s_n$ are linearly independent at each point.

Conversely, suppose that the sections $s_1, \ldots, s_n$ exist. Then we define a bundle morphism $f\colon B \times \mathbf{R}^n \longrightarrow E$ in the following manner. For a point $(b, (a_1, \ldots, a_n)) \in B \times \mathbf{R}^n$, define

$$f\left(b, (a_1, \ldots, a_n)\right) \stackrel{\text{def}}{=} \left(b, \sum_{i=1}^{n} a_i s_i(b)\right).$$

This map is continuous since the s_i are, and a fiberwise isomorphism by definition. Thus it is an isomorphism of vector bundles. □

Corollary 2.23. *The Mobius strip has no everywhere-nonzero section.*

Proof. Suppose that Mobius strip had an everywhere-nonzero section. Then by Proposition 2.22 it would be trivial. However, we just proved that it is not. □

This suggests the following (not terribly interesting) partial answer to the Hairy Ball Question:

Partial Answer. When $\dim M = 1$, an everywhere-nonzero section exists exactly when the tangent bundle is trivializable. Since the circle is the only connected 1-dimensional closed manifold and TS^1 is trivializable, this is true for all 1-dimensional closed manifolds.

2.4 Grassmannians and the universal bundle

We now turn to an example of a vector bundle which will play a significant role in our future discussions.

Definition 2.24. The *Grassmannian* $G_n(\mathbf{R}^k)$ is the quotient of $V_n(\mathbf{R}^k)$ by the action of $O(n)$ defined in Definition 2.15, topologized via the quotient topology. In particular, $G_n(\mathbf{R}^k)$ is also compact. A point $\omega \in G_n(\mathbf{R}^k)$ represents a k-dimensional linear subspace of $\mathbf{R}^k$, so that the Grassmannian can be thought of as a "space of n-planes in $\mathbf{R}^k$." By an abuse of notation, we will sometimes use ω for both the point of $G_n(\mathbf{R}^k)$ and the subset of $\mathbf{R}^k$ that it represents.

Below, we will see that $G_n(\mathbf{R}^k)$ is a manifold when $k < \infty$.

Example 2.25. If $n > k$ then $V_n(\mathbf{R}^k)$ is empty and so is $G_n(\mathbf{R}^k)$. We also have

$$G_n(\mathbf{R}^n) = V_n(\mathbf{R}^n)/O(n) \cong *$$

and

$$G_1(\mathbf{R}^k) = S^{n-1}/O(1) \cong \mathbf{R}P^{n-1}.$$

The action of $O(n)$ on $V_n(\mathbf{R}^k)$ is compatible with the inclusion $V_n(\mathbf{R}^k) \hookrightarrow V_n(\mathbf{R}^{k+1})$, and thus there is an induced inclusion $G_n(\mathbf{R}^k) \hookrightarrow G_n(\mathbf{R}^{k+1})$. As for the Stiefel manifold V_n, we write

$$G_n \overset{\text{def}}{=} G_n(\mathbf{R}^\infty) \cong \operatorname*{colim}_k G_n(\mathbf{R}^k).$$

In all of the examples in Example 2.25, $G_n(\mathbf{R}^k)$ was a manifold; in fact, this will always be the case for finite k.

Lemma 2.26. *$G_n(\mathbf{R}^k)$ is Hausdorff.*

Proof. To prove that $G_n(\mathbf{R}^k)$ is Hausdorff it suffices to show that for any two points $\omega_1, \omega_2 \in G_n(\mathbf{R}^k)$ there exists a continuous function $f\colon G_n(\mathbf{R}^k) \longrightarrow \mathbf{R}$ such that $f(\omega_1) \neq f(\omega_2)$. For any point $p \in \mathbf{R}^k$, let $f_p\colon G_n(\mathbf{R}^k) \longrightarrow \mathbf{R}$ set $f_p(\omega)$ to be the Euclidean distance from p to the plane represented by ω. This is continuous because for any n-frame $(v_1, \dots, v_n) \in V_n(\mathbf{R}^k)$ representing ω,

$$f_p(\omega) = \sqrt{p \cdot p - (p \cdot v_1)^2 - \cdots - (p \cdot v_n)^2}.$$

This function gives the same value on each preimage of ω, so it is a continuous function $G_n(\mathbf{R}^k) \longrightarrow \mathbf{R}$. Now let p be any point in ω_1 which is not in ω_2. Then $f_p(\omega_1) = 0$ but $f_p(\omega_2) \neq 0$, as desired. □

We are ready to prove:

Proposition 2.27. *$G_n(\mathbf{R}^k)$ is a manifold of dimension $n(k-n)$.*

Proof. Fix a point $\nu \in G_n(\mathbf{R}^k)$. This is an n–plane in $\mathbf{R}^k$; write $\nu^\perp$ for its orthogonal complement (of dimension $k-n$). We claim that

$$U_\nu \overset{\text{def}}{=} \{\omega \in G_n(\mathbf{R}^k) \,|\, \omega \cap \nu^\perp = 0\}$$

is a neighborhood of ν which is homoeomorphic to $\mathbf{R}^{n(k-n)}$. We leave checking that U_ν is open to the reader. The fact that $\omega \cap \nu^\perp = 0$ implies that ω is the graph of a linear map $\nu \longrightarrow \nu^\perp$. Such maps are determined by $n \times (k-n)$ matrices, so the set U_ν is isomorphic to $\mathbf{R}^{n(k-n)}$. Since ν was arbitrary these give an atlas of $G_n(\mathbf{R}^k)$, showing that it is a manifold. □

Since $G_n(\mathbf{R}^k)$ is Hausdorff we can try to construct a CW structure on it. This is relatively simple once we figure out the correct cells to look at. Let $p_i: \mathbf{R}^k \longrightarrow \mathbf{R}^i$ be the projection onto the first i coordinates; thus p_k is the identity and p_0 is the trivial map to the point. As i goes from k to 0 the dimension of the image of an n-plane $\omega \in G_n(\mathbf{R}^k)$ drops from n to 0; let σ_i be the smallest integer such that $\dim p_i(\omega) = i$. The sequence $\sigma = (\sigma_1, \dots, \sigma_n)$ is called a *Schubert symbol*. If we let $e(\sigma)$ be the subset of $G_n(\mathbf{R}^k)$ having σ as their Schubert symbold we note that these are spaces whose n-frames, after reducing into Eschelon form, have columns $\sigma_1, \dots, \sigma_n$ as the pivots; all entries which are not pivot columns can have any real number they wish as the entries, so this subspace is homeomorphic to a Euclidean space (open cell) of dimension $\sum_{i=1}^{n}(k - (\sigma_i - 1) - (n - i + 1))$. The maximal value of this is $n(k-n)$. This decomposition is compatible with the inclusion $G_n(\mathbf{R}^k) \hookrightarrow G_n(\mathbf{R}^{k+1})$, and thus in the colimit they induce a CW structure on G_n.

Since a Grassmannian is a space encoding information about vector subspaces it naturally carries with it a vector bundle.

Definition 2.28. The *universal bundle* γ_{nk} is a bundle $p: \gamma_{nk} \longrightarrow G_n(\mathbf{R}^k)$ which has over every point the plane that the point represents. More concretely,

$$\gamma_{nk} = \{(\omega, v) \mid \omega \in G_n(\mathbf{R}^k),\ v \in \omega \subseteq \mathbf{R}^k\} \subseteq G_n(\mathbf{R}^k) \times \mathbf{R}^k.$$

The map $p: \gamma_{nk} \longrightarrow G_n(\mathbf{R}^k)$ is projection onto the first coordinate.

When $k = \infty$ we simply write γ_n instead of $\gamma_{n\infty}$.

The example of universal bundles is fundamental to the classification of vector bundles. We will show (in Chapter 5) that all universal bundles are not trivializable, for now we focus on the case $n = 1$.

Lemma 2.29. *For any $k > 1$ (including $k = \infty$) γ_{1k} is not trivializable.*

Proof. By Proposition 2.22, in order to show that γ_{1k} is nontrivial it suffices to show that it has no everywhere-nonzero sections. Let $s: G_1(\mathbf{R}^k) \longrightarrow \gamma_{1k}$ be any section, and consider the composition $S^{k-1} \longrightarrow \mathbf{R}P^{k-1} = G_1(\mathbf{R}^k) \longrightarrow \gamma_{1k}$ where the first map is the usual double cover. This takes a point x to a pair $(\{\pm x\}, t(x)x)$ for some continuous $t: S^{k-1} \longrightarrow \mathbf{R}$. By definition, $t(-x) = -t(x)$. Since S^{k-1} is connected we must have $t(x_0) = 0$ for some $x_0 \in S^{k-1}$. □

We finish up this section with two important observations about the structure of Grassmannians. There is a natural map $\iota\colon O(n) \longrightarrow O(n+1)$ given by

$$M \longmapsto \begin{pmatrix} M & \\ & 1 \end{pmatrix}.$$

Definition 2.30. The map $\tau_V\colon V_n(\mathbf{R}^k) \longrightarrow V_{n+1}(\mathbf{R}^{k+1})$ is $O(n)$-equivariant (where the action on the codomain is via ι), and in the limit as $k \longrightarrow \infty$ induces a map

$$\tau\colon G_n \longrightarrow G_{n+1}.$$

This map can be directly obtained from ι via the notion of a *classifying space.*

Definition 2.31. Let G be a topological group, and let EG be a weakly contractible space with a free G-action. In other words, the map $EG \longrightarrow *$ is an isomorphism on homotopy groups, and all nonidentity elements of G have no fixed points. The *classifying space* BG is defined to be the quotient of EG by G.

When G is discrete, BG has $\pi_0 BG = *$, $\pi_1 BG = G$ and $\pi_n BG = 0$ for $n > 1$. However, for groups with other topologies this is often not the case. In particular, when $G = O(n)$ it is still an open question what the homotopy groups are. Although the above construction is not intrinsically functorial, it turns out that it can be made functorial.[c]

Theorem 2.32.

$$G_n \simeq BO(n).$$

Proof. This follows directly from the definition of the classifying space $BO(n)$ and Theorem 2.14. □

This result will play a key role in the development of K-theory in Chapter 9. The map τ above is homotopic to the map

$$B\iota\colon BO(n) \longrightarrow BO(n+1);$$

by an abuse of notation we will sometimes refer to $B\iota$ as τ in later parts of the text.

[c]This uses the *bar construction*, and is far beyond the scope of this book. The interested reader can see, for example, [Rie14, Chapter 4].

Further reading

The reader interested in learning more about fibrations, fiber bundles, and long exact sequences of homotopy groups should begin with [Hat02, Chapter 4]. For those interested in advanced computational techniques using these approaches the extra topics in that chapter are especially recommended.

The geometry, topology, and combinatorics of Grassmannians is studied deeply in Schubert calculus, which allows deep computations with the cohomology of general Grassmannians and intersection theory. A recommended article for those interested in these kinds of calculations is [KL72].

Exercises and Extensions

2.1 We have not used any property of $\mathbf{R}$ when defining vector bundles. Thus we could define complex vector bundles in exactly the same way as we defined real vector bundles, but using the structure of complex vector spaces instead of real ones. Which of the examples in this chapter still work? Which do not?

Extend the proofs in this chapter to show that for complex Grassmannians $G_n(\mathbf{C}^\infty) \simeq BU(n)$.

2.2 As a corollary of Lemma 2.2, prove that if F and F' have the structure of vector spaces and $\mathrm{Lin}(F, F')$ is the space of linear maps from F to F', then there is a natural bijection

$$\left\{ \begin{array}{c} \text{maps} \\ f\colon U \times F \longrightarrow U \times F \\ \text{s.t. } \mathrm{proj}_1 \circ f = \mathrm{proj}_1 \text{ and} \\ \mathrm{proj}_2 \circ f|_{u \times F} \text{ linear } \forall\, u \in U \end{array} \right\} \leftrightsquigarrow \left\{ \begin{array}{c} \text{maps} \\ f'\colon U \longrightarrow \mathrm{Lin}(F, F') \end{array} \right\}.$$

2.3 Prove that the definition of the tangent bundle to a smooth manifold given in Definition 1.4 is isomorphic to the construction in Example 2.5, and conclude that the definition from Definition 1.4 is independent of the choice of embedding.

2.4 Prove that

$$V_n(\mathbf{R}^{k+n}) \cong O(k+n)/O(k).$$

Here, the quotient is as a space with a group action, *not* as a quotient of groups.

2.5 The goal of this exercise is to prove that $O(n)$ has two connected components.

(a) Prove that it suffices to show that $SO(n)$ is connected.

(b) Prove that if $H < G$ is a closed subgroup of a topological group and both H and G/H are connected then G is connected.

(c) Consider $SO(n)$ as a topological space with an action of $SO(n-1)$ (via the standard inclusion). Prove that

$$SO(n)/SO(n-1) \cong S^n.$$

(d) Conclude that $SO(n)$ is connected.

Chapter 3

Classification of vector bundles

In order to begin to answer the Hairy Ball Question, we must first develop some tools to work with vector bundles. As is usual for mathematicians, once we define a class of objects we would like a complete classification of examples in that class.

Goal. Given a space B, classify all vector bundles of dimension n over B up to isomorphism.

This goal will be realized in Theorem 3.29 for all compact spaces. In fact, we will be able to do this for all spaces for which there exists a "partition of unity": a good way of weighing locally-defined functions on the space to extend them to the entire space. These exist for all compact spaces and all CW complexes, which are the primary spaces of interest in this book.

The main idea of the classification is that there is exists an object, called the *universal bundle*, which contains all possible vector bundles inside of it (in a precise sense). Any other bundle can be extracted from this universal bundle using a construction called a *pullback*. Such an approach has several formal advantages:

- Pullbacks preserve many structures. This implies that if a property is preserved under pullbacks, it suffices to prove that property for the universal bundle in order to prove it for all bundles.
- Operations on vector bundles (such as addition, multiplication, dualization, etc.) can be cleanly defined on the universal bundle. Appropriate versions of these will tranfer along the pullback construction

and will allow us to define ideas such as addition and multiplication of general bundles.

- The definition of a fiber bundle itself can be rewritten using pullbacks. This rephrasing allows us to easily generalize to structures other than just vector bundles (such as including orientations, further algebraic structures, generalizing the notion of fiber bundle outside of topology, etc.) in order to see how the ideas transfer to other fields of mathematics.

In this chapter, the first section gives an overview of pullbacks, starting with the general categorical definition and its properties, and then using it to restate the definition of a fiber bundle. The reader familiar with pullbacks in general categories can omit this section, with possibly the exception of Definition 3.4. Sections 3.2 and 3.3 state and prove the classification theorem for compact base spaces, and Section 3.4 extends it to well-partitionable (see Definition 3.28) spaces.

3.1 Pullbacks

We begin with the general categorical definition of a pullback:

Definition 3.1. Let $\mathcal{C}$ be a category, and consider a diagram of the form

$$X \xrightarrow{f} Z \xleftarrow{g} Y$$

in $\mathcal{C}$. (Diagrams of this form are sometimes called *cospans*.) The *pullback of this diagram* is a commutative square as on the left

$$\begin{array}{ccc} P & \xrightarrow{f'} & Y \\ {\scriptstyle g'}\downarrow & & \downarrow{\scriptstyle g} \\ X & \xrightarrow{f} & Z \end{array} \qquad\qquad \begin{array}{ccc} A & \xrightarrow{F} & Y \\ {\scriptstyle G}\downarrow & & \downarrow{\scriptstyle g} \\ X & \xrightarrow{f} & Z \end{array}$$

such that for any other square as on the right there exists a unique morphism $\phi\colon A \longrightarrow P$ making the diagram

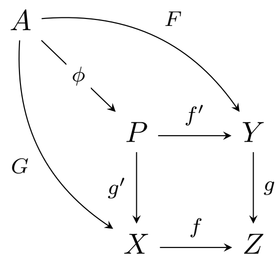

commute. When clear from context, P is often referred to as the *pullback*, with the morphisms down to X and Y omitted.

A commutative square

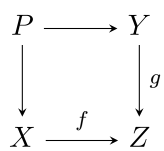

is called a *pullback square* if it is the pullback of the diagram

$$X \xrightarrow{f} Z \xleftarrow{g} Y.$$

This is a particular case of the categorical notion of a "limit." As usual for limits, pullbacks are not unique; any two objects which satisfy the universal property will be the "pullback." However, once both the object and the morphisms in the square are chosen, this becomes unique *up to unique isomorphism*: for any two choices, there is a unique isomorphism between them that is compatible with all of the structure. This implies that for any two choices of pullback there is a unique recipe for translating between them. This is the perspective that we take to justify using the phrase "the pullback" instead of "a pullback."

Example 3.2. Suppose that $\mathcal{C}$ is a partial order, with a unique morphism $A \longrightarrow B$ if $A \leq B$. Then for any cospan

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

the pullback is the greatest lower bound of A and B. This is independent of C—although this is generally not the case (as seen in the next example) it works this way in a partial order because the fact that the pullback square needs to compose does not impose any extra conditions on f' and g'.

Although universal properties are useful, it is also important to have a direct construction that can be used for computation.

Example 3.3. Consider a cospan in any set-based category,[a] such as spaces, groups, graphs, etc:

$$X \xrightarrow{f} Z \xleftarrow{g} Y.$$

The pullback of this diagram (also referred to as the *fiber product*) has underlying set

$$X \times_Z Y \overset{\text{def}}{=} \{(x, y) \in X \times Y \mid f(x) = g(y)\}.$$

(Observe that this is the pullback in the category of sets.) The extra structure on $X \times_Z Y$ (i.e. the topology in spaces, or the group structure in groups, etc.) is inherited from the corresponding structure on $X \times Y$.

Many important constructions are special cases of fiber products. Several examples:

- If Z a single point[b] then $X \times_Z Y \cong X \times Y$. This is because, as in the previous example, the commutativity of the pullback square imposes no extra conditions on f' and g'.

- Suppose that both f and g are injections of sets, so that we can consider X and Y to be subsets of Z. Then the pullback is isomorphic to $X \cap Y$. Indeed, for any $x \in X$, there is a point $y \in Y$ such that $(x, y) \in X \times_Z Y$ exactly if $x \in Y$, and vice versa.

- Suppose that $f\colon X \longrightarrow Z$ is an injection of sets. Then $P \cong g^{-1}(f(X))$, and the morphism $g^{-1}(f(X)) \longrightarrow X$ is just

  $$g|_{g^{-1}(f(X))}\colon g^{-1}(f(X)) \longrightarrow X.$$

  For example, given a fiber bundle $p\colon E \longrightarrow B$ the restriction of the bundle to a subset $S \subseteq B$ can be modeled as the fiber product

  $$\{(s, e) \in S \times E \mid f(y) = s\} \cong f^{-1}(S),$$

  together with the projections onto the two coordinates. The projection onto the second coordinate is exactly $f|_{f^{-1}(S)}$. An important example

[a] More formally, we need a category $\mathcal{C}$ in which the pullback under consideration exists, and which is endowed with a free-forgetful adjunction $F\colon \mathcal{C} \rightleftarrows \mathbf{Set}\colon U$. As these properties hold in most examples of interest, we omit a more detailed discussion; the relevant theorem is that "right adjoints preserve limits"; see for example [Rie17, Section 4.5].

[b] Or more precisely the terminal object of $\mathcal{C}$.

of this is the case where S is a single point, where this is isomorphic to the fiber above S.

This second example motivates the following recasting of the definition of a fiber bundle:

Definition 3.4. A *fiber bundle with fiber* F is a map $p\colon E \longrightarrow B$ for which there exists an open cover $\mathcal{U}$ of B such that for all $U \in \mathcal{U}$ there is a pullback square

$$\begin{array}{ccc} U \times F & \longrightarrow & E \\ {\scriptstyle \mathrm{proj}_1}\big\downarrow & & \big\downarrow{\scriptstyle p} \\ U & \hookrightarrow & B. \end{array}$$

The induced homeomorphism $\varphi_U\colon p^{-1}(U) \longrightarrow U \times F$ is the *trivialization homeomorphism*. Given such a pullback square, we say that p is *trivial over* U.

Pullback squares have an important composition and "partial inverse" property. Consider a diagram

$$\begin{array}{ccccc} A & \longrightarrow & B & \longrightarrow & C \\ \big\downarrow & & \big\downarrow & & \big\downarrow \\ D & \longrightarrow & E & \longrightarrow & F. \end{array}$$

If both of the squares in this diagram are pullback squares, so is the "composed" outer rectangle. In a partial converse to this, if the outside rectangle is a pullback and so is the right-hand square, the left-hand square must also be a pullback. This not only shows that the pullback of the diagram

$$D \longrightarrow E \longleftarrow B$$

exists, but also works as a construction of this pullback. (For an example of this perspective, see the proof of Lemma 3.7.) This simple observation implies an important property for fibers of maps produced using pullbacks.

Lemma 3.5. *Let*

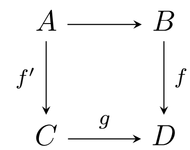

be a pullback square of spaces. For any point $c \in C$*, the fiber of* f' *over* c *is canonically isomorphic to the fiber of* f *over* $g(c)$*. In particular, if the fibers of a map* $f\colon B \longrightarrow D$ *have some associated structure, the fibers of the map* $f'\colon A \longrightarrow C$ *naturally inherit this structure.*

In particular, this lemma implies that if the fibers of f are groups, topological spaces, vector spaces, etc. then the fibers of f' naturally inherit this structure.

Proof. Let $c \in C$ be any point. Write $*$ for the single-point space. Consider the following solid-arrow diagram:

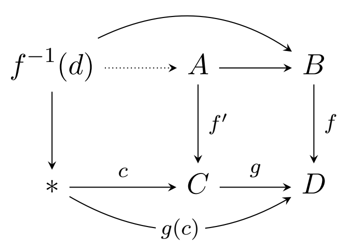

By assumption the right-hand square is a pullback square, as is the outer rectangle. Thus there exists a unique morphism $f^{-1}(d) \longrightarrow A$ (drawn dotted) which makes the diagram commute. Thus the left-hand square is a pullback square, and we see that the fiber over c (which is the pullback of $* \xrightarrow{c} C \xleftarrow{f'} A$) is uniquely isomorphic to the fiber $f^{-1}(d)$. □

We will now introduce two new constructions of fiber bundles. With the above lemma in mind, we can also see that if the original fiber bundles are actually vector bundles then the fibers naturally inherit vector space structures. The proof that these constructions work for general vector bundles is left to the Exercise 3.6.

Definition 3.6. Let $p\colon E \longrightarrow B$ be a fiber bundle with fiber F, and let $f\colon B' \longrightarrow B$ be any map. The *pullback of* p *along* f (also called the *pullback bundle*) is the fiber bundle $p'\colon B' \times_B E \longrightarrow B'$ given by the pullback square

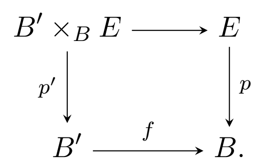

This fiber bundle is often written $f^*p\colon f^*E \longrightarrow B'$; it also has fiber F.

Unlike the definition of pullback, this construction considers the two arguments to be *asymmetric*, emphasizing the relationship of E and B over the relationship of B' and B. The proof that f^*p is a fiber bundle fairly straightforward using the fiber product description of the pullback. However, we present an alternate proof as it illustrates a common technique used in category theory. As an advertisement for category theory, we want to draw the reader's attention to the presence of only one diagram in the proof; all of the text is simply analyzing the properties of the diagram. This is what is known as a "diagram chase" (analogously to an "angle chase" in classical geometry).

Lemma 3.7. *f^*p is well-defined.*

Proof. Let $\mathcal{U}$ be the trivialization cover of p. We claim that the cover $\{f^{-1}(U) \mid U \in \mathcal{U}\}$ works to show that f^*p is a fiber bundle with fiber F. This is a well-defined open cover, since f is continuous. For $U \in \mathcal{U}$, the following black solid-arrow diagram commutes:

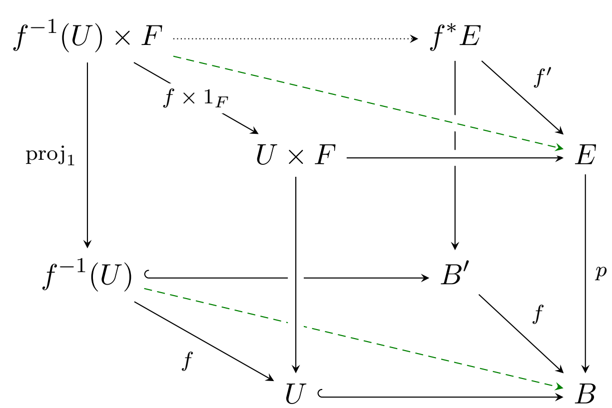

In this diagram, the front face is a pullback square because p is a fiber bundle The left-hand face is also a pullback square. By the composition property of pullback squares, the green-and-black square is also a pullback square.

Since the right-hand square is a pullback square by the definition of f^*E, this implies that the dotted arrow exists and makes the back and top faces commute. By the partial inverse property of pullback squares, the back square must also be a pullback, showing that f^*E is trivial over $f^{-1}(U)$, as desired. □

Example 3.8. Suppose that the map $B' \longrightarrow B$ is an inclusion of a subspace. Then $f^*E \longrightarrow E$ is also an inclusion of a subspace, consisting of exactly the fibers that lie above the points of B'. In other words, $f^*E \cong p^{-1}(B')$. This is also called the *restriction to B' of E*, and is denoted $E|_{B'}$.

Another way to use pullbacks to define a new fiber bundle is to take products of fiber bundles along one another.

Definition 3.9. For $i = 1, 2$, let $p_i\colon E_i \longrightarrow B$ for $i = 1, 2$ be a fiber bundle with fiber F_i. The *product bundle* $E_1 \times_B E_2 \longrightarrow B$ is the diagonal map in the pullback square

$$\begin{array}{ccc} E_1 \times_B E_2 & \longrightarrow & E_1 \\ \downarrow & & \downarrow{\scriptstyle p_2} \\ E_2 & \xrightarrow{p_1} & B. \end{array}$$

This is a fiber bundle with fiber $F_1 \times F_2$. It is often denoted, somewhat ambiguously, by $p_1 \times p_2$.

When p_1 and p_2 are vector bundles, the product bundle is also a vector bundle with fiber $F_1 \oplus F_2$. This is called the *Whitney sum* of the two vector bundles, written $p_1 \oplus p_2$.

The proof that this is well-defined is left to the reader.

As pullbacks are defined using universal properties, they cannot distinguish between isomorphic copies of the same object. This directly implies the following:

Lemma 3.10. *If p_1 and p_2 are isomorphic bundles over B and $f\colon B' \longrightarrow B$ is any map, then f^*p_1 and f^*p_2 are isomorphic bundles over B'. If p is any bundle over B then $p \times p_1$ and $p \times p_2$ are isomorphic.*

Moreover, defining bundles using universal properties allows us to commute constructions past one another. As an example of such a construction, consider the following lemma:

Lemma 3.11. *Let* $p\colon E \longrightarrow B$ *and* $p'\colon E' \longrightarrow B$ *be vector bundles, and let* $f\colon B' \longrightarrow B$ *be a map. Then*

$$f^*(E \oplus E') \cong f^*E \oplus f^*E'.$$

Proof. Consider the commutative cube

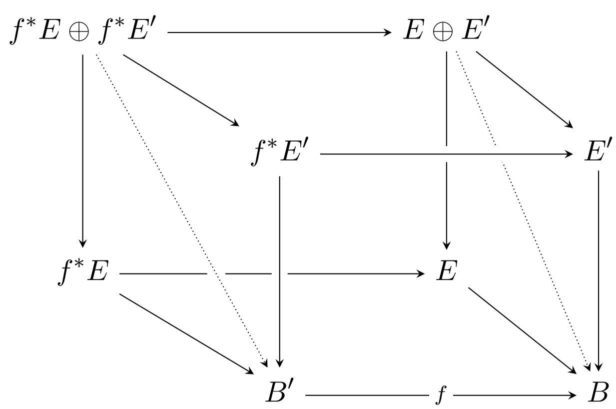

In this cube, the front, bottom, right, and left faces are pullbacks by definition. Since the left and bottom are pullbacks, the composition of them is, as well; since the right is a pullback it therefore follows (by the partial-inverse property of pullbacks) that the top is also a pullback. But this means that the composition of the top and the front is a pullback (given by the square containing the dotted maps), which implies that $f^*E \oplus f^*E'$ satisfies the universal property for $f^*(E \oplus E')$; i.e., they must be isomorphic, as desired. □

A benefit of this technique is that the proof gives more than what is stated in the lemma. Not only are these two bundles isomorphic, but they are *uniquely* isomorphic, once one takes into account their projections to B', f^*E and f^*E'.

We finish up this section with an explicit example of constructing a bundle as the pullback of another bundle.

Example 3.12. The tangent bundle to any smooth manifold is a pullback of a map from the manifold into a Grassmannian. To see this, let M be a smooth n-manifold, and let $f\colon M \hookrightarrow \mathbf{R}^N$ be a smooth embedding. Define a map $g\colon M \longrightarrow G_n(\mathbf{R}^N)$ by taking a point x to the linear subspace of $\mathbf{R}^N$ parallel to the tangent plane to M at x.

In this case,

$$g^*\gamma_{nN} = \{(x, \omega, v) \in M \times \gamma_{nN} \,|\, v \text{ is in the tangent space to } M \text{ at } f(x)\}.$$

In other words,

$$g^*\gamma_{nN} \cong TM.$$

The goal of this chapter is to prove that this is not an accident; in fact, over a nice base, *all* isomorphism classes of vector bundles can be constructed as pullbacks of the universal bundle.

3.2 The classification theorem

We can now state the classification theorem:

Theorem 3.13 (Classification theorem for vector bundles over compact spaces)**.** *Suppose that B is compact. Let $\mathrm{Vect}_n(B)$ be the set of isomorphism classes of n-dimensional vector bundles over B, and write $[B, G_n]$ for the set of homotopy classes of maps $B \longrightarrow G_n$. Then the map*

$$\rho\colon [B, G_n] \longrightarrow \mathrm{Vect}_n(B) \qquad \textit{given by} \qquad f \longmapsto f^*\gamma_n$$

is a bijection.

This is a very nifty statement: it says that the "geometrical data" of vector bundles *up to isomorphism* as the same as the "homotopical data" of *homotopy classes of maps into Grassmannians*. This is the first indication that homotopical invariants can contain information about *geometry*.

Before proving the theorem, let us explore several of its consequences; for the reader who wishes to skip directly to the proof, see page 38. As a first example, we see that the set of vector bundles depends only on the homotopy type of the base.

Corollary 3.14. *If X and Y are homotopy equivalent finite CW complexes then $\mathrm{Vect}_n(X)$ and $\mathrm{Vect}_n(Y)$ are in bijection. In particular, if X is contractible then all bundles over X are trivial.*

Proof. Let $f\colon X \longrightarrow Y$ be a homotopy equivalence, so that precomposition by f induces a bijection $[Y, G_n] \longrightarrow [X, G_n]$; by the theorem, $\mathrm{Vect}_n(X)$ and $\mathrm{Vect}_n(Y)$ are in bijection. In the special case when X is contractible, this implies that $\mathrm{Vect}_n(X) \cong [*, G_n]$. Since G_n is connected, this is a singleton, as desired. □

Approaching the theorem from the other direction, we see that constructions on G_n can induce operations on the set of vector bundles over a space. We will need an extra definition:

Definition 3.15. We say that a homotopy class $\alpha \in [B, G_n]$ *classifies* $p\colon E \longrightarrow B$ if for any $f \in \alpha$, $E \cong f^*\gamma_n$. Any such choice of f is a *classifying map* for E.

Example 3.16. Consider the map

$$\oplus\colon \mathbf{R}^\infty \times \mathbf{R}^\infty \longrightarrow \mathbf{R}^\infty$$

induced by

$$(a_1, a_2, \ldots) \oplus (b_1, b_2, \ldots) = (a_1, b_1, a_2, b_2, \ldots).$$

A point (ν, ν') in $G_m \times G_n$ is a pair of subspaces of $\mathbf{R}^\infty$. Then $\nu \times \nu'$ is a subspace of $\mathbf{R}^\infty \times \mathbf{R}^\infty$; write $\nu \oplus \nu'$ for its image under $\oplus$; this is an $n+m$-plane in $\mathbf{R}^\infty$, which is a point in G_{m+n}. This construction induces a map

$$\oplus\colon G_m \times G_n \longrightarrow G_{m+n}.$$

Let $f\colon B \longrightarrow G_m$ and $f'\colon B \longrightarrow G_n$ be classifying maps for E and E', respectively. Consider the map

$$f_\oplus\colon B \xrightarrow{\text{diag}} B \times B \xrightarrow{f \times f'} G_m \times G_n \xrightarrow{\oplus} G_{m+n}.$$

Then $f_\oplus^*\gamma_n$ is isomorphic to $E \oplus E'$. In other words, the map $\oplus$ as defined on Grassmannians classifies Whitney sum of vector bundles.

Example 3.17. Consider the map $\tau\colon G_n \longrightarrow G_{n+1}$ introduced in Example 2.30. It is represented by maps $G_n(\mathbf{R}^k) \longrightarrow G_{n+1}(\mathbf{R}^{k+1})$ adding to each n-plane the same new coordinate. Thus if we are given a map $f\colon B \longrightarrow G_n$ classifying a vector bundle p, the map $\tau \circ f$ classifies the bundle $p \oplus \epsilon^1$. In fact, the map τ classifies the bundle $\gamma_n \oplus \epsilon^1$, but the theorem does not apply as currently stated in this case, since G_n is not compact. We will later extend the theorem to all CW complexes (see Theorem 3.29).

Unfortunately, this theorem is not as powerful as it appears it first glance. In a perfect world it would be possible to classify all vector bundles on B by computing $[B, G_n]$, and this computation would be effective enough that we could use it to determine when two vector bundles are isomorphic. While this does not happen, as later chapters will show, this theorem still gives us a method for constructing interesting invariants of vector bundles.

3.3 The proof of the classification theorem

The rest of this chapter is dedicated to the proof of Theorem 3.13 and generalizing it beyond to compact spaces.

The proof of each step proceeds by first proving the necessary statements for trivializable bundles, and then by using the fact that there is a finite cover by opens over which the bundle is trivial. The main technical tool for doing this gluing is a *partition of unity*.

Definition 3.18. A *partition of unity* for X subordinate to a finite open cover $\{U_i\}_{i=1}^n$ is n functions $\varphi_i\colon X \longrightarrow I$ such that for every $x \in X$,

$$\sum_{i=1}^{n} \varphi_i(x) = 1,$$

and such that the support of each φ_i is contained inside U_i.

For a finite cover these always exist. (For a more in-depth discussion of partitions of unity, see for example [HatB, Appendix to Chapter 1].)

Step 1: ρ is well-defined

The function ρ is clearly well-defined as a function

$$\mathrm{Hom}(B, G_n) \longrightarrow \mathrm{Vect}_n(B),$$

from the set of maps $B \longrightarrow G_n$ to the isomorphism classes of n-bundles. The question of well-definedness therefore hangs on whether or not homotopic maps produce isomorphic vector bundles. This is implied by the following more general statement:

Lemma 3.19. *Let X be compact. Let $p\colon E \longrightarrow X \times I$ be a n-dimensional vector bundle. Let $f\colon X \longrightarrow I$ be any map, and write X_f for its graph inside $X \times I$. Then the isomorphism type of the restriction of E to X_f is independent of the choice of f. In particular, letting f be the constant map at 0 or 1 it follows that the restrictions of E to $X \times \{0\}$ and $X \times \{1\}$ are isomorphic.*

Proof. For any $f\colon X \longrightarrow I$, write E_f for the restriction of E to X_f. We will show that E_f is isomorphic to E_0, the case where f is the constant map at 0.

To begin, consider the case of a trivial bundle $(X \times I) \times F$. The bundle E_f is isomorphic to the space

$$\{(x, t, y) \in X \times I \times F \,|\, t = f(x)\}.$$

There is an explicit isomorphism $E_f \longrightarrow E_0$ by

$$(x, f(x), y) \longmapsto (x, 0, y).$$

We generalize this approach to a somewhat stronger statement. Let $f, f' \colon X \longrightarrow I$ be two functions, and suppose that

$$\{x \in X \mid f(x) \neq f'(x)\} \subseteq U$$

for some open $U \subseteq X$ such that $U \times I$ in the trivialization cover of p. In other words, f and f' are the same except inside a patch over which p is trivial. Then we can define an isomorphism $g \colon E_f \longrightarrow E_{f'}$ by

$$g(e) = \begin{cases} e & \text{if } p(e) \notin U \times I \\ \varphi^{-1}(x, f'(x), y) & \text{if } \varphi(e) = (x, f(x), y) \in U \times I \times F \end{cases}$$

where $\varphi \colon p^{-1}(U \times I) \longrightarrow U \times I \times F$ is the trivialization of p over $U \times I$. This is continuous because the points where f and f' are distinct are contained inside $U \times I$.

To glue these into a global isomorphism, he key observation is that the trivialization cover contains a subcover of sets of the form $\{U_\alpha \times I\}_{\alpha \in A}$, where the U_α cover X. Using the compactness of X we can then reduce to working within each of these sets separately, which is exactly the special case handled above.

To show that we can always trivialize over sets of the form $U \times I$ we first need the following observation: if E is trivializable over $U \times [a, b]$ and $U \times [b, c]$ then it is trivalizable over $[a, c]$. If the two trivialization isomorphisms agree on $U \times \{b\}$, we are done since we can just glue them together. Otherwise, given $\varphi_1 \colon E|_{U \times [a,b]} \longrightarrow U \times [a, b] \times F$ and $\varphi_2 \colon E|_{U \times [b,c]} \longrightarrow U \times [b, c] \times F \times F$, there is an induced automorphism $h \colon U \times \{b\} \times F \longrightarrow U \times \{b\} \times F$ given by $\varphi_1 \varphi_2^{-1}$. Extend this to an automorphism $h \colon U \times [b, c] \times F \longrightarrow U \times [b, c] \times F$ by ignoring the $[b, c]$-coordinate. This gives an alternate trivialization $h \circ \varphi \colon E|_{U \times [b,c]} \longrightarrow U \times [b, c] \times F$. This agrees with φ_1 on $U \times \{b\}$, and thus the earlier case applies.

Using this we show that the trivialization cover of p contains a subcover $\{U_\alpha \times I\}_{\alpha \in A}$ where $\{U_\alpha\}_{\alpha \in A}$ is a cover of X. Indeed, for any $(x, t) \in X \times I$

there exists an open subset $U_{xt} \times V_{xt}$ over which E is trivializable. Fixing x, since I is compact there exist $0 = t_0 < \cdots < t_n = 1$ such that $[t_{i-1}, t_i] \subseteq V_{xt'_i}$ for some $t'_i \in [0,1]$. Define $U_x \stackrel{\text{def}}{=} \bigcap_{i=1}^n U_{xt'_i}$, and note that E is trivializable over $U_x \times [t_{i-1}, t_i]$ for all i. Using the above observation, we conclude that U must be trivializable over $U_x \times I$, as desired.

Since X is compact, $\{U_\alpha\}_{\alpha \in A}$ contains a finite subcover $\{U_i\}_{i=1}^n$. Since this is finite, it has associated with it a subordinate partition of unity $\{\varphi_i\}_{i=1}^n$.

Define $f_i \colon X \longrightarrow I$ by

$$f_i(x) = f(x) \sum_{j=i+1}^{n} \varphi_j(x), \tag{3.20}$$

so that $f_0 = f$ and $f_n = 0$. Thus to prove the lemma it suffices to prove that $E_{f_{i-1}} \cong E_{f_i}$ for all $i \geq 1$; this is exactly the special case considered above, since f_{i-1} and f_i differ only inside the support of φ_i. □

Corollary 3.21. *If $f, g \colon X \longrightarrow Y$ are homotopic and $E \longrightarrow Y$ is a vector bundle over Y then f^*E and g^*E are isomorphic.*

Proof. Let $h \colon X \times I \longrightarrow Y$ from f to g. Let $c_t \colon X \longrightarrow I$ be the constant function at t. By Lemma 3.19, the restrictions of h^*E to the graph of c_t is independent of t. By definition, f^*E (resp. g^*E) is isomorphic to the restriction of h^*E to X_{c_0} (resp. X_{c_1}). Since these restrictions are isomorphic, $f^*E \cong g^*E$. □

This completes Step 1 of the proof.

Step 2: Rephrasing as fiberwise-injective maps

In order to analyze ρ an alternate way of looking at vector bundles will be useful. We will need to be able to both construct an arbitrary bundle as a pullback of the universal bundle, and also show that, up to homotopy, this is unique. The goal is to construct a representation of vector bundles which is easier to compute with than the one we currently have. The key points here will be that fiberwise-injective maps $E \longrightarrow \mathbf{R}^\infty$ will correspond exactly to representations of E as a pullback of the universal bundle. Moreover, homotopic maps will correspond to homotopic representations. Given this representation, checking that ρ is surjective will correspond to a representation as a fiberwise-injective map existing, and checking that ρ is injective will correspond to checking that all fiberwise-injective maps are appropriately homotopic. In this step we develop the details of this representation.

Definition 3.22. Let $p\colon E \longrightarrow B$ be a vector bundle. A *fiberwise-injective map* $E \longrightarrow \mathbf{R}^\infty$ is a map which is a linear injection when restricted to $p^{-1}(b)$ for any $b \in B$.

Two fiberwise injections $g, g'\colon E \longrightarrow \mathbf{R}^\infty$ are *homotopic through fiberwise-injective maps* if there exists a homotopy $G\colon E \times I \longrightarrow \mathbf{R}^\infty$ such that for all $t \in I$, $G(\cdot, t)$ is fiberwise-injective.

Example 3.23. There is a fiberwise-injective map $\mathrm{proj}\colon \gamma_n \longrightarrow \mathbf{R}^\infty$ given by taking the composition

$$\gamma_n \subseteq G_n \times \mathbf{R}^\infty \xrightarrow{\mathrm{proj}_2} \mathbf{R}^\infty.$$

Example 3.24. Suppose that B is compact and let $p\colon E \longrightarrow B$ be a rank-n fiber bundle. We can construct a fiberwise-injective map $E \longrightarrow \mathbf{R}^\infty$ as follows. Let $\{U_i\}_{i=1}^m$ be a finite subcover of the trivialization cover, and let $\{\varphi_i\}_{i=1}^m$ be a subordinate partition of unity. Over each U_i we can define a fiberwise-injective map

$$\widetilde{g}_i\colon p^{-1}(U_i) \xrightarrow{\tau_i} U_i \times \mathbf{R}^n \xrightarrow{pr_2} \mathbf{R}^n$$

and then extend it to a map $g_i\colon E \longrightarrow \mathbf{R}^n$ (which will *not* be fiberwise-injective) by setting

$$g_i(e) = \begin{cases} \varphi_i(p(e))\widetilde{g}_i(e) & \text{if } e \in U_i \\ 0 & \text{otherwise.} \end{cases}$$

To assemble all of these to a fiberwise-injective map $g\colon E \longrightarrow \mathbf{R}^\infty$, we simply define

$$g(e) = (g_1(e), g_2(e), \ldots, g_m(e)) \in (\mathbf{R}^n)^m \subseteq \mathbf{R}^\infty.$$

Thus we see that all bundles can be represented using fiberwise-injective maps. In fact, fiberwise injections represent more than just the bundle data: the set of fiberwise injections is in bijection with the representations of a bundle as a pullback of the universal bundle.

Proposition 3.25. *Let* proj *be the fiberwise-injective map* $\gamma_n \longrightarrow \mathbf{R}^\infty$ *defined in Example 3.23. Let* $p\colon E \longrightarrow B$ *be a vector bundle. There is a bijection*

$$\left\{ \begin{array}{c} \textit{pullback squares} \\ \begin{array}{ccc} E & \xrightarrow{f'} & \gamma_n \\ {\scriptstyle p}\downarrow & & \downarrow \\ B & \xrightarrow{f} & G_n \end{array} \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \textit{fiberwise-injective maps} \\ E \longrightarrow \mathbf{R}^\infty \end{array} \right\}.$$

sending the square on the left to $\mathrm{proj} \circ f'$.

Proof. The map $\mathrm{proj} \circ f'$ is fiberwise-injective, since the restriction of f' to any fiber is an isomorphism, and the composition of an isomorphism and an injection is an injection. On the other hand, suppose we are given a fiberwise-injective map $g\colon E \longrightarrow \mathbf{R}^\infty$. Define $f\colon B \longrightarrow G_n$ and $g'\colon E \longrightarrow \gamma_n$ by

$$f(b) = g(p^{-1}(p(e))) \qquad \text{and} \qquad g'(e) = (f(p(e)), g(e)).$$

The map g factors as $\mathrm{proj} \circ g'$. Thus we obtain a commutative square

$$\begin{array}{ccc} E & \xrightarrow{g'} & \gamma_n \\ {\scriptstyle p}\downarrow & & \downarrow{\scriptstyle p_n} \\ B & \xrightarrow{f} & G_n. \end{array}$$

To check that it's a pullback square, it suffices to construct an isomorphism between E and the fiber product (see Example 3.3)

$$B \times_{G_n} \gamma_n = \{(b, (\omega, x)) \in B \times \gamma_n \subseteq B \times (G_n \times \mathbf{R}^\infty) \mid f(b) = \omega\}$$

compatible with the projection maps; this is the map

$$e \longmapsto (p(e), (f \circ p(e), g(e))) .$$

This is a fiberwise isomorphism because g is fiberwise-injective and the fibers are finite-dimensional vector spaces, and is therefore a bundle isomorphism, as desired.

This gives functions in both directions which are mutually inverse, so it is a bijection, as desired. □

Moreover, if two fiberwise-injective maps are homotopic through fiberwise-injective maps, the bottom maps in the corresponding pullback squares are also homotopic:

Lemma 3.26. *Let* $G\colon E \times I \longrightarrow \mathbf{R}^\infty$ *be a homotopy through fiberwise-injective maps. Then the maps* $B \longrightarrow G_n$ *corresponding to* $G(\cdot, 0)$ *and* $G(\cdot, 1)$ *are homotopic.*

Proof. Consider the map $F\colon B \times I \longrightarrow G_n$ given by $F(b, t) = G(p^{-1}(b) \times \{t\})$. By definition, $G(p^{-1}(b) \times \{t\})$ is an n-dimensional subspace of $\mathbf{R}^\infty$, and thus gives a point in G_n. This is continuous because G is, and thus gives a homotopy as desired. □

Thus to show that there is exactly one homotopy class of maps $B \longrightarrow G_n$ corresponding to a vector bundle, it suffices to show that all fiberwise-injective maps are homotopic through fiberwise-injective maps.

Lemma 3.27. *Let $p: E \longrightarrow B$ be a vector bundle. Any two fiberwise-injective maps $E \longrightarrow \mathbf{R}^\infty$ are homotopic through fiberwise-injective maps.*

Proof. Let $g_0, g_1: E \longrightarrow B$ be the two fiberwise-injective maps. Whenever $g_0(e) \neq 0$ it must also be the case that $g_1(e) \neq 0$, since g_i can only map the 0 in each fiber to 0 (since the restriction to each fiber is a *linear* injection).

It is tempting to define G by setting $G(e,t) = g_0(e)t + g_1(e)(1-t)$. However, in the case when $g_0(e), g_1(e) \neq 0$ but $g_1(e) = \lambda g_0(e)$ for some negative scalar λ, this will have a problem: when $t = -\lambda/(1-\lambda)$ this will be 0, and $G(-,t)$ will not be injective on the fiber containing e. Luckily, this is the *only* thing that can go wrong, and thus this formula shows that g_0 and g_1 are homotopic through fiberwise injections if it is never the case that $g_0(e) = \lambda g_1(e)$ for any e with $g_0(e) \neq 0$.

Moreover, the relation "homotopic through fiberwise-injective maps" is an equivalence relation, so it suffices to construct a chain of maps which are each homotopic to each other through fiberwise-injective maps.

Consider the injection $L_0: \mathbf{R}^\infty \longrightarrow \mathbf{R}^\infty$ defined by

$$(a_1, a_2, a_3, \ldots) \longmapsto (a_1, 0, a_2, 0, a_3, \ldots).$$

Both g_0 and $L_0 \circ g_0$ are fiberwise-injective, and there does not exist an e such that $g_0(e) = \lambda L_0(g_0(e))$ for a negative λ. Thus the above formula shows that g_0 and $L_0 \circ g_0$ are homotopic through fiberwise-injective maps. Analogously define $L_1: \mathbf{R}^\infty \longrightarrow \mathbf{R}^\infty$ to send $(a_1, a_2, \ldots)$ to $(0, a_1, 0, a_2, 0, \ldots)$, so that g_1 and $L_1 \circ g_1$ are homotopic through fiberwise-injective maps. The maps $L_0 \circ g_0$ and $L_1 \circ g_1$ are homotopic through fiberwise-injective maps, since they never share any nonzero coordinates Thus g_0 and g_1 are homotopic through fiberwise-injective maps, as desired. □

Step 3: Checking bijectivity

With the above results, we can now prove that ρ is a bijection.

Proof that ρ is a bijection. First, consider surjectivity. Given a rank-n vector bundle $p: E \longrightarrow B$, it can be represented as a pullback of the universal bundle if and only if (by Proposition 3.25) the set of fiberwise injections $E \longrightarrow \mathbf{R}^\infty$ is nonempty. By Example 3.24, these always exist, so ρ is surjective.

Now consider injectivity. Suppose that $\rho([f]) = \rho([f'])$, so that there exists an isomorphism $\alpha\colon f^*(\gamma_n) \longrightarrow (f')^*(\gamma_n)$. In particular, this means that there exists a diagram

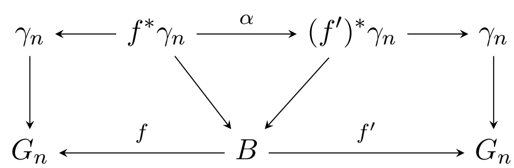

In this diagram, both the left square and the right square are pullback squares, and thus correspond to fiberwise injections $f^*\gamma_n \longrightarrow \mathbf{R}^\infty$ and $(f')^*\gamma_n \longrightarrow \mathbf{R}^\infty$. Since all fiberwise-injective maps are homotopic through fiberwise-injective maps, (by Lemma 3.27) applying Lemma 3.26 shows that f and f' are homotopic. Thus $[f] = [f']$, as desired. □

3.4 Beyond compactness

Theorem 3.13 is beautiful, but somewhat unsatisfying. Firstly, although compact spaces arise often, we often want to work with more general spaces. Moreover, the space G_n is itself not compact, and so it appears that we are classifying all vector bundles on compact spaces using a structure on a noncompact space (which is aesthetically unsatisfying). It turns out that the above proof actually works in a much wider class of spaces, which will in particular include all CW-complexes (and thus also G_n).

In order to do this, let us inspect the proof above to see where compactness was used:

(a) In the proof of Lemma 3.19 it is used to ensure that the partition of unity $\{\varphi_i\}_{i=1}^n$ exists, in order to define the maps f_i (see (3.20)).

(b) In the same proof it is also used because the final isomorphism is a composition of m isomorphisms $E_{f_{i-1}} \longrightarrow E_{f_i}$.

(c) In Example 3.24 it is used to ensure that the partition of unity $\{\varphi_i\}_{i=1}^m$ exists, in order to extend the maps $\widetilde{g}_i$ continuously to all of E.

(d) The finiteness of m is also used in order to have $(\mathbf{R}^n)^m \subseteq \mathbf{R}^\infty$. This portion of the proof will still work if m is *countably* infinite, although not if it is uncountably infinite, since almost all of the coordinates in the function we construct will be 0.

Thus if we wish our proof to work in a more general family of spaces, we must either explain why moving away from finite covers does not pose a problem, or else rework the proof to explain why the potential problem does not arise.

The key idea is to focus on the existence of the partition of unity, rather than on compactness. Before we were considering partitions of unity subordinate to finite covers. The following definition widens the definition somewhat to allow for countable partitions of unity:

Definition 3.28. Given a cover $\mathcal{U}$ of X, a *countable subordinate partition of unity* is a countable family of functions $\{\varphi_i\colon X \longrightarrow \mathbf{R}_{\geq 0}\}_{i=1}^{\infty}$ such that:

- for all i, the support of φ_i lies inside some element of $\mathcal{U}$,
- for all $x \in X$, there is a neighborhood of x such that only finitely many of the φ_i are nonzero in that neighborhood, and
- for all $x \in X$,
$$\sum_{i=1}^{\infty} \varphi_i(x) = 1,$$
which makes sense since all but finitely many of these are nonzero.

If X is such that every open cover admits a countable subordinate partition of unity, we say that X is *well-partitionable.*[c]

The fact that the partition of unity is countable ensures that the proof in points (c) and (d) above work. The fact that the partition is indexed over the naturals means that in point (a) the functions f_i in (3.20) is well-defined. Moreover, in point (b) the fact that near every point of B only finitely many of the functions φ_i are nonzero implies that only finitely many of the functions g_i are nonidentity. Thus near every point the function g is well-defined, and thus g is well-defined everywhere. This implies that the proof of Theorem 3.13 directly generalizes to prove the following:

Theorem 3.29 (Classification theorem for vector bundles)**.** *Suppose that B is well-partitionable. Let $\mathrm{Vect}_n(B)$ be the set of isomorphism classes of n-dimensional vector bundles over B. Then the map*

$$\rho\colon [B, G_n] \longrightarrow \mathrm{Vect}_n(B) \qquad \textit{given by} \qquad f \longmapsto f^*\gamma_n$$

is a bijection.

[c]The well-partitionable spaces are exactly the paracompact spaces. We focus on the viewpoint of partitions of unity in order to emphasize the relevant properties to our approach.

Some examples of well-partitionable spaces:

- All compact spaces are well-partitionable.
- All metric spaces are well-partitionable.
- Given a countable sequence

  $$X_0 \hookrightarrow X_1 \hookrightarrow X_2 \hookrightarrow \cdots$$

  of closed embeddings between well-partitionable spaces, the union $\bigcup_{n\geq 0} X_n$ is well-partitionable. In particular, as the spaces $G_n(\mathbf{R}^N)$ are compact for all finite N, the space G_n is well-partitionable.
- All CW complexes are well-partitionable.

It is not the case that all spaces are well-partitionable, or satisfy the condition of Theorem 3.29. For example, the long line (also known as the Alexandroff line) is not well-partitionable.

Further reading

For more on classification of vector bundles, see [HatB, Section 1.2] or [MS74, Section 5]. For a more in-depth discussion of paracompact spaces, see [Eng89, Chapter 5].

Exercises and extensions

3.1 Look up the definition of paracompactness and prove rigorously that all CW complexes are paracompact. Use this to show that the long line cannot be given the structure of a CW complex.

3.2 Let $p\colon E \longrightarrow D^n$ be any vector bundle over the n-disk. Use the proof of Theorem 3.29 to construct an isomorphism between E and the trivial bundle.

3.3 Verify that an analogous classification using complex Grassmannians works for complex vector bundles.

There are further results along these lines for other structured bundles. For example, for any topological group G, the space BG classifies the principal G-bundles (the bundles whose fibers are G, together with continuous G-action)

3.4 Verify that the bundle in Example 2.3 is well-defined..

3.5 Let $p: E \longrightarrow B$ be a vector bundle of dimension n. Define a fiber bundle $\widehat{p}: \widehat{E} \longrightarrow B$ which has fiber S^n and for which there exists a morphism of fiber bundles $p \longrightarrow \widehat{p}$ which is a dense inclusion on each fiber. This is the "fiberwise one-point compactification bundle." Explain why it is different from both the Thom space and the sphere bundle inside p.

3.6 Verify that the constructions in Definitions 3.6 and 3.9, when applied to vector bundles, produce vector bundles.

3.7 Describe the map $\otimes: G_n \times G_m \longrightarrow G_{nm}$ that classifies the *tensor product* of vector bundles.

3.8 Let $V^1(B)$ be the subset of isomorphism classes of vector bundles over B. Prove that $V^1(B)$ is a group under $\otimes$. (Hint: for inverses, consider the bundle made by composing with the inverse inside $GL_1(\mathbf{R})$.)

3.9 Let **Vect** be the category of finite-dimensional vector spaces. Consider a functor $\star: \mathbf{Vect} \times \mathbf{Vect} \longrightarrow \mathbf{Vect}$. What conditions are required on $\star$ to make it possible to define a vector bundle $E_1 \star E_2$? More generally, what if $\star$ has as its domain a subcategory of $\mathbf{Vect}^k$? (Examples you may want to work: dual bundle, quotient bundle, orthogonal complements of subbundles.)

3.10 Prove that

$$T\mathbf{R}P^n \cong \mathrm{Hom}(\gamma_{1n}, \gamma_{1n}^{\perp}).$$

(Hint: consider the map $S^n \longrightarrow \mathbf{R}P^n$; it can be used to construct a map $TS^n \longrightarrow T\mathbf{R}P^n$.)

3.11 Let $p: E \longrightarrow B$ be an n-dimensional bundle over B. Generalize Proposition 2.22 to state that if there exist m everywhere-linearly-independent sections, then

$$E \cong E' \oplus (B \times \mathbf{R}^m),$$

where E' is an $n-m$-dimensional bundle. (The difficulty here lies in defining the fiber of E': it is either necessary to describe a family of cases in which a notion of "orthogonal complement" is defined, or to give a rigorous description of a quotient bundle.) With this rephrasing, the Hairy Ball Question can be restated as follows:

Question (Hairy Ball Question)**.** *When is a vector bundle isomorphic to a trivial bundle added to a lower-dimensional bundle?*

We will give a partial answer to this question for the case of tangent bundles to manifolds in the Poincaré–Hopf Theorem (Theorem 5.10) and Section 6.3.

3.12 Let E, E', E'' be bundles over a common base B. Prove that, as B-bundles,

$$\mathrm{Hom}(A, B) \oplus \mathrm{Hom}(A, C) \cong \mathrm{Hom}(A, B \oplus C).$$

Chapter 4

Some crucial players

The first part of this chapter contains a "review" of cohomology theory, as well as an introduction to several other important concepts which will be necessary later. This is a conceptual reintroduction from the perspective which will be most useful for the rest of this book; it is intended to be neither a complete introduction nor a tutorial on computations. The definition of cohomology is introduced via Eilenberg–Steenrod Axioms and representing spectra. For a more classical introduction to cohomology, see for example [Hat02, Chapter 3] or [May99, Chapter17].

Section 4.1 introduces pointed spaces, Ω-spectra and their relationship to cohomology theories. Section 4.2 discusses the special case of Eilenberg–MacLane spectra and singular cohomology. Section 4.3 discusses the cohomology of unpointed spaces and their relationship to pointed spaces. Lastly, Section 4.4 introduces Thom spaces, Thom classes, and the Thom Isomoprhism Theorem.

In order to help students get up to speed with the types of topological tools and approaches we will be using frequently in this book, several of the proofs in this section have technical steps omitted and left as an exercises for the reader. These steps are also stated explicitly in the "Exercises" section of the chapter.

4.1 Spaces, Ω-spectra, and reduced cohomology theories

We begin with a review of the definitions of pointed spaces and reduced cohomology theories.[a]

Definition 4.1. A *pointed space* is a space X together with a distinguished basepoint $* \in X$. The category of pointed spaces, denoted $\mathbf{Top}_*$, has as objects pointed spaces and as morphisms maps which preserve basepoints.

The category of pointed spaces has a closed symmetric monoidal structure given by the *smash product*.[b] This is defined as follows:

$$X \wedge Y = X \times Y/(X \times *) \cup (* \times Y),$$

pointed via the image of $(X \times *) \cup (* \times Y)$. The right adjoint to $X \wedge \cdot$ is the pointed mapping space functor $\mathrm{Map}_*(X, \cdot)$, whose points are maps $X \longrightarrow Y$ which preserve the basepoint.

In the category of pointed spaces, the *suspension* functor $\Sigma\colon \mathbf{Top}_* \longrightarrow \mathbf{Top}_*$ is given by $S^1 \wedge \cdot$, the smash product with the (pointed) circle. Its right adjoint is the *loop space* functor $\Omega\colon \mathbf{Top}_* \longrightarrow \mathbf{Top}_*$. Explicitly,

$$\Omega X \overset{\text{def}}{=} \{f \in \mathrm{Hom}(S^1, X) \mid f(*) = *\}.$$

A *pointed homotopy* between maps $f, g\colon X \longrightarrow Y$ of pointed spaces is a map $H\colon X \wedge I_+ \longrightarrow Y$ such that $H|_{X\times\{0\}} = f$ and $H|_{X\times\{1\}} = g$. We write $[X, Y]$ for the set of pointed homotopy classes of maps $X \longrightarrow Y$.

Example 4.2. For all nonnegative m, n, $S^{n+m} \cong S^n \wedge S^m$. Thus in particular, $\Sigma\Sigma X = S^1 \wedge S^1 \wedge X \cong S^2 \wedge X$, and this extends to general n.

The functor Σ is left adjoint to the functor Ω. This adjunction is unusual, as it descends to pointed homotopy classes of maps. More rigorously, there is an isomorphism, natural in both X and Y,

$$[\Sigma X, Y] \cong [X, \Omega Y].$$

Definition 4.3. An Ω-*spectrum* X is a sequence of pointed spaces $X_0, X_1, \ldots$ together with weak equivalences $X_n \longrightarrow \Omega X_{n+1}$ for all $n \geq 0$.

For an Ω-spectrum X we define, for all $n < 0$,

$$X_n = \Omega^{-n} X_0.$$

[a] We omit all discussion of the compact-open topology on mapping spaces. For a discussion of this topology, see for example [Hat02, Appendix, "Compact-open topology"].

[b] This is not true in general topological spaces.

Ω-spectra are generally extremely difficult to construct. Even if X is finite-dimensional, or a finite CW-complex, ΩX is infinite dimensional. In order to construct an Ω-spectrum it is necessary to construct a space X_0 together with an infinite sequence of "deloopings" of X_0. Despite this complexity, Ω-spectra are fundamental to homotopy theory. In the course of this book we will meet several of them. The first is discussed in detail in Section 4.2.

Just like Grassmannians, Ω-spectra represent important topological invariants.

Definition 4.4. A *reduced cohomology theory* a sequence of functors

$$h^i\colon \mathbf{Top}_*^{\mathrm{op}} \longrightarrow \mathbf{AbGp} \qquad \text{for all } i \in \mathbb{Z},$$

together with natural *suspension isomorphisms*

$$\sigma_i\colon h^i(X) \xrightarrow{\cong} h^{i+1}(\Sigma X)$$

such that the following axioms hold:

homotopy invariance If $f_1, f_2\colon X \longrightarrow Y$ are two maps which are homotopic (relative to the basepoint) then the induced homomorphisms

$$h^i(f_1), h^i(f_2)\colon h^i(X) \longrightarrow h^i(Y)$$

are equal.

exactness Let $\alpha\colon A \hookrightarrow X$ be an inclusion of pointed spaces. Write $CA = A \wedge I$, where I is pointed at 1, and define

$$C_\alpha \overset{\text{def}}{=} X \sqcup CA/(\alpha(a) \sim a \wedge 0)$$

to be the mapping cone of α with $\beta\colon X \longrightarrow C_\alpha$ the natural inclusion. Then the induced sequence

$$h^i(C_\alpha) \xrightarrow{h^i\beta} h^i(X) \xrightarrow{h^i\alpha} h^i(A)$$

is exact.

additivity For any set of pointed spaces $\{X_i\}_{j\in J}$, the universal comparison

$$h^i\Big(\bigvee_{j\in J} X_j\Big) \longrightarrow \prod_{j\in J} h^i(X_j)$$

is an isomorphism.

When discussing the data of h^i for all i we sometimes write $h^*\colon \mathbf{Top}_*^{\mathrm{op}} \longrightarrow \mathbf{AbGp}^{\mathbb{Z}}$.

Generalized cohomology theories are extremely useful. The axioms, especially exactness, allow us to compute them for finite CW complexes "one cell at a time," which creates a powerful computational tool. Singular cohomology theory is an example of a generalized cohomology theory, although there are many others.

Theorem 4.5. *Let X be an Ω-spectrum, and define a sequence of functors $h^n\colon \mathbf{Top}_*^{\mathrm{op}} \longrightarrow \mathbf{AbGp}$ by*

$$h^n(Y) \overset{\mathrm{def}}{=} [Y, X_n].$$

This is a reduced cohomology theory.

The reduced cohomology theory produced using the above formula is called the *generalized cohomology theory represented by X*.

We provide a sketch of the proof of Theorem 4.5; some details are left to the reader, as they are good exercises in homotopy theory. The necessary steps are written out in the exercise section.

Proof Sketch. By definition, each h^n is a functor $\mathbf{Top}_* \longrightarrow \mathbf{Set}$. In order to check that this actually lands in $\mathbf{AbGp}$ it is necessary to show that $[Y, X_n] = [Y, \Omega^2 X_{n+2}]$ has a natural abelian group structure, and that the induced morphisms are actually homomorphisms. This is left as an exercise to the reader.

The suspension isomorphism is induced by the adjunction

$$h^i(Y) = [Y, X_i] \xrightarrow{\sim} [Y, \Omega X_{i+1}] \cong [\Sigma Y, X_{i+1}] = h^{i+1}(\Sigma Y).$$

This is natural by definition, since the adjunction isomorphism is natural and postcomposition with the weak equivalence $X_n \xrightarrow{\sim} \Omega X_{n+1}$ is as well.

It remains to check the axioms. Homotopy invariance is direct from the definition. Additivity follows from the definition and from the universal property of wedges (which is the coproduct in the category of pointed spaces). For exactness, consider any inclusion $\alpha\colon A \hookrightarrow Y$ and let C_α be the mapping cone, with inclusion $\beta\colon Y \hookrightarrow C_\alpha$. We must show that the sequence

$$[C_\alpha, X_n] \xrightarrow{\circ\beta} [Y, X_n] \xrightarrow{\circ\alpha} [A, X_n]$$

is exact. To prove this it suffices to check that those maps $f\colon Y \longrightarrow X_n$ which are homotopic to a constant map when restricted to A are exactly those maps which can be extended to C_α. This is left as an exercise for the reader. $\square$

It turns out that the converse of this theorem also holds: all reduced cohomology theories arise from Ω-spectra. This is the *Brown representability theorem* [Hat02, Section 4.E].

Definition 4.6. Let E be an Ω-spectrum. For any pointed space X, define

$$\widetilde{E}^n(X) \stackrel{\text{def}}{=} [X, E_n].$$

An an example of the use of the axioms for a generalized cohomology theory we will prove the long exact sequence in cohomology. Let $\alpha\colon A \hookrightarrow X$ be an inclusion; this induces an inclusion $\beta\colon X \hookrightarrow C_\alpha$. Applying the mapping cone construction to this inclusion gives an inclusion $\gamma\colon C_\alpha \hookrightarrow C_\beta$, and repeating this again gives an inclusion $C_\beta \hookrightarrow C_\gamma$. Thus there is a sequence of inclusions

$$A \xhookrightarrow{\alpha} X \xhookrightarrow{\beta} C_\alpha \xhookrightarrow{\gamma} C_\beta \hookrightarrow C_\gamma,$$

Applying h^i to this sequence and using the exactness axioms produces an exact sequence

$$h^i(C_\gamma) \longrightarrow h^i(C_\beta) \longrightarrow h^i(C_\alpha) \longrightarrow h^i(X) \longrightarrow h^i(A). \tag{4.7}$$

From the fact that cones are contractible it follows that

$$C_\beta \simeq \Sigma A \qquad \text{and} \qquad C_\gamma \simeq \Sigma X;$$

by the homotopy invariance axiom this implies that $h^i(C_\beta) \cong h^i(\Sigma A)$ and $h^i(C_\gamma) \cong h^i(\Sigma X)$. Thus the sequence (4.7) is isomorphic to the sequence

$$h^{i-1}(X) \longrightarrow h^{i-1}(A) \longrightarrow h^i(C_\alpha) \longrightarrow h^i(X) \longrightarrow h^i(A),$$

and is therefore exact at each of the three middle terms. As i was arbitrary this holds for all i, producing the usual long exact sequence in cohomology.

4.2 Example: Singular cohomology

In this section we discuss the example of singular cohomology in detail. We assume that the reader has seen the definition of singular or CW cohomology using chain complexes; the reader who needs a refresher should consult the beginning of [Hat02, Chapter 3]. We assume no further familiarity with cohomology. All material in this section is discussed in far more detail in [Hat02].

In order to construct the representing spectrum of singular cohomology, we first introduce Eilenberg–MacLane spaces.

Example 4.8. For an abelian group A, the *Eilenberg–MacLane space* $K(A,n)$ is defined to be CW complex with $\pi_n K(A,n) \cong A$ and all other homotopy groups (and π_0 if $n > 0$) trivial. When $n = 0$ we set $K(A,0)$ to be A with the discrete topology. The space $K(A,n)$ exists for any abelian group and is uniquely defined by its homotopy groups, in the sense that any two CW complexes satisfying this criterion are homotopy equivalent. For a more in-depth discussion, see for example [Hat02, Section 4.2].

The *Eilenberg–MacLane spectrum for* A, usually denoted HA, is the spectrum with n-th space equal to $K(A,n)$. Since $\pi_i \Omega X \cong \pi_{i+1} X$, we see that $\Omega K(A,n) \simeq K(A,n-1)$. Thus HA is an Ω-spectrum.

Singular cohomology with coefficients in an abelian group A is represented by HA. (For a more in-depth discussion of this, see for example [Hat02, Theorem 4.57].) For a pointed space Y write

$$\widetilde{H}^n(Y;A) \overset{\text{def}}{=} [Y, K(A,n)].$$

Define

$$\widetilde{H}^*(Y;A) \overset{\text{def}}{=} \bigoplus_{n\in\mathbb{Z}} \widetilde{H}^n(Y;A).$$

When A is clear from context (or when it is equal to $\mathbb{Z}$) it is omitted from the notation. To prove that this is ordinary singular cohomology, we use the following theorem:

Theorem 4.9 (Eilenberg–Steenrod)**.** *Let h_0, h_1 be two cohomology theories. Suppose that there is a natural transformation $\alpha\colon h_0^* \Longrightarrow h_1^*$ such that α_{S^0} is an isomorphism. Then for all pointed CW complexes X, $h_0^*(X) \cong h_1^*(X)$.*

Proof. First, note that for any cohomology theory h, $h^i(D^n) = 0$ for all i and all n. By the homotopy property, $h^i(D^n) = h^i(\text{point})$ for all n. Since

$$\text{point} \cong \bigvee_{\emptyset} S^0,$$

by the additivity property

$$h^i(\text{point}) \cong \prod_{\emptyset} h^i(S^0) \cong 0.$$

Now we will show that the homomorphism $h_0^i(S^n) \longrightarrow h_1^i(S^n)$ is an isomorphism for all i and all n. When $n = 0$ this is assumed in the statement of the theorem. Now suppose that this is true up to $n-1$, and consider the long exact sequence associated to the cofiber sequence $S^{n-1} \hookrightarrow D^n \longrightarrow S^n$. (Here, the basepoint of D^n is chosen to agree with the basepoint of S^{n-1}, so that the maps are well-defined.) This cofiber sequence, together with the natural transformation α, gives rise to the following commutative diagram:

$$\begin{array}{ccccccc}
h_0^{i-1}(D^n) & \longrightarrow & h_0^{i-1}(S^{n-1}) & \longrightarrow & h_0^i(S^n) & \longrightarrow & h_0^i(D^n) \\
\downarrow & & \downarrow \star & & \downarrow * & & \downarrow \\
h_1^{i-1}(D^n) & \longrightarrow & h_1^{i-1}(S^{n-1}) & \longrightarrow & h_1^i(S^n) & \longrightarrow & h_1^i(D^n)
\end{array}$$

The two outer groups in each row are equal to 0, by the above, so that the two horizontal middle morphisms are isomorphisms, and the homomorphism marked with $\star$ is an isomorphism by the inductive hypothesis. Thus the homomorphism marked with $*$ is also an isomorphism, as desired.

Let X^n be the n-skeleton of X. We will prove, by induction on n, that for all i, α induces an isomorphism $h_0^i(X^n) \cong h_1^i(X^n)$. We have $X^0 \cong \bigvee_{j \in J} S^0$, where J is the set of all non-basepoint 0-cells in X. In particular, for all i,

$$h_0^i(X^0) \cong \prod_{j \in J} h_0^i(S^0) \overset{\alpha}{\cong} \prod_{j \in J} h_1^i(S^0) \cong h_1^i(X^0).$$

Here the first and last isomorphisms are given by the additivity axiom. Thus the claim holds for $n = 0$. Now suppose that the claim holds up to X^{n-1}, and consider X^n. We have $X^n/X^{n-1} \cong \bigvee_{j \in J} S^n$, where J is the set of all n-cells in X. Thus by an analogous argument to the one for S^0 above, α induces an isomorphism $h_0^i(X^n/X^{n-1}) \longrightarrow h_1^i(X^n/X^{n-1})$ for all i. For all i there is a commutative diagram

$$\begin{array}{ccccccccc}
h_0^{i-1}(X^{n-1}) & \longrightarrow & h_0^i(X^n/X^{n-1}) & \longrightarrow & h_0^i(X^n) & \longrightarrow & h_0^i(X^{n-1}) & \longrightarrow & h_0^{i+1}(X^n/X^{n-1}) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
h_1^{i-1}(X^{n-1}) & \longrightarrow & h_1^i(X^n/X^{n-1}) & \longrightarrow & h_1^i(X^n) & \longrightarrow & h_1^i(X^{n-1}) & \longrightarrow & h_1^{i+1}(X^n/X^{n-1})
\end{array}$$

where the vertical morphisms are all given by α. The first and fourth morphisms are isomorphisms by the inductive hypothesis, and we have just shown that the second and fifth morphisms are also isomorphisms. Thus by the Five Lemma, the middle morphism is an isomorphism, as desired.

If X is finite-dimensional then we are done, since $X = X^n$ for some n. For an outline of the infinite-dimensional case, see Exercise 4.4 $\square$

Corollary 4.10. *If h is a cohomology theory on pointed CW complexes such that $h^n(S^0) = 0$ for all $n \neq 0$ then h is isomorphic to $\widetilde{H}^*(X; h^0(S^0))$.*

Proof Sketch. By the hypothesis in the lemma, $h^*(S^n)$ is 0 if $* \neq n$, and equal to $h^0(S^0)$ otherwise. Using the theorem it suffices to check that there is a natural transformation $h^* \Longrightarrow H^*(\cdot; h^0(S^0))$. Let X^n be the n-skeleton of X. Then

$$h^*(X^n/X^{n-1}) \cong h^*\Bigg(\bigvee_{n-\text{cells}} S^n\Bigg) \cong \prod_{n \text{ cells}} h^*(S^n).$$

This cohomology is concentrated in degree n and equal to $\prod_{n \text{ cells}} h^0(S^0)$; these are exactly the groups in the CW cohomology of X with coefficients in $h^0(S^0)$. The induced boundary homomorphism

$$h^n(X^n/X^{n-1}) \longrightarrow h^n(X^n) \longrightarrow h^{n+1}(X^{n+1}/X^n)$$

is isomorphic to the boundary homomorphism in the CW cohomology. On the other hand, by the long exact sequence in cohomology the cohomology groups of the cochain complex

$$\longrightarrow h^{n-1}(X^{n-1}/X^{n-2}) \longrightarrow h^n(X^n/X^{n-1}) \longrightarrow h^{n+1}(X^{n+1}/X^n) \longrightarrow \cdots$$

are isomorphic to $h^*(X)$. This gives the desired natural transformation, which is an isomorphism on S^0, as desired. $\square$

The upshot of the corollary is that we can now freely use either definition of ordinary cohomology.

Before we end the section, there are two important theorems about singular cohomology that we will need. The first of these is the Kunneth Theorem:

Theorem 4.11 (Kunneth Theorem, [Hat02, Theorem 3.18])**.** *Let R be a commutative ring. There is a group homomorphism, natural in both X and Y,*

$$\widetilde{H}^*(X; R) \otimes_R \widetilde{H}^*(Y; R) \longrightarrow \widetilde{H}^*(X \wedge Y; R).$$

When R is a field, or when the cohomology of either X or Y is free, this is an isomorphism.

The second theorem is the cup product $\smile$ on cohomology. This can be described simply in the following manner. Let R be a ring. For a CW complex X, the diagonal map $\Delta\colon X \longrightarrow X \wedge X$, composed with the homomorphism from the Kunneth Theorem defines a product

$$\smile\colon \widetilde{H}^*(X;R) \otimes_R \widetilde{H}^*(X;R) \longrightarrow \widetilde{H}^*(X \wedge X;R) \xrightarrow{\Delta^*} \widetilde{H}^*(X;R).$$

Theorem 4.12 (The cup product, [Hat02, Chapter 3, "The Cohomology Ring"])**.** *The homomorphism $\smile$ induces a unital graded commutative ring structure on $\widetilde{H}^*(X;R)$ for all pointed CW complexes X.*

4.3 Unpointed spaces and unreduced cohomology

This section contains a miscellany of results about (singular or CW) cohomology of unpointed spaces, and how it interacts with the pointed structures we have discussed thus far.

There is a functor $\cdot_+\colon \mathbf{Top} \longrightarrow \mathbf{Top}_*$ which takes a space and adds a disjoint basepoint. This functor is monoidal with respect to the cartesian monoidal structure on **Top**, so that

$$(X \times Y)_+ \cong X_+ \wedge Y_+.$$

For an unpointed space Y' we write

$$H^n(Y';A) \stackrel{\text{def}}{=} \widetilde{H}^n((Y')_+;A);$$

this is called the *unreduced* cohomology of Y'.

Definition 4.13. Let E be an Ω-spectrum. For an unpointed space X, define

$$E^n(X) \stackrel{\text{def}}{=} \widetilde{E}^n(X_+) = [X_+, E_n].$$

When discussing the fundamental group of a space X it is common to discuss "basepoint-independence": if X is connected then $\pi_1(X)$ is independent of the choice of basepoint. Although this is true, it is not true *naturally*. Given two different basepoints x_0, x_1, although it is true that there exists an isomorphism $\pi_1(X,x_0) \longrightarrow \pi_1(X,x_1)$, this isomorphism depends on a choice of path from x_0 to x_1. Thus, if X is not simply connected, there can be multiple different isomorphisms. This means that if we wish to work functorially we need a way of choosing a basepoint "naturally." Adding a disjoint basepoint is one way to do this, and although it cannot produce any interesting homotopy groups it does produce useful structures on cohomology.

Consider the homomorphism from the Kunneth homomorphism for the spaces X_+ and Y_+. Since $X_+ \wedge Y_+ \cong (X \times Y)_+$, the homomorphism becomes

$$H^*(X\colon R) \otimes_R H^*(Y;R) \longrightarrow H^*(X \times Y;R).$$

This is the form in which the Kunneth homomorphism is commonly encountered. Moreover, consider the case of a quotient X/A of unpointed spaces. This is naturally pointed by the point $[A]$. Moreover, the diagonal map $X_+ \longrightarrow X_+ \wedge X_+$ induces a map $\Delta\colon X/A \longrightarrow X_+ \wedge X/A$: a point $x \in X \backslash A$ maps to (x,x), while a point $a \in A$ maps to (a,a), which is equal to the basepoint in $X_+ \wedge X/A$. Thus we can define a cup product

$$\smile\colon H^*(X;R) \otimes \widetilde{H}^*(X/A) \longrightarrow \widetilde{H}^*(X_+ \wedge X/A) \xrightarrow{\Delta^*} \widetilde{H}^*(X/A).$$

Sometimes algebraic structions on cohomology are simpler to describe in the unpointed case, rather than the pointed one.

Example 4.14. Let $G = \mathbb{Z}/2$. Take the CW structure on S^∞ which has 2 cells ("hemispheres") in each dimension. $\mathbb{Z}/2$ acts on this by swapping opposite hemispheres. This induces a CW structure on $\mathbf{R}P^\infty$, which is $S^\infty/(\mathbb{Z}/2)$, with one cell in each dimension. The boundary of an even-dimensional cell is twice the cell in one dimension lower; the boundary of an odd-dimensional cell is 0. Thus $H^n(\mathbf{R}P^\infty;\mathbb{Z}/2) \cong \mathbb{Z}/2$ for all $n \geq 0$.

The ring structure of $H^*(\mathbf{R}P^\infty;\mathbb{Z}/2)$ is surprisingly simple: it is just a polynomial ring in one generator of grading 1. For a proof of this, see for example [Hat02, Theorem 3.12] or the exercises.

An analogous proof can also show that, as a graded ring,

$$H^*(\mathbf{C}P^\infty;\mathbb{Z}) \cong \mathbb{Z}[x]$$

with $|x| = 2$.

The spaces $\mathbf{R}P^\infty$ and $\mathbf{C}P^\infty$ are exactly the Grassmannians of lines in real and complex space, respectively.

Definition 4.15. The *one-point compactification*[c] of a locally compact space X, denoted $\widehat{X}$, has as its underlying set $X \cup \{\infty\}$, with the topology induced by the topology of X, in addition to declaring neighborhoods

[c]This terminology is somewhat controversial, since it is not necessarily a "compactification" of X, as generally the assumption is that X should embed densely into a compactification. However, in the case of compact X, X does not embed densely into $\widehat{X}$. This functor is sometimes called the "Alexandroff extension."

of ∞ to be the complements of compact closed subsets of X. $\widehat{X}$ is always compact.

One-point compactification is functorial, and produces a pointed space (with the basepoint given by $\{\infty\}$); this is called the *point at infinity* in $\widehat{X}$.

We now have a second way considering the cohomology of a non-pointed space. We can define

$$H_c^*(X) \stackrel{\text{def}}{=} \widetilde{H}^*(\widehat{X}).$$

This is called the *compactly-supported cohomology of* X. Moreover, because the new point ∞ "seals up" the "holes" left by the noncompactness of X, it can produce interesting homotopical information, as well (although this perpective is beyond the scope of this book).

Example 4.16. If X is compact, then $\widehat{X} \cong X_+$. This is because, since X is compact, the set $\{\infty\}$ is the complement of a compact subset of X and is therefore open. In particular, $H^*(X) \cong H_c^*(X)$.

If $X = \mathbf{R}^n$ then $\widehat{X} \cong S^n$. The ordinary unpointed cohomology $H^*(\mathbf{R}^n)$ is $\mathbb{Z}$ if $* = 0$, and 0 otherwise. On the other hand, $H_c^*(\mathbf{R}^n) \cong \widetilde{H}^*(S^n)$, which is $\mathbb{Z}$ if $* = n$ and 0 otherwise.

As an example of how classes in cohomology can represent geometric information, we consider the simple case of a sphere.

Example 4.17. Consider S^n, modeled via the one-point compactification of $\mathbf{R}^n$. We claim that a generator of $\widetilde{H}^n(S^n)$ corresponds to an orientation of $\mathbf{R}^n$. A generator of $\widetilde{H}^n(S^n)$ is represented by a one-to-one map $\mathbf{R}^n \longrightarrow \mathbf{R}^n$, up to homotopy. Up to homotopy we may assume that it fixes the origin and is linear. The image of a basis in the domain gives a basis of the codomain. Since $SL_n(\mathbf{R})$ is connected, any two bases that are related by an element in $SL_n(\mathbf{R})$ will be related by a homotopy. This is exactly the data of an orientation of $\mathbf{R}^n$, as desired.

4.4 Thom spaces and the Thom isomorphism

At this point we are ready to introduce one more important character in the story of vector bundles: the Thom space.

Definition 4.18. Let $p: E \longrightarrow B$ be an n-dimensional vector bundle. The fiber bundle $\widetilde{p}: \widetilde{E} \longrightarrow B$ is defined by taking the fiberwise 1-point compactification of E, so that the fiber of $\widetilde{E}$ above $b \in B$ is a copy of S^n. This produces a new section s_∞, the *section at infinity*, given by taking each

$b \in B$ to the point at infinity in its fiber. The *Thom space of* E, denoted $\mathrm{Th}(E)$, is defined by

$$\mathrm{Th}(E) \stackrel{\mathrm{def}}{=} \widetilde{E}/s_\infty(B).$$

When B is compact, this is the one-point compactification of E.

There is an alternate, more geometric, definition of Thom spaces which is equivalent for bundles over a well-partitionable base.

Definition 4.19. Let $p: E \longrightarrow B$ be a vector bundle, and suppose that it is possible to continuously define a positive definitive metric on each fiber. More concretely, such a metric is a map $\mu: E \longrightarrow \mathbf{R}$ such that its restriction to each fiber is a positive definite quadratic form. (This is possible if B is well-partitionable; the proof is left to the exercises.) This has two important fiber bundles sitting inside it:

disk bundle: Let $D(E) = \{e \in E \mid \mu(e) \leq 1\}$; then the preimage of any point in B is a copy of D^n.

sphere bundle: Let $S(E) = \{e \in E \mid \mu(e) = 1\}$; then the preimage of any point in B is a copy of S^{n-1}.

Then

$$\mathrm{Th}(E) \cong D(E)/S(E).$$

The Thom space is an example of a place where the pointed and unpointed notions of cohomology collide in an interesting manner. The Thom space is a fundamentally *pointed* object: it is defined as either a quotient (which is pointed by the image of the subspace that is quotiented by) or a one-point compactification (which is pointed by the "point at infinity"). However, the cohomology of the Thom space is naturally related to the cohomology of B, which is *unpointed*. The following example illustrates both the conflict and the natural solution.

Example 4.20. Suppose that $E \cong B \times \mathbf{R}^n$ is a trivial bundle. Then

$$D(E) \cong B \times D^n \qquad \text{and} \qquad S(E) \cong B \times S^{n-1}.$$

The Thom space is then

$$\begin{aligned} \mathrm{Th}(E) = D(E)/S(E) &\cong B \times D^n / B \times S^{n-1} \cong B \times S^n / B \times \{*\} \\ &\cong (B_+) \wedge S^n. \end{aligned}$$

As smashing with S^n gives the n-fold suspension, when E is trivial

$$\mathrm{Th}(E) \cong \Sigma^n(B_+).$$

In particular, for each integer i,

$$\widetilde{H}^{n+i}(\mathrm{Th}(E)) \cong H^i(B). \tag{4.21}$$

As suspensions destroy cup products, it is not the case that (4.21) holds as rings. The Kunneth theorem implies that

$$H^*(\Sigma^n B) \cong \widetilde{H}^*(S^n) \otimes H^*(B);$$

since $\widetilde{H}^*(S^n)$ is a copy of $\mathbb{Z}$ concentrated in degree n, the isomorphism $H^{i-n}(B) \cong \widetilde{H}^i(\mathrm{Th}(E))$ is given by multiplying by a generator in $\widetilde{H}^n(S^n)$.

More generally, Thom spaces work well with products of vector bundles:

Lemma 4.22. *For vector bundles $E \longrightarrow B$ and $E' \longrightarrow B'$,*

$$\mathrm{Th}(E \times E') \cong \mathrm{Th}(E) \wedge \mathrm{Th}(E').$$

Proof. The key observation is that there is a homeomorphism

$$D(E \times E') \cong D(E) \times D(E')$$

which restricts on $S(E \times E')$ to a homeomorphism

$$S(E \times E') \cong (S(E) \times D(E')) \cup (D(E) \times S(E')).$$

Then

$$\begin{aligned}\mathrm{Th}(E \times E') &= D(E \times E')/S(E \times E') \\ &\cong \big(D(E) \times D(E')\big)/\left(\big(S(E) \times D(E')\big) \cup \big(D(E) \times S(E')\big)\right).\end{aligned}$$

□

Corollary 4.23.

$$\mathrm{Th}(E \oplus \epsilon^k) \cong \mathrm{Th}(E) \wedge S^k.$$

Proof. $E \oplus \epsilon^k \cong E \times \epsilon^k$, where in the right-hand side we think of ϵ^k as a bundle over the point. In particular, in this case $\mathrm{Th}(\epsilon^k) \cong S^k$. Thus

$$\mathrm{Th}(E \oplus \epsilon^k) \cong \mathrm{Th}(E) \wedge \mathrm{Th}(\epsilon^k).$$

□

For other constructions, including the Whitney sum of bundles, the Thom spaces are harder to analyze.

The zero section $s_0 \colon B \longrightarrow E$ gives rise to a map (also by abuse of notation called s_0)

$$s_0 \colon B \xrightarrow{s_0} E \hookrightarrow \widehat{E} \cong \mathrm{Th}(E).$$

There is no projection $\mathrm{Th}(E) \longrightarrow B$ for which this map is a section. However, the "fibers" above various points in B sit inside $\mathrm{Th}(E)$ nicely: for every $b \in B$ there is a map $S^n \hookrightarrow \mathrm{Th}(E)$ induced by the inclusion $(p|_{D(E)})^{-1}(b) \cong D^n \hookrightarrow D(E)$. This induces a restriction $\widetilde{H}^n(\mathrm{Th}(E)) \longrightarrow \widetilde{H}^n(S^n)$, called the *restriction of u to the fiber above b*. This also works for cohomology with coefficients in any abelian group A.

Definition 4.24. A class $c \in \widetilde{H}^n(\mathrm{Th}(E); A)$ is called a *Thom class with coefficients in A* if for every $b \in B$ the restriction of c to the fiber above b is a generator for $\widetilde{H}^n(S^n; A)$.

The usual cases of interest are when $A = \mathbb{Z}/2$ and when $A = \mathbb{Z}$. When $A = \mathbb{Z}$ this is often simply called a *Thom class*.

A bundle with a choice of Thom class with coefficients in A is called an *A-oriented vector bundle*.

More informally, using the intuition from Example 4.17, the restriction of c to the fiber $\widehat{E}_b$ over b gives a consistent choice of orientation of the fiber E_b. This is exactly the intuitive notion of an orientation of a vector bundle: an orientation on each fiber that assemble in a reasonable manner to a global orientation.

Thom classes do not always exist; for example, the Möbius bundle has no Thom class if we take $\mathbb{Z}$ coefficients. Thom classes are natural with respect to pullbacks, in the following sense:

Proposition 4.25. *Let $p \colon E \longrightarrow B$ be an oriented vector bundle, let $c \in \widetilde{H}^n(\mathrm{Th}(E))$ be the chosen Thom class, and let $f \colon B' \longrightarrow B$ be any map. There is an induced map $\mathrm{Th}(f) \colon \mathrm{Th}(f^*(E)) \longrightarrow \mathrm{Th}(E)$ which takes c to a Thom class for f^*E.*

Heuristically speaking, this states that an oriented bundle induces an orientation on a pullback bundle.

Proof. The map $(f')^* \colon \mathrm{Th}(f^*(E)) \longrightarrow \mathrm{Th}(E)$ is induced from the map $f^*(E) \longrightarrow E$, which is locally the one-point compactification. It thus remains to check that $(f')^*(c)$ is a Thom class. For any $b' \in B$, by the definition of the pullback bundle, $(p')^{-1}(b') = p^{-1}(f(b'))$, and the inclusion of the fiber factors through the map f'. Thus since the pullback of c to the cohomology of any fiber in E is a generator, this must also be true for $(f')^*(c)$. □

The most important feature of the Thom space is the ThomIsomorphism Theorem, which states that the cohomology of the Thom space is simply a "shift" of the cohomology of the base B. Although this may appear at first glance to imply that the Thom space does not possess interesting information about the bundle, this turns out to not be the case.

Theorem 4.26 (Thom Isomorphism Theorem)**.** *Let c be a Thom class with coefficients in A for the n-dimensional bundle $p\colon E \longrightarrow B$. The homomorphism*

$$\Phi\colon H^i(B;\mathbb{Z}/2) \longrightarrow \widetilde{H}^{i+n}(\mathrm{Th}(E);\mathbb{Z}/2) \qquad b \longmapsto p^*(b) \smile c$$

is an isomorphism for all i.

We prove the theorem for compact base spaces. It holds for general base spaces, as well, but a proper explanation of the proof requires an exploration of cohomology for arbitrary unions of spaces. See [MS74, Section 10] for a detailed discussion.

Before we begin the proof, let us define precisely what "$\smile c$" actually means. The diagonal map $E_+ \longrightarrow E_+ \wedge E_+$ induces a map $\mathrm{Th}(E) \longrightarrow E_+ \wedge \mathrm{Th}(E)$. Thus the cup product on $H^*(E)$ induces a cup product

$$H^*(E) \otimes \widetilde{H}^*(\mathrm{Th}(E)) \xrightarrow{\smile} \widetilde{H}^*(\mathrm{Th}(E)).$$

As $p^*\colon H^*(B) \longrightarrow H^*(E)$ is an isomorphism, precomposing with $p^* \otimes 1$ produces a homomorphism

$$H^*(B) \otimes \widetilde{H}^*(\mathrm{Th}(E)) \longrightarrow H^*(\mathrm{Th}(E)).$$

The map Φ is simply the restriction to the subgroup where the second coordinate in the tensor product is set to be c.

Proof of Thom Isomorphism Theorem. If E is trivial, then by Example 4.20 the Thom class exists: in the notation of that example, it is $1 \otimes \alpha$, where 1 is the generator of $H^0(B)$. The choice of α is dictated by the chosen orientation on the fibers. Moreover, by the Kunneth Theorem multiplication by it induces an isomorphism between the homologies.

Now suppose that $B = U \cup V$, with U and V open subsets of B such that the Thom Isomorphism Theorem holds for the bundles $E|_U$, $E|_V$, and $E|_{U\cap V}$ with chosen orientations given by restriction. Call the Thom classes associated to these bundles $c_U, c_V, c_{U\cap V}$, respectively. Since the Thom Isomorphism Theorem holds for $E|_{U\cap V}$, in particular it is the case that $\widetilde{H}^{n-1}(\mathrm{Th}(E|_{U\cap V})) = 0$. We therefore have a Mayer–Vietoris sequence, which begins:

$$\widetilde{H}^n(\mathrm{Th}(E)) \hookrightarrow \widetilde{H}^n(\mathrm{Th}(E|_U)) \oplus \widetilde{H}^n(\mathrm{Th}(E|_V)) \overset{u-v}{\longrightarrow} \widetilde{H}^n(\mathrm{Th}(E|_{U\cap V})) \longrightarrow \cdots .$$

The inclusion $\mathrm{Th}(E|_{U\cap V}) \longrightarrow \mathrm{Th}(E|_U)$ induces a homomorphism

$$\widetilde{H}^n(\mathrm{Th}(E|_U)) \longrightarrow \widetilde{H}^n(\mathrm{Th}(E|_{U\cap V});$$

since the theorem holds for $E|_{U\cap V}$ and Thom classes are preserved under pullbacks (and are usnique because of the specified orientation), the image of c_U is $c_{U\cap V}$. Thus the image of $c_U \oplus c_V$ in the middle term is 0; since the sequence is exact there is a unique element $c \in \widetilde{H}^n(\mathrm{Th}(E))$ which hits $c_U \oplus c_V$.

The Mayer–Vietoris sequence above receives a homomorphism from the Mayer–Vietoris sequence

$$\cdots \longrightarrow \widetilde{H}^i(E|_{U\cap V}) \longrightarrow \widetilde{H}^i(E) \longrightarrow \widetilde{H}^i(E|_U) \oplus \widetilde{H}^i(E|_V) \longrightarrow \widetilde{H}^i(E|_{U\cap V}) \longrightarrow \cdots$$

induced by cupping with the appropriate Thom class. By the induction hypothesis this is an isomorphism on all terms other than the terms $\widetilde{H}^i(E)$; by the five lemma these must also be isomorphisms, and the proof of this case is complete.

We can now use the above arguments to prove the Thom Isomorphism theorem for compact bases B. We proceed by induction on the number of patches n needed to cover B such that the bundle is trivial over each patch. When $n = 1$ the bundle is trivial, which is the first case we handled above. If the theorem holds for $n-1$ then when $B = U_1 \cup \cdots \cup U_n$ split it up as $U = U_1 \cup \cdots \cup U_{n-1}$ and $V = U_n$. Then the theorem holds for U and V, and also for $U \cap V = \bigcup_{i=1}^{n-1}(U_i \cap V)$. Thus the above argument applies, and it also holds for $U \cup V = B$, as desired. □

Further reading

For a detailed introduction to cohomology, see [Hat02, Chapter 4]. For a discussion of Eilenberg–MacLane spaces in particular, see [Hat02, Section 4.3].

A proof that the cup product is Poincare dual to the intersection of submanifolds can be found in [GM80]. The question of which homology classes can be represented by submanifolds is in itself interesting, and is answered in [Tho54, Section II]; it is known that in small dimensions and codimensions everything is realizable, but in general there are classes that are not realizable.

A complete proof of the Thom isomorphism theorem is discussed in [Coc62]. There is also a good discussion of the difficulties of extending the proof presented in the text for compact bases to general bases in [MS74, Section 10].

Exercises and Extensions

Exercises 1-3 fill in the details of Theorem 4.5.

4.1 Let X, X', Y, Z, Z' be pointed spaces. Prove that if $Y \simeq \Omega Z$ then the set $[X, \Omega Z]$ has the structure of a group. For any map $f\colon X' \longrightarrow X$, show that the induced function $[X, Y] \longrightarrow [X', Y]$ is a group homomorphism. If in addition $Z \simeq \Omega Z'$, show that the set $[X, Y]$ has the structure of an abelian group.

4.2 Prove that Σ and Ω are adjoint functors $\mathbf{Top}_* \longrightarrow \mathbf{Top}_*$, and that this adjunction descends to (pointed) homotopy classes of maps. More concretely, show that there exists a bijection, natural in both X and Y,

$$[\Sigma X, Y] \cong [X, \Omega Y].$$

4.3 Let $A \hookrightarrow X$ be an inclusion of pointed spaces. For any pointed space Z, show that the sequence

$$[X \cup_A CA, Z] \longrightarrow [X, Z] \longrightarrow [A, Z]$$

is exact as a sequence of pointed sets, in the sense that a map $X \longrightarrow Z$ is homotopic to the constant map when restricted to A if and only if it extends to a map $X \cup_A CA \longrightarrow X$.

4.4 This exercise completes the proof of Theorem 4.9. (Based on [Hat02, Proof of Lemma 2.34].)

(a) Suppose that it is known that if the n-skeleton of X is a point, then $h^i(X) \cong 0$ for $k \leq n$. Complete the proof of Theorem 4.9.

(b) Give the half-line $[0, \infty)$ a CW structure by putting a 0-cell at every integer and a 1-cell connecting i and $i+1$. Prove that $X \times [0, \infty)$ is homotopy equivalent to X.

(c) Let X be a CW complex whose n-skeleton is a point. Let $T = \bigcup_{i \geq 0} X^i \times [i, \infty)$, considered as a subcomplex of $X \times [0, \infty)$. Prove that X is homotopy equivalent to T.

(d) Let X be as in the previous part, and let

$$Z = \left(X^0 \times [0, \infty)\right) \cup \bigcup_{i \geq 0} X^i \times \{i\}.$$

Prove that $h^i(Z) \cong h^i(T/Z) \cong 0$ for $i \leq n$.

(e) Conclude part (a) from part (d).

4.5 The goal of this exercise is to compute the cohomology ring structure of real projective spaces. We will be using the map $r\colon \mathbf{R}P^{n-1} \times \mathbf{R}P^1 \longrightarrow \mathbf{R}P^{2n-1}$ given by

$$r([x_0\colon \cdots \colon x_{n-1}], [y_0, y_1]) = [x_0 y_0\colon \cdots \colon x_{n-1} y_0\colon x_0 y_1\colon \cdots x_{n-1} y_1].$$

(a) Construct a CW structure on S^n which is compatible with the $\mathbb{Z}/2$-action given by multiplying by -1. (Here, by "compatible" we mean that applying -1 to a cell gives a homeomorphism to another cell.) Use this CW structure to prove that $H^i(\mathbf{R}P^n; \mathbb{Z}/2)$ is $\mathbb{Z}/2$ for $0 \leq i \leq n$ and 0 otherwise.

(b) Prove that the inclusion $\mathbf{R}P^k \longrightarrow \mathbf{R}P^n$ induces a ring homomorphism which is an isomorphism on cohomology up to degree k.

(c) Let $\alpha \in H^1(\mathbf{R}P^n; \mathbb{Z}/2)$ be a generator. Let $\beta \in H^1(\mathbf{R}P^{2n-1})$ be a generator. Prove that the map on cohomology induced by

$$p\colon \mathbf{R}P^{n-1} \hookrightarrow \mathbf{R}P^{n-1} \times \mathbf{R}P^1 \xrightarrow{r} \mathbf{R}P^{2n-1}$$

sends β to α.

(d) Prove that the map on cohomology induced by

$$q\colon \mathbf{R}P^1 \hookrightarrow \mathbf{R}P^{n-1} \times \mathbf{R}P^1 \xrightarrow{r} \mathbf{R}p^{2n-1}$$

sends β to a generator of $H^1(\mathbf{R}P^1; \mathbb{Z}/2)$.

(e) Suppose that α^{n-1} is a generator in $H^{n-1}(\mathbf{R}P^{n-1})$. Use the Kunneth theorem to prove that

$$p_1^*(p^*(\beta)^{n-1}) \times p_2^*(q^*(\beta)) \in H^n(\mathbf{R}P^{n-1} \times \mathbf{R}P^1; \mathbb{Z}/2)$$

is nonzero. Here, p_1 and p_2 are the two projections from $\mathbf{R}P^{n-1} \times \mathbf{R}P^n$.

- Use induction to conclude that $H^*(\mathbf{R}P^n; \mathbb{Z}/2) \cong \mathbb{Z}/2[\alpha]/\alpha^{n+1}$ and thus that $H^*(\mathbf{R}P^\infty; \mathbb{Z}/2) \cong \mathbb{Z}/2[\alpha]$.

4.6 Let $p: E \longrightarrow B$ be an n-dimensional vector bundle over B. Show that if B is well-partitionable then there exists a Euclidean metric on E: a function $\mu: E \longrightarrow \mathbf{R}$ such that its restriction to each fiber is a positive definite quadratic form.

Chapter 5

The Euler class

In this chapter we construct our first example of a characteristic class: the Euler class. This class is closely connected to the Euler characteristic, and will bring us to an answer to the Hairy Ball Question.

5.1 The construction

In this section, all cohomology will be with A-coefficients, with the usual examples to keep in mind being $A = \mathbb{Z}/2$ or $\mathbb{Z}$. In the interest of space we omit the coefficients from the notation. As usual, $p\colon E \longrightarrow B$ is an n-dimensional A-oriented vector bundle with a Thom class $c \in \widetilde{H}^n(\mathrm{Th}(E))$.

Since $\mathrm{Th}(E) \cong D(E)/S(E)$, there is a long exact sequence in cohomology

$$\cdots \longrightarrow \widetilde{H}^i(\mathrm{Th}(E)) \xrightarrow{j^*} H^i(D(E)) \longrightarrow H^i(S(E)) \longrightarrow \widetilde{H}^{i+1}(\mathrm{Th}(E)) \xrightarrow{j^*} \cdots, \tag{5.1}$$

where we write $j\colon D(E) \longrightarrow \mathrm{Th}(E)$ for the quotient map. From the Thom isomorphism theorem and the fact that the map $s_0\colon B \longrightarrow D(E)$ is a homotopy equivalence, we obtain the following diagram

$$\begin{array}{ccccc} \widetilde{H}^i(\mathrm{Th}(E)) & \xrightarrow{j^*} & H^i(D(E)) & \longrightarrow & H^i(S(E)) \\ \uparrow{\scriptstyle \cdot \smile c} & & \downarrow{\scriptstyle s_0^*} & \nearrow{\scriptstyle (p|_{S(E)})^*} & \\ H^{i-n}(B) & \dashrightarrow & H^i(B) & & \end{array}$$

where the two vertical arrows are isomorphisms. Thus the dotted map exists. Moreover, the cup product is natural in the sense that the diagram

$$\begin{array}{ccc} H^*(D(E)) \otimes H^*(\mathrm{Th}(E)) & \xrightarrow{\smile} & H^*(\mathrm{Th}(E)) \\ \downarrow{\scriptstyle 1\otimes j^*} & & \downarrow{\scriptstyle j^*} \\ H^*(D(E)) \otimes H^*(D(E)) & \xrightarrow{\smile} & H^*(D(E)) \end{array}$$

commutes. Thus the dotted arrow in the larger diagram can be rewritten as $\smile (s_0^* j^* c)$.

Definition 5.2. Let $p\colon E \longrightarrow B$ be an n-dimensional vector bundle with Thom class $c \in \widetilde{H}^n(\mathrm{Th}(E))$. Let $s_0\colon B \longrightarrow \mathrm{Th}(E)$ be the zero-section. The *A-Euler class* $e(E)$ is the cohomology class

$$e(E) \overset{\text{def}}{=} s_0^* c \in H^n(B; A).$$

The long exact sequence (5.1) can therefore be rewritten as

$$\cdots \longrightarrow H^{i-n}(B) \xrightarrow{\smile e(E)} H^i(B) \xrightarrow{(p|_{S(E)})^*} H^i(S(E)) \longrightarrow H^{i-n+1}(B) \longrightarrow \cdots.$$

This sequence is called the *Gysin sequence*, and it exists for any sphere bundle.

Since Thom classes are preserved under pullbacks, so are Euler classes.

Lemma 5.3. *If $p\colon E \longrightarrow B$ is an n-dimensional vector bundle with Euler class $e(E) \in \widetilde{H}^n(B)$, and $f\colon B' \longrightarrow B$ is any map, then $f^*(e(E)) = e(f^*(E))$.*

The following proposition is stated somewhat loosely, as a properly rigorous description of it requires delving more deeply into transversality, which is beyond the scope of this book. Informally, a "transverse" intersection is one where there is no tangency or singularity at the intersection points. For a proper introduction to the differential viewpoint necessary for this result, see for example [BT13, Chapter 4,5].

Proposition 5.4. *Let $A = \mathbb{Z}/2$ or $\mathbb{Z}$, and let B be a smooth n-manifold such that TB is A-orientable. Then $H^n(B) \cong A$. For an A-oriented n-dimensional vector bundle E over B, the $e(E) = n[1]$, where n is the number of points*[a] *in the intersection of a generic section $s\colon B \longrightarrow E$ and the 0-section $s_0\colon B \longrightarrow E$. Here, by "generic section" we mean a section which intersects the 0-section transversally.*

[a] Counted correctly with signs depending on the local orientations near the intersection points

Corollary 5.5. *If E has an everywhere-nonzero section then $e(E)$ is zero for any coefficients.*

Proof. Suppose that E has an everywhere-nonzero section. Then the intersection of this with the 0-section is empty, so $e(A) = 0$. □

In particular, if there exists any bundle with a nonzero Euler class then there exist bundles with no everywhere-nonzero section. (I.e. these are bundles it is impossible to "comb.")

Corollary 5.6. *For all n there exists an n-bundle with a nonzero $\mathbb{Z}/2$-Euler class. Consequently, for the universal bundle γ_n over G_n, the $\mathbb{Z}/2$-Euler class e is nonzero.*

Proof. Consider the Möbuis bundle $\mu \longrightarrow S^1$. This has a section which intersects 0 at exactly one point (for example, by considering S^1 to be $[0,1]/\{0,1\}$ and defining $s(t) = (t, \sin(2\pi t - \frac{\pi}{2}) + 1)$;[b] this will be 0 at 0 and 1 but is nonzero everywhere else). Now consider the bundle $\mu^n \longrightarrow (S^1)^n$. This has a section s^n, which is also zero at exactly one point. Thus the $\mathbb{Z}/2$-Euler class of this bundle (which is dual to a single point) is nonzero, as desired.

Since the $\mathbb{Z}/2$-Euler class is preserved under pullbacks, if $e(\gamma_n) = 0$ then for any n-bundle E, $e(E) = 0$. Since we have constructed an example of a bundle where this is not the case, $e(\gamma_n) \neq 0$.

□

Remark 5.7. This result should not be surprising, since a single nonvanishing section allows us to "slice off" a line bundle. Thus if a universal n-plane bundle had a nonvanishing section this would mean that *all* n-bundles had a line bundle that could be sliced off.

As an illustration of how Euler classes can be computed, we show that the Euler class commutes with products.

Proposition 5.8. *Let $p\colon E \longrightarrow B$ and $p'\colon E' \longrightarrow B'$ be vector bundles. Then the $\mathbb{Z}/2$-Euler class for the bundle $p \times p'\colon E \times E' \longrightarrow B \times B'$ is*

$$e(E \times E') = e(E) \otimes e(E').$$

The formula in the proposition above makes sense: the Kunneth theorem implies that $H^*(B \times B'; \mathbb{Z}/2) \cong H^*(B; \mathbb{Z}/2) \otimes H^*(B'; \mathbb{Z}/2)$. In particular, the Euler class of $E \times E'$ should live in $H^{\dim E + \dim E'}(B \times B')$, which contains the element $e(E) \otimes e(E') \in H^{\dim E}(B) \otimes H^{\dim E'}(B')$.

[b] The formula $t(1-t)$ would work equally well, but this construction is used to make clear that it can be chosen to be smooth.

Proof. There are isomorphisms

$$H^*(B) \xrightarrow{\smile c} H^*(\mathrm{Th}(E)) \qquad \text{and} \qquad H^*(B') \xrightarrow{\smile c'} H^*(\mathrm{Th}(E')).$$

These induce an isomorphism

$$\begin{aligned} H^*(B \times B') \cong H^*(B) \otimes H^*(B') &\xrightarrow{\smile (c \otimes c')} \\ &\longrightarrow H^*(\mathrm{Th}(E)) \otimes H^*(\mathrm{Th}(E')) \cong H^*(\mathrm{Th}(E \times E')), \end{aligned}$$

where the first and last step use the Kunneth theorem. Thus $c \smile c'$ is exactly a Thom class for $E \times E'$. By Lemma 4.22, $\mathrm{Th}(E \times E') \cong \mathrm{Th}(E) \wedge \mathrm{Th}(E')$. Since the Euler class is the pullback of the Thom class along the 0-section, we conclude that the Euler class of $E \times E'$ is the tensor product of the Euler classes of the bundles, as desired. □

Using the fact that the cup product is induced using an external product and then pulling back along the diagonal map, we can conclude the following:

Corollary 5.9. *Let $p\colon E \longrightarrow B$ and $p'\colon E' \longrightarrow B$ be two vector bundles. Then*

$$e(E \oplus E) = e(E)e(E') \in H^*(B).$$

We complete this section with the statement of the the Poincaré–Hopf Theorem:

Theorem 5.10 (Poincaré–Hopf Theorem)**.** *It is possible to comb a closed Riemannian $\mathbb{Z}$-oriented manifold M if and only if the Euler characteristic of M is zero.*

The proof of this theorem depends heavily on differential geometry, and is therefore beyond the scope of this book. The interested reader should see [MS74, Section 12], which contains an outline of the proof and references to the relevant sources.

We give a (very short) sketch of the theoretical underpinnings of the "only if" direction.

Basic idea. By Corollary 5.6, it suffices to check that if the Euler characteristic of a smooth closed $\mathbb{Z}$-oriented manifold M is nonzero then the $\mathbb{Z}$-Euler class is nonzero. In fact, a stronger statement is true: the Euler class will be $\chi(M)[1]$, where $\chi(M)$ is the Euler characteristic of M. To prove this a shift in perspective is necessary. In the cohomology of manifolds, some classes in H^k can be represented by codimension-k submanifolds. In $\mathrm{Th}(TM)$ the

Thom class can be represented[c] by the 0-section; via a homotopy, we can replace it by a nearby section which intersects the 0-section transversely. Using Corollary 5.6, the Euler class is the number of times this nearby section intersects the 0-section. (This is the "self-intersection number of the 0-section.") That this number is equal to the Euler characteristic requires a significant amount of differential geometry, which is beyond the scope of this book. See [MS74, Chapter 11] or [BT13]. □

5.2 The Cohomology of Grassmannians

We are now almost ready to compute the cohomology of Grassmannians with $\mathbb{Z}/2$-coefficients.

To begin, a lemma which is vital to relating the structure of G_n and G_{n-1}:

Lemma 5.11. *Let $S(\gamma_n)$ be the unit sphere bundle in γ_n (under the standard metric on $\mathbf{R}^\infty$). Then*

$$S(\gamma_n) = \{(v, \omega) \in \mathbf{R}^\infty \times G_n \mid v \in \omega,\ \|v\| = 1\},$$

and the map $p' \colon S(\gamma_n) \longrightarrow G_{n-1}$ given by $(v, \omega) \longmapsto v^\perp \cap \omega$ is a homotopy equivalence.

Proof. The key observation is that p' gives $S(\gamma_n)$ the structure of a fiber bundle with fiber S^∞. Since S^∞ is contractible, by the long exact sequence of homotopy groups (Theorem 2.8), p' induces isomorphisms on all homotopy groups, and thus also on cohomology rings. The fact that it is a homotopy equivalence follows from the fact that any fiber bundle with connected base and CW structures on the base and the fiber has the type of a CW complex.[d] □

Definition 5.12. Let $\eta = \tau^* \colon H^*(G_n) \longrightarrow H^*(G_{n-1})$ be the ring homomorphism on homology induced by the map $\tau \colon G_{n-1} \longrightarrow G_n$ (see Example 3.17).

We want to compute the cohomology of Grassmannians by induction. The natural map which provides the inductive step is η, but the current definition is not easy enough to work with to make this induction straightforward. The lemma above allows us to construct another model for η which will make this computation simpler.

[c] Th(TM) is not a manifold, but it turns out to be close enough for this to work

[d] See for example [FP90, Theorem 5.4.2].

Lemma 5.13. *Pick a map $G_{n-1} \longrightarrow S(\gamma_n)$ which is a homotopy inverse to p'. Then the composition $G_{n-1} \longrightarrow S(\gamma_n) \longrightarrow G_n$ also classifies the bundle $\gamma_n \oplus \epsilon^1$. It is thus homotopic to τ and induces the homomorphism η on cohomology.*

Proof. The second part of the lemma is a direct application of the Classification Theorem for vector bundles, so we focus on the first part.

Define

$$E = \left\{ (\omega, v, w) \in G_{n+1} \times \mathbf{R}^\infty \times \mathbf{R}^\infty \,\middle|\, v, w \in \omega \right\},$$

which, with the natural projection $E \longrightarrow S(\gamma_{n+1})$ forgetting the w-coordinate, is an $n+1$-bundle on $S(\gamma_{n+1})$. We will show that E is isomorphic to both the pullback of γ_{n+1} along $p\colon S(\gamma_{n+1}) \longrightarrow G_{n+1}$ and the pullback of $\gamma_n \oplus \epsilon^1$ along $p'\colon S(\gamma_{n+1}) \longrightarrow G_n$. This will imply, by Theorem 3.29 that the diagram

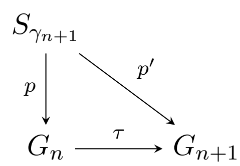

commutes up to homotopy, proving the desired statement.

E is the pullback of γ_{n+1} along $S(\gamma_{n+1}) \longrightarrow G_{n+1}$ by definition. We can also write $E \cong E^\perp \oplus E^\parallel$, where

$$E^\perp \stackrel{\text{def}}{=} \left\{ (\omega, v, w) \in E \,\middle|\, w \perp v \right\}, \text{ and}$$
$$E^\parallel \stackrel{\text{def}}{=} \left\{ (\omega, v, w) \in E \,\middle|\, w \| v \right\}.$$

The pullback of γ_n along p' is isomorphic to $E^\perp$ by definition, and the pullback of ϵ^1 is isomorphic to $E^\parallel$. Thus E is the pullback of $\gamma_n \oplus \epsilon^1$ along p', as desired. □

We are now ready to compute the cohomology of Grassmannians.

Theorem 5.14.

$$H^*(G_n; \mathbb{Z}/2) \cong \mathbb{Z}/2[w_1, \ldots, w_n] \qquad \text{with } |w_i| = i.$$

Proof. We prove this by induction on n using the Gysin sequence for the universal bundle $\gamma_n \longrightarrow G_n$. We can start the induction at $n = 0$, where

$G_0 = *$ and the result holds trivially.[e] Thus we can assume that $H^*(G_{n-1}) \cong \mathbb{Z}/2[w_1, \ldots, w_{n-1}]$.

The Gysin sequence for γ_n is

$$\cdots \longrightarrow H^i(G_n) \xrightarrow{\smile e} H^{i+n}(G_n) \xrightarrow{\eta} H^{i+n}(G_{n-1}) \longrightarrow H^{i+1}(G_n) \longrightarrow \cdots.$$

For $-n \leq i < -1$ this says that $\widetilde{H}^{i+n}(G_n) \cong \widetilde{H}^{i+n}(G_{n-1})$; thus for each generator $w_j \in H^*(G_{n-1})$ (with $j < n-1$) there exists a unique generator $w'_j \in \widetilde{H}^*(G_n)$ such that $\eta(w'_j) = w_j$.

When $i = 0$, $\smile e\colon H^0(G_n) \longrightarrow H^n(G_n)$ is injective by Lemma 5.6, and thus $H^{n-1}(G_n) \cong H^{n-1}(G_{n-1})$; thus w'_{n-1} exists in $H^*(G_n)$ as well. In addition, since $H^*(G_{n-1})$ is generated by the w_j and η is a ring homomorphism, it is surjective in each degree. Thus the Gysin sequence above splits for all i into short exact sequences

$$0 \longrightarrow H^i(G_n) \xrightarrow{\smile e} H^{i+n}(G_n) \xrightarrow{\eta} H^{i+n}(G_{n-1}) \longrightarrow 0.$$

Define $w'_n = e \in H^n(G_n)$.

We claim that every element in $H^{i+n}(G_n)$ can be written uniquely as a polynomial in the $w'_1, \ldots, w'_n$. We do this by induction on i. For $-n \leq i < 0$ this follows because $H^{i+n}(G_n) \cong H^{i+n}(G_{n-1})$. Now let $x \in H^{i+n}(G_n)$. Then $\eta(x) = p(w_1, \ldots, w_{n-1})$ for a unique p. Then $x = p(w'_1, \ldots, w'_{n-1}) + w'_n \cdot y$, where y comes from $H^i(G_n)$. However, since $i = (i-n)+n$, by the inductive hypothesis for $i-n$ we know that y can be written as a unique polynomial $q(w'_1, \ldots, w'_n)$. □

Remark 5.15. This uniquely characerizes the w_i in terms of the Thom classes: each w_i appears as the image of a Thom class in $H^*(G_i)$, which is then uniquely translated into the w_i for G_n via the sphere bundle construction.

An analogous theorem holds for the complex case:

Theorem 5.16. *Let GU_n be the Grassmannian of n-dimensional subspaces in $\mathbf{C}^\infty$. Then*

$$H^*(GU_n) \cong \mathbb{Z}[c_1, \ldots, c_n] \qquad \textit{with } |c_i| = 2i.$$

[e]We could also start it at $n = 1$, since we have already proved this result for projective space; however, this way we get an alternate proof of that, as well.

The moral of this story: cohomology is computable, and the cohomology of Grassmannians has a very nice universal characterization.

Moreover, if we analyze the proof we'll see that we proved a somewhat stronger statement than we were originally going for.

Corollary 5.17. *The homomorphism* $\eta\colon H^*(G_n) \longrightarrow H^*(G_{n-1})$ *takes* w_n *to* 0 *and maps all other generators correspondingly.*

Further Reading

A discussion of this material from the point of view of differential topology can be found in [BT13, Chapter 6, 12]; this exposition also fills in all of the details about manifolds that we have skipped. The differential geometry that is necessary to prove the Poincaré–Hopf Theorem is in [BT13, Chapters 1-5]. A proof of the Poincaré–Hopf Theorem can also be found in [MS74, Chapter 12].

Exercises and Extensions

5.1 Prove Theorem 5.16 by following the same outline as the proof of the real case.

5.2 Let M be a smooth manifold of dimension n. For pairs of smooth submanifolds α, β that intersect transversely (informally: tangent spaces interesect as little as possible) we can define

$$[\alpha] \cdot [\beta] = [\alpha \cap \beta].$$

(a) Let $M = S^1 \times S^1$. Prove that every homology class in $H_*(M; \mathbb{Z}/2)$ can be represented by a submanifold. Prove that the structure above induces a commutative ring structure on $H_*(M; \mathbb{Z}/2)$ which is isomorphic to the ring $H^*(M; \mathbb{Z}/2)$.

(b) Repeat part (a) for $\mathbf{R}P^2$.

5.3 In fact, a much more general statement than Exercise 5.2 is true: for homology classes represented by the image of the fundamental class of a manifold, the cup product of their Poincare duals is the Poincare dual of their intersection. (But proving this statement is beyond the scope of this book.) Assuming this statement prove that $H^*(\mathbf{R}P^n) \cong \mathbb{Z}/2[x]/x^{n+1}$.

5.4 Suppose that $p\colon E \longrightarrow B$ is a vector bundle, that B is compact and that E has an everywhere-nonzero section. Prove that the Gysin sequence splits into short exact sequences

$$0 \longrightarrow H^i(B) \longrightarrow H^i(S(E)) \longrightarrow H^{i-n+1}(B) \longrightarrow 0.$$

(Hint: start with a trivial bundle.)

5.5 The Euler class and the Gysin sequence exists for any sphere bundle over B. (This requires spectral sequences to prove; see for example [BT13, Chapter 14].) Consider the fiber bundle

$$S^{n-1} \longrightarrow V_2(\mathbf{R}^{n+1}) \longrightarrow S^n.$$

(a) Define the map $\rho\colon \mathbf{R}P^{n-1} \longrightarrow SO(n)$ by mapping a vector $v \in \mathbf{R}P^{n-1}$ to the composition $r(v)r(e_1)$, where $r(v)$ is the reflection in the plane orthogonal to v and e_1 is the standard unit vector. For any tuple $I = (i_1, \ldots, i_m)$ with each $i_j < n$, define the map

$$\rho_I\colon \mathbf{R}P^{i_1} \times \cdots \times \mathbf{R}P^{i_m} \longrightarrow SO(n)$$

to map $(v_1, \ldots, v_m)$ to $\rho(v_1) \cdots \rho(v_m)$. Prove that, as I ranges over all such tuples, this gives a CW structure on $SO(n)$.

(b) Prove that $V_2(\mathbf{R}^{n+1})$ is homeomorphic to the coset space $SO(n)/SO(2)$ and that the above cell structure descends to a cell structure on $V_2(\mathbf{R}^{n+1})$.

(c) Prove that $H^*(V_2(\mathbf{R}^{n+1}))$ is $\mathbb{Z}$ in degrees 0 and $2n-1$ and $\mathbb{Z}/2$ in degree n.

(d) Use the Gysin sequence to prove that the Euler class is $2 \in \mathbb{Z} \cong H^n(S^n)$.

Chapter 6

Characteristic classes

Consider again the formula $\mathrm{Vect}_n(B) \cong [B, G_n]$ from the Classification Theorem. In an ideal world, we would be able to identify the set $[B, G_n]$ for all B, and this would give a complete characterization of vector bundles. Unfortunately, at the moment this is beyond our mathematical capabilities. But just because we cannot compute this set explicitly does not mean that we cannot use it to construct some invariants of elements of vector bundles: here, by an *invariant* we mean a function $\mathrm{Vect}_n(B) \longrightarrow S$ for some set S.

The general goal of invariants is for them to be *computable* and *powerful*[a]. These are often properties that need to be played against one another: the more powerful the invariant, the more difficult it will be to compute, and vice versa. To give one extreme example, the bijection $\mathrm{Vect}_n(B) \longrightarrow [B, G_n]$ gives an invariant on $[B, G_n]$ which always distinguishes isomorphism classes of vector bundles; however, as this is not any more computable than the original problem, this is generally not considered a good invariant. On the other extreme, we could consider the constant function $\mathrm{Vect}_n(B) \longrightarrow \mathbb{Z}$ taking all bundles to 0. This is extremely computable, but not powerful enough to distinguish any isomorphism classes of vector bundles. Our goal, therefore, is to exchange some amount of power for some amount of computability.

The idea for the construction we will use is simple. Suppose that there existed a space Z such that we could compute both $[B, Z]$ and $[G_n, Z]$. For any map $f \colon B \longrightarrow G_n$ there exists a *precomposition* function

$$[G_n, Z] \longrightarrow [B, Z] \qquad g \longmapsto g \circ f.$$

[a] I.e. able to distinguish more types of bundles

If we fix a class $[g] \in [G_n, Z]$ then this precomposition gives an invariant

$$\mathrm{Vect}_n(B) \cong [B, G_n] \xrightarrow{[g]} [B, Z] \qquad [f] \longmapsto [g \circ f].$$

Thus a choice of Z and $[g]$ exactly determines an invariant. When Z comes from a cohomology theory, this is a *characteristic class*: an element in a cohomology group that encodes a fundamental property of the vector bundle.

In this chapter, Section 6.1 defines characteristic classes formally. Section 6.2 gives an exposition of some of the classical computations with characteristic classes. Section 6.3 proves the first result relating to parallelizability of projective spaces, showing that they can be parallelizable only when their dimension is one less than a power of 2.

6.1 The definition of characteristic classes

Definition 6.1. Let $h^*\colon \mathbf{Top}_*^{\mathrm{op}} \longrightarrow \mathbf{AbGp}^{\mathbb{Z}}$ be a reduced cohomology theory and let $\xi \in h^*(G_n)$ be a fixed element of $h^*(G_n)$. For every n-dimensional real vector bundle $p\colon E \longrightarrow B$ this determines an element $f^*\xi \in h^*(B)$ by pulling back along a classifying map f for E; this is denoted, by an abuse of notation, $\xi(E)$. This assignment is called a *characteristic class with values in h^**.

When h^* is singular cohomology, with any coefficients, this is simply called a *characteristic class*.

Directly from this definition we can conclude two properties of all characteristic classes with values in h^*:

(I) $\xi(E) \in h^*(B)$ depends only on the isomorphism class of E.

(P) For any map $f\colon B' \longrightarrow B$, $\xi(f^*(E)) = f^*(\xi(E))$.[b]

For the categorically-minded, a function ξ satisfying (I) and (P) can be described functorially as follows. Let $\mathbf{Bund}_n$ be the category with

objects n-dimensional real vector bundles $p\colon E \longrightarrow B$, and

morphisms $(p\colon E \longrightarrow B) \longrightarrow (p'\colon E' \longrightarrow B')$ consist of a map $f\colon B' \longrightarrow B$ together with an isomorphism of bundles over B', $\phi\colon E' \xrightarrow{\cong} f^*E$.

[b]The two f^* denote different things: in the first case it denotes the pullback bundle and in the second the map induced on cohomology.

Let $\mathbf{AbGp}^{\mathbb{Z},e}$ be the category with

objects pairs (A^*, ξ) of a $\mathbb{Z}$-graded abelian group A^* and an element $\xi \in A^*$, and

morphisms $(A^*, \xi) \longrightarrow (A'^*, \xi')$ consist of a graded homomorphism $\varphi\colon A^* \longrightarrow A'^*$ such that $\varphi(\xi) = \xi'$.

There is a forgetful functor $\mathbf{AbGp}^{\mathbb{Z},e} \longrightarrow \mathbf{AbGp}^{\mathbb{Z}}$ given by forgetting ξ. In this language, a characteristic class with values in h^* is just a functor $\mathbf{Bund}_n \longrightarrow \mathbf{AbGp}^{\mathbb{Z},e}$, which, after composing with the forgetful functor, is the functor taking $p\colon E \longrightarrow B$ to $h^*(B)$.

Remark 6.2. There is an alternate categorical viewpoint on this, using natural transformations. From this viewpoint, Lemma 6.3 can be proved using the Yoneda lemma.

It turns out that this perspective is exactly equivalent to the homotopical perspective above, in the sense that functors with values in cohomology groups correspond to characteristic classes:

Lemma 6.3. *Functors $\mathbf{Bund}_n \longrightarrow \mathbf{AbGp}^{\mathbb{Z},e}$ which are equal to h^* after composing with the forgetful functor are in natural bijection with characteristic classes with values in h^* (and thus with elements of $h^*(G_n)$). In other words, functions on vector bundles satisfying (I) and (P) are in natural bijection with characteristic classes.*

Proof. The lemma describes a function

$$\{F \in \mathrm{Fun}(\mathbf{Bund}_n, \mathbf{AbGp}^{\mathbb{Z},e}) \mid U \circ F = h^*\} \longrightarrow \{\text{char. classes}\}$$

given by $F \longmapsto F(\gamma_n) \in h^*(G_n)$. We claim that this function is a bijection.

Let $\xi \in h^*(G_n)$ be any element. Define $F(p\colon E \longrightarrow B)$ to be $f^*\xi \in h^*(B)$, where $f\colon B \longrightarrow G_n$ is a classifying map for E. This gives a well-defined functor, as desired, showing that the above function is surjective.

Let $F, F'\colon \mathbf{Bund}_n \longrightarrow \mathbf{AbGp}^{\mathbb{Z},e}$ be two functors which after composing with the forgetful functor are equal to h^*. Suppose that $F(\gamma_n) = F'(\gamma_n)$; we claim that this implies that $F = F'$. Indeed, for any other bundle $p\colon E \longrightarrow B$ there exists a map $f\colon B \longrightarrow G_n$ such that there exists an isomorphism $\varphi\colon E \longrightarrow f^*\gamma_n$. This gives a morphism in $\mathbf{Bund}_n$, and the image of this morphism is the morphism taking $F(\gamma_n) \in h^*(G_n)$ to $f^*F(\gamma_n) \in h^*(B)$. Since f is unique up to homotopy, the image of this morphism in $\mathbf{AbGp}^{\mathbb{Z},e}$ is unique. Thus F is uniquely determined by $F(\gamma_n)$, and must therefore be equal to F', as desired. □

Remark 6.4. It is possible to generalize this definition to bundles with extra structure, such as oriented bundles. This restricts the source category of the definition to only allow bundles and morphisms which have (resp. preserve) this extra structure. Otherwise, the definition is unchanged.

We have already constructed one example of a characteristic class:

Example 6.5. The $\mathbb{Z}/2$-Euler class is a characteristic class with values in $H^n(\cdot;\mathbb{Z}/2)$ for n-dimensional vector bundles. If we restrict to oriented bundles, the Euler class is a characteristic class with values in $H^n(\cdot;\mathbb{Z})$ for n-dimensional vector bundles.

For now we only have one cohomology theory to work with: singular cohomology. By Theorem 5.14, with $\mathbb{Z}/2$-coefficients this cohomology is a polynomial ring on the variables w_i with $|w_i| = i$. This gives us a good collection of characteristic classes to work with, right off the bat.

In fact, the structure of the cohomology of Grassmannians gives us more than we had hoped. Once w_i appears in $H^*(G_n)$ it is uniquely determined for $H^*(G_{n+k};\mathbb{Z}.2)$ for all $k \geq 0$. Thus these characteristic classes for bundles of dimension n determines a characteristic class for all bundles of dimension *at least* n. Given that we know good generators for $H^*(G_n)$ and these come with nice well-defined gradings, it makes sense to name them.

Definition 6.6. The i-th *Stiefel–Whitney class* is the characteristic class associated to $w_i \in H^*(G_n;\mathbb{Z}/2)$ (where when $n < i$ we set $w_i = 0$).

This makes sense for *all* vector bundles. Let us explore some consequences of this definition. The first is direct from the definition:

Lemma 6.7. *For an n-dimensional vector bundle $p\colon E \longrightarrow B$, $w_n(E) = e(E)$, where $e(E)$ is the $\mathbb{Z}/2$-Euler class.*

Unlike the Euler class, we can compare Stiefel–Whitney classes across different dimensions. This makes certain calculations much simpler:

Lemma 6.8. *For any vector bundle $p\colon E \longrightarrow B$ and all i,*

$$w_i(E \oplus \epsilon^k) = w_i(E).$$

In particular, $w_i(\epsilon^k) = 0$ for $i > 0$.

Proof. It suffices to check this for $k = 1$ (the general case then follows directly by induction). It also suffices to prove this for $E = \gamma_n$. To see this, let E be any rank n vector bundle, and let $f\colon B \longrightarrow G_n$ be its classifying map. Then there is a diagram

$$\begin{array}{ccc} E \oplus \epsilon^1 & \longrightarrow & \gamma_n \oplus \epsilon^1 \\ \downarrow & & \downarrow \\ B & \xrightarrow{f} & G_n \end{array}$$

Thus $w_i(E \oplus \epsilon^1) = f^*(w_i(\gamma_n \oplus \epsilon^1))$ (using Lemma 3.11). By Lemma 5.13 $w_i(\gamma_n \oplus \epsilon^1) = w_i(\gamma_n)$. But by definition, $w_i \in H^i(G_{n+1}; \mathbb{Z}/2)$ corresponds to $w_i \in H^i(G_n; \mathbb{Z}/2)$ under pullback along this map, proving the lemma. □

In particular, all Stiefel–Whitney classes associated to trivial bundles are 0.

Apart from being nicely defined for all vector bundles, the Stiefel–Whitney classes satisfy a beautiful relation called the *Whitney sum formula*.

Theorem 6.9 (Whitney sum formula)**.** *For vector bundles E and E' over a common base B,*

$$w_i(E \oplus E') = \sum_{j+k=i} w_j(E) \smile w_k(E').$$

Proof. Let $f\colon B \longrightarrow G_m$ and $f'\colon B \longrightarrow G_n$ be the classifying maps of E and E', respectively. Then the classifying map of $E \oplus E'$ is

$$B \xrightarrow{\Delta} B \times B \xrightarrow{f \times f'} G_m \times G_n \xrightarrow{\oplus} G_{m+n}.$$

Since characteristic classes are preserved under pullbacks, it suffices to understand how pulling back along this composition transforms the generator $w_i \in H^*(G_{m+n})$. By the Kunneth Formula, $H^*(G_m \times G_n) \cong H^*(G_m) \otimes H^*(G_n)$ and $H^*(B \times B) \cong H^*(B) \otimes H^*(B)$. Thus

$$\oplus^* w_i = \sum_{j+k=i} a_{jk}(w_j \otimes w_k) \in H^*(G_m) \otimes H^*(G_n) \qquad (6.10)$$

for some $a_{jk} \in \mathbb{Z}/2$. Pulling back w_j along f (resp. f') gives $w_j(E)$ (resp. $w_j(E')$), and pulling back along Δ^* takes tensor products to cup products. It therefore follows from (6.10) that

$$w_i(E \oplus E') = \sum_{j+k=i} a_{jk} w_j(E) w_k(E)' \in H^*(B);$$

in order to prove the theorem it suffices to check that $a_{jk} = 1$ for all j, k with $j + k = i$.

We prove this by induction on $m+n$. Recall the map τ from Example 3.17 which induces $\eta\colon H^*(G_n) \longrightarrow H^*(G_{n-1})$; thus τ classifies the bundle $\gamma_{n-1} \oplus \epsilon^1$ on G_{n-1}. Then

$$(1 \times \tau)^*(\oplus^* w_i) = w_i(\gamma_m \oplus \gamma_{n-1} \oplus \epsilon^1) = w_i(\gamma_m \oplus \gamma_{n-1}),$$

by Lemma 6.8. The right-hand bundle is now over $G_m \times G_{n-1}$. By induction, the right-hand side is equal to $\sum_{j+k=i} w_j \otimes w_k$. Since $\tau^*(w_k) = w_k$ for $k < n$ and to 0 otherwise,

$$\oplus^* w_i \equiv \sum_{j+k=i,\ k<n} w_j \otimes w_k \pmod{1 \otimes w_n}.$$

Symmetrically, we can also conclude that

$$\oplus^* w_i \equiv \sum_{j+k=i,\ j<m} w_j \otimes w_k \pmod{w_m \otimes 1}.$$

Thus by the Chinese Remainder Theorem,

$$\oplus^* w_i \equiv \sum_{j+k=i} w_j \otimes w_k \pmod{w_m \otimes w_n}.$$

When $i < m+n$ we know that this must be equality, since $w_m \otimes w_n$ has grading $m+n$. When $i > m+n$ $w_i = 0$ and so its pullback must also be 0; this also follows from this formula. Therefore it simply remains to check that $\oplus^* w_{m+n} = w_m \otimes w_n$.

The element w_m is exactly the $\mathbb{Z}/2$-Euler class of γ_m, and the element w_n is exactly the $\mathbb{Z}/2$-Euler class of γ_n; thus, by Proposition 5.8 the Euler class of $\gamma_m \times \gamma_n$ is $w_m \otimes w_n$. On the other hand, $\oplus^* w_{m+n}$ is the pullback of the $\mathbb{Z}/2$-Euler class of γ_{m+n} along the classifying map, so it is the $\mathbb{Z}/2$-Euler class of $\oplus^* w_{m+n}$, and the two are equal, as desired. □

Remark 6.11. We could define the *total Whitney class* of a vector bundle as an element of $H^*(B)$, to be $w = 1 + w_1 + w_2 + \cdots$; note that this is well-defined because for every bundle only finitely many of these are nonzero. Then the Whitney sum formula says that $w(E \oplus E') = w(E) \smile w(E')$. This involves adding up elements in different cohomology degrees. This class therefore does not contain geometric information, as it mixes different-dimensional invariants in a geometrically incoherent manner. In light of this, we will not use the total Whitney class in this book.

We have shown that the Steifel–Whitney classes satisfy the following properties:

(SW1) For every $j \geq 0$ you can assign a Steifel–Whitney class $w_j(E) \in H^j(B)$. (We define $w_0(E) = 1$.) $w_i(E) = 0$ if i is greater than the rank of E.

(SW2) Given any map $f: B' \longrightarrow B$ and any bundle $E \longrightarrow B$, $w_j(f^*E) = f^*w_j(E)$.

(SW3) For any bundles E and E' over B,

$$w_i(E \oplus E') = \sum_{j+k=i} w_j(E)w_k(E').$$

(SW4) For the universal bundle $\gamma_n \longrightarrow G_n$, $w_n(\gamma_n) \neq 0$.

It turns out that these four properties uniquely characterize the Steifel–Whitney classes. Although we do not prove this here, the interested reader should see [MS74, Chapter 7]. In [MS74] the authors define Stiefel–Whitney classes as classes satsifying properties (SW1)-(SW4), and the referenced chapter proves that these are unique. Proving that they exist in this narrative turns out to be far more complicated, and is done in [MS74, Chapter 8].

Remark 6.12*.* As before, this works exactly the same way for complex vector bundles, except that all of the above classes live in $\mathbb{Z}$-coefficient cohomology. They are called the *Chern classes*.

6.2 Some computations with Steifel–Whitney classes

In the Hairy Ball Question we asked if the tangent bundle of a manifold has an everywhere-nonzero section. More generally, we could ask *how many* everywhere-linearly-independent sections a tangent bundle could have. In the same way that the Euler class being nonzero implies that there can be *no* nonvanishing sections, Steifel–Whitney classes can be used to put an upper bound on the number of linearly independent nonvanishing sections that can be constructed for a bundle simultaneously.

Proposition 6.13. *If E is a rank-n bundle over B with an everywhere-nonzero section s then $w_n(E) = 0$. If E has k everywhere-independent sections then*

$$w_n(E) = w_{n-1}(E) = \cdots = w_{n-k+1}(E) = 0.$$

Proof. Given k everywhere-independent sections $s_1, \ldots, s_k$, E contains a trivial subbundle E' given by the spans of $s_1, \ldots, s_k$. Letting $E'' = (E')^{\perp}$, we see that $E \cong E' \oplus E''$. Then $w_i(E) = w_i(E'')$ by the Whitney sum formula and the fact that $w_i(E') = 0$ for $i > 0$. Since E'' has rank $n-k$, all Steifel–Whitney classes above dimension $n-k$ must be 0. □

The set of all isomorphism classes of bundles over B forms a monoid under the operation $\oplus$, with the identity being the 0-dimensional bundle given by the identity map. This operation does not have an inverse, as it always increases the dimension of a bundle. However, as trivial bundles are particularly simple, it makes sense to ask if there is an "almost-inverse," up to trivial bundles, to a given bundle.

Definition 6.14. A bundle E' is *complementary* to E if $E \oplus E' \cong \epsilon^n$ for some n.

With the tools currently at our disposal, we can show that over compact (but not well-partitionable!) bases such bundles always exist. To prove this, we first formalize the relationship between the Stiefel–Whitney classes of complementary bundles.

Proposition 6.15. *For every k there exists a polynomial $q_k(x_1, \ldots, x_k)$, such that whenever $E \oplus E' \cong \epsilon^n$, $w_k(E') = q_k(w_1(E), \ldots, w_k(E))$.*

Proof. We construct q_k inductively. When $k = 0$ setting $q_k = 1$ works. For $k = 1$, by the Whitney sum formula, $w_1(E) + w_1(E') = w_1(E \oplus E') = 0$. Thus $q_1(x) = -x$ works. Now suppose that polynomials $q_1, \ldots, q_{k-1}$ exist. By the Whitney sum formula,

$$w_k(E \oplus E') = w_k(E) + w_{k-1}(E)w_1(E') + \cdots + w_1(E)w_{k-1}(E') + w_k(E') = 0.$$

Solving for $w_k(E')$ gives

$$w_k(E') = -\sum_{i=1}^{k} w_i(E) q_{k-i}(w_1(E), \ldots, w_{k-i}(E)),$$

which shows that the polynomial

$$q_k(x_1, \ldots, x_k) = -\sum_{i=1}^{k} x_i q_{k-i}(x_1, \ldots, x_{k-i})$$

works. □

Since the polynomial q_k is independent of E, the class $q_k(w_1(E), \dots, w_k(E))$ exists even when E' does not.

Definition 6.16. Write $\overline{w}_k(E)$ for $q_k(w_1(E), \dots, w_i(E))$ for the *dual Steifel–Whitney classes* of E.

As a special case, we get the following:

Lemma 6.17 (Whitney duality theorem)**.** *Let TM be the tangent bundle of a manifold M with a chosen embedding $\iota\colon M \hookrightarrow \mathbf{R}^N$ into Euclidean space, and let ν be the normal bundle. Then*

$$w_i(\nu) = \overline{w}_i(TM).$$

In particular, note that the characteristic classes of the normal bundle are independent of the choice of embedding.

The fact that it is possible to find a bundle that adds to a trivial bundle is not restricted to tangent bundles of manifolds: such a bundle exists for any bundle over a compact base.

Proposition 6.18. *For any bundle $p\colon E \longrightarrow B$ where B is compact there exists a bundle E' such that $E \oplus E' \cong \epsilon^k$ for some k.*

Proof. As discussed in the proof of Theorem 3.29, it suffices to construct a map $g\colon E \longrightarrow \mathbf{R}^N$ which is a linear injection on fibers. This produces an embedding of E into $B \times \mathbf{R}^n$, and the orthogonal complement of each fiber of E is the desired bundle.

We proceed as in the proof of Theorem 3.29, except that we must have a finite open cover so that we can embed into a finite-dimensional space. For each point $x \in B$ there exists a U_x over which E is trivial. By Urysolhn's Lemma there is a map $\varphi_x\colon B \longrightarrow [0,1]$ which is 0 outside U_x and nonzero at x. Then $\{\varphi_x^{-1}(0,1]\}$ is an open cover of B; since B is compact, it has a finite subcover $\{\varphi_{x_i}^{-1}(0,1]\}_{i=1}^k$. Write $h_i\colon p^{-1}(U_{x_i}) \longrightarrow U_{x_i} \times \mathbf{R}^n \xrightarrow{pr_2} \mathbf{R}^n$ for the composition of the local trivialization at x_i and the projection onto the second coordinate. Multiplying by $\varphi_i(p(e))$ extends this to all of E by 0. We then define

$$g(e) = (\varphi_1(p(e))h_1(e), \dots, \varphi_k(p(e))h_k(e)) \in \mathbf{R}^{nk}.$$

□

It is important to note that Steifel–Whitney classes are *not* a complete invariant. To illustrate this we give an example of a nontrivial bundle all of whose Steifel–Whitney classes are 0.

Example 6.19. Consider TS^n. When S^n is embedded in $\mathbf{R}^{n+1}$, the normal bundle is trivial. Thus $w_i(\nu) = 0$ for $i > 0$. Since the sum of the two bundles is also trivial, we must have $w_i(TS^n) = 0$ for $i > 0$. However, by the Poincare–Hopf theorem, the bundle TS^n is in general nontrivial.

It is also important to keep in mind that complementary bundles do not necessarily exist.

Example 6.20. Consider the universal line bundle γ_{1n} over $\mathbf{R}P^n$. From before we know that $H^*(\mathbf{R}P^n) \cong \mathbb{Z}/2[x]/(x^{n+1})$. Let $\iota\colon \mathbf{R}P^n \longrightarrow G_1$ be the inclusion, and note that $w_i(\gamma_{1n}) = \iota^* w_i(\gamma_1) = 0$ for $i > 1$. Consider the inclusion $j\colon \mathbf{R}P^1 \longrightarrow \mathbf{R}P^n$. As shown before, $j^* w_1(\mathbf{R}P^n) = w_1(\mathbf{R}P^1) \neq 0$. Thus $w_1(\mathbf{R}P^n) \neq 0$, so it must be x.

From this we can conclude that there is no bundle E such that $\gamma_1 \oplus E$ is trivial. Indeed, if such a bundle existed we must have $w_i(E) = x^i$ for all i—which means that the bundle is infinite-dimensional, contradicting the definion of a vector bundle.

Remark 6.21. The previous example is one reason that the "total class" is such a prevalent object. If we were to speak of it in those terms, we could simply write that $w(\gamma_{1n})w(\gamma_{1n}^{\perp}) = 1$ and note that $w(\gamma_{1n}) = 1 + x$ and

$$(1 + x)(1 + x + \cdots + x^n) = 1 \in \mathbb{Z}/2[x]/(x^{n+1}).$$

Our next example is an interesting exploration of how vector bundles can become simpler when trivial bundles are added in.

Example 6.22. Let $T\mathbf{R}P^n$ be the tangent bundle to $\mathbf{R}P^n$. As discussed in Exercise 3.10, $T\mathbf{R}P^n \cong \mathrm{Hom}(\gamma_{1n}, \gamma_{1n}^{\perp})$.

We do not have a method for computing Steifel–Whitney classes of Hom-bundles. It turns out that adding a trivial bundle can make such a computation much simpler, while not changing the value of the classes.

Consider the sum $T\mathbf{R}P^n \oplus \epsilon^1$. We can write $\epsilon^1 \cong \mathrm{Hom}(\gamma_{1n}, \gamma_{1n})$, since the latter is a line bundle with an everywhere-nonzero section. Thus

$$\begin{aligned} T\mathbf{R}P^n \oplus \epsilon^1 &\cong \mathrm{Hom}(\gamma_{1n}, \gamma_{1n}^{\perp}) \oplus \mathrm{Hom}(\gamma_{1n}, \gamma_{1n}) \cong \mathrm{Hom}(\gamma_{1n}, \gamma_{1n} \oplus \gamma_{1n}^{\perp}) \\ &\cong \mathrm{Hom}(\gamma_{1n}, \epsilon^{n+1}) \cong \mathrm{Hom}(\gamma_{1n}, \epsilon^1)^{\oplus n+1}. \end{aligned}$$

By the Whitney sum formula it therefore remains to consider $\mathrm{Hom}(\gamma_{1n}, \epsilon^1)$.

By definition,

$$\gamma_{1n} = \{(\ell, x) \in \mathbf{R}P^n \times \mathbf{R}^{n+1} \mid x \in \ell\}.$$

There is a morphism of bundles $\mathrm{Hom}(\gamma_{1n}, \epsilon^1) \longrightarrow \gamma_{1n}$ given by

$$(\ell, T\colon \ell \longrightarrow \mathbf{R}) \longmapsto \left(\ell, \frac{T^{-1}(1)}{\|T^{-1}(1)\|^2}\right).$$

Here, the norm is taken using the embedding $\gamma_{1n} \subseteq G_1(\mathbf{R}^n) \times \mathbf{R}^n$. This is a linear injection on fibers, so it is an isomorphism. Thus in fact

$$T\mathbf{R}P^n \oplus \epsilon^1 \cong \gamma_{1n}^{\oplus n+1}.$$

Since adding a trivial bundle does not change Steifel–Whitney classes, we can use induction to conclude that

$$w_i(T\mathbf{R}P^n) = \binom{n+1}{i} x^i \pmod 2.$$

As a final application we show that in every dimension there exist bundles which are "atomically" of that dimension; i.e. which cannot be split as a Whitney sum of lower-dimensional bundles. While this is not surprising, it is still a useful reminder of the structure of vector bundles.

Example 6.23. As a generalization of Proposition 6.13, we can use Steifel–Whitney classes to check that the universal bundle γ_n cannot be written as a sum of lower-dimensional bundles. Suppose that $\gamma_n \cong E \oplus E'$, with $\dim E, \dim E' < \gamma_n$. Then

$$\begin{aligned} w_n = w_n(\gamma_n) &= \sum_{i+j=n} w_i(E) w_j(E') \\ &= \sum_{\substack{i+j=n \\ i \le \dim E \\ j \le \dim E'}} \in \mathbb{Z}/2[w_1, \dots, w_{n-1}]. \end{aligned}$$

This is a contradiction, since w_n is independent of $w_1, \dots, w_{n-1}$. Thus for every dimension there exist "irreducible" bundles which cannot be written as the sum of lower-dimensional bundles.

6.3 Parallelizability of $\mathbf{R}P^n$

Example 6.22 above has enough interesting geometric consequences that are consequential enough that they are worth discussing in detail. Whenever $\mathbf{R}P^n$ is parallelizable, S^n must be, as well. Thus limiting the dimensions in which $\mathbf{R}P^n$ is parallelizable will help narrow down the dimensions in which S^n can be parallelizable.

Proposition 6.24. $\mathbf{R}P^n$ *is parallelizable only if* $n = 2^k - 1$ *for some* k*.*

Proof. If M is parallizable then $w_i(M) = 0$ for all $i > 0$. From Example 6.22 we know that $w_i(M) = \binom{n+1}{i} x^i$; these are all 0 exactly when $n + 1 = 2^k$ for some k. □

In fact, the parallelizable real projective spaces are exactly $\mathbf{R}P^1$, $\mathbf{R}P^3$ and $\mathbf{R}P^7$; we will prove that these *are* parallelizable as a consequence of the next theorem, and we will show that the rest are not later when we introduce K-theory in Chapter 9.

Theorem 6.25 (Steifel)**.** *Suppose that there exists a bilinear product operation* $p\colon \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}^n$ *without zero divisors. Then the projective space* $\mathbf{R}P^{n-1}$ *is parallelizable.*

The map p is not required to be associative or to have a unit.

Proof. Our goal is to construct a trivialization of $T\mathbf{R}P^{n-1}$ using p. Thus we want to construct $n - 1$ linearly independent sections of $T\mathbf{R}P^{n-1} \cong \mathrm{Hom}(\gamma_{1(n-1)}, \gamma_{1(n-1)}^{\perp})$.

We begin by using p to construct linear maps $T_1, \ldots, T_n\colon \mathbf{R}^n \longrightarrow \mathbf{R}^n$ such that T_1 is the identity and, for all nonzero x, $T_1(x), \ldots, T_n(x)$ are linearly independent. Let $e_1, \ldots, e_n$ be the standard basis of $\mathbf{R}^n$, and consider the map $S = p(-, e_1)$. This is a linear map $\mathbf{R}^n \longrightarrow \mathbf{R}^n$ with trivial kernel (since there are no zero divisors), so it is an isomorphism. We let

$$T_i(z) = p(S^{-1}(z), e_i).$$

Then

$$T_1(x) = p(S^{-1}(x), e_1) = x,$$

by definition. Moreover, for a nonzero x and any $\lambda_1, \ldots, \lambda_n \in \mathbf{R}$,

$$\sum_{i=1}^{n} \lambda_i T_i(x) = \sum_{i=1}^{n} \lambda_i p(S^{-1}(x), e_i) = p\left(S^{-1}(x), \sum_{i=1}^{n} \lambda_i e_i\right).$$

Since $S^{-1}(x) \neq 0$ and p has no zero divisors, the only way that the left-hand side can equal 0 is if $\lambda_1 = \cdots = \lambda_n = 0$; in other words, $T_1(x), \ldots, T_n(x)$ are linearly independent, as desired.

Given a linear map $T\colon \mathbf{R}^n \longrightarrow \mathbf{R}^n$ define the map $\overline{T}\colon \ell \longrightarrow \ell^{\perp}$ for any line ℓ through the origin by taking $x \in \ell$ to the orthogonal projection of $T(x)$ onto $\ell^{\perp}$. Then $\overline{T}$ is a section of $T\mathbf{R}P^{n-1}$. Moreover, if $T(x)$ and x are linearly independent for all nonzero x, this is an *everywhere-nonzero* section. Similarly, for all nonzero x, $\overline{T}_2(x), \ldots, \overline{T}_n(x)$ are everywhere linearly independent, and thus give $n - 1$ linearly independent sections of $T\mathbf{R}P^{n-1}$. □

Putting this together with Proposition 6.24, we see that $\mathbf{R}^n$ can be given a skew field structure only when n is a power of 2. When $n = 1, 2, 4, 8$ such structures exist: they are the real, complex, quaternionic and octonic structures. When we prove that $\mathbf{R}P^n$ is not parallelizable for $n > 7$ we will simultaneously show that these are the only skew field structures on $\mathbf{R}^n$.

For the last application of this section, we turn to a geometric consideration: when can a manifold M be immersed into $\mathbf{R}^n$? (Recall that an immersion, unlike an embedding, can self-intersect, but only transversely.)

Suppose that an n-manifold M is immersed into $\mathbf{R}^{n+k}$. Let ν be the normal bundle, so that $\nu \oplus TM = \epsilon^{n+k}$. Thus $\overline{w}_i(TM) = w_i(\nu)$; since ν is k-dimensional we must have $\overline{w}_i(TM) = 0$ for $i > k$.

For example, when $M = \mathbf{R}P^9$ the equation $w_i(TM) = x^i$ holds exactly when $i = 2, 8$, and thus $\overline{w}_i(TM) = x^i$ exactly when $i = 2, 4, 6$. Therefore if $\mathbf{R}P^9$ can be immersed into $\mathbf{R}^{9+k}$ we must have $k \geq 6$. When $M = \mathbf{R}P^{2^k}$ the equation $w_i(TM) = x^i$ holds exactly when $i = 1, n$; thus $\overline{w}_i(TM) = x^i$ exactly when $i = 1, 2, \ldots, 2^k - 1$. In particular, $\mathbf{R}P^{2^k}$ cannot be immersed in $\mathbf{R}^{2^k+\ell}$ unless $\ell \geq 2^k - 1$.

On the other hand, Whitney's immersion theorem [Whi44] states that any n-manifold can be immersed in $\mathbf{R}^{2n-1}$. We have thus shown that this bound is sharp: there is no $k > 1$ such that an n-manifold can always be immersed in $\mathbf{R}^{2n-k}$.

Further reading

The discussion in the previous two chapters can be done in the opposite order. The four properties of characteristic classes (SW1)-(SW4) suffice to determine them uniquely. Thus one can define characteristic classes axiomatically and use these axiomatics to compute the cohomology of Grassmannians. It is then necessary to prove that characteristic classes exist using a separate construction. See, for example, [MS74, Section 7], which uses Steenrod squares to construct the Stiefel–Whitney classes. The computations in this section are based off of the computations in [MS74, Section 4].

The question of exactly how many linearly independent vector fields there are on a sphere is extremely deep and intersting. The classic paper is [Ada62], and unfortunately there are not many modern expositions of it.

Exercises and Extensions

6.1 Assume that for every bundle $p\colon E \longrightarrow B$ it is possible to define classes $w_i(E) \in H^i(B;\mathbb{Z}/2)$ (for $i \geq 1$) such that properties (SW1)-(SW4) on page 84 hold. Prove that these must be equal to our definition of Stiefel–Whitney classes, and conclude that these four properties uniquely determine the Stiefel–Whitney classes.

6.2 Use Axioms (SW1)-(SW4) to compute the cohomology of Grassmannians. Thus the behavior of Stiefel–Whitney classes exactly reflects the structure of the cohomology of Grassmannians.

6.3 The *Chern classes* are defined analogously to Stiefel–Whitney classes for complex vector bundles, with generators $c_i \in H^{2i}(B;\mathbb{Z})$. Verify that these satisfy a (suitably modified) set of axioms (SW1)-(SW4) and are therefore also characterized by analogous properties.

6.4 Verify directly that the Euler class satisfies the Whitney sum formula.

6.5 (a) Let γ_1 be the canonical line bundle. Prove that $w_1(\gamma_1 \otimes \gamma_1) = w_1(\gamma_1) + w_1(\gamma_1)$. (Although the right hand side of the equation is 0, since we are working over $\mathbb{Z}/2$, it is important for the next part to prove the formula in this form.) Your proof should also work for c_1 from Exercise 6.3. Hint: since w_1 lives inside H^1 it suffices to prove this when restricted to the 2-skeleton of $G_1 \times G_1$. How does this allow us to simplify the space?

(b) Let L, L' be line bundles over X. Use the previous part to prove that $w_1(L \otimes L') = w_1(L) + w_1(L')$.

6.6 (For those comfortable with category theory.) Prove Lemma 6.3 using the Yoneda Lemma.

6.7 Prove Theorem 5.16. The proof follows analogously to the proof of Theorem 5.14 in a more straightforward manner, since the long exact sequence in the Gysin sequence will be concentrated in even degrees and thus split into short exact sequences.

Chapter 7

Cobordism invariants

In this chapter we take a detour into cobordism theory. We have three goals for this chapter:

(a) to explore an interesting new geometric invariant that turns out to be classifiable by a homotopical invariant,

(b) to construct a new example of a cohomology theory, and

(c) to completely classify cobordism classes using *Stiefel–Whitney numbers*, which are derived from Stiefel–Whitney classes.

In this chapter we focus on unoriented cobordism. There are many other types of cobordism theories, including oriented cobordism, framed cobordism, complex cobordism, and so on. All of these take manifolds with some sort of structure and require the cobordism to carry this structure as well. Thus, for example, with oriented cobordism we require an orientation on W that restricts to the correct orientations on M and N. For more on this, see the "Further Reading" section.

The word "cobordism" is often used in two different ways. One is as a witness to the relationship between two manifolds, as in "let W be a cobordism between M and N." This is the sense in which we mostly use the word in this chapter. This is also the definition one finds upon looking up the word. However, people will also say "cobordism is a cohomology theory," and this sentence does not make sense with the previous definition. When people say this, it is a shortcut for a more complicated statement: "there is a cohomology theory whose value on a manifold of large enough dimension is the groups of cobordism classes of submanifolds." As any spectrum that represents a cohomology theory also represents a homology theory, people will sometimes use the word "bordism" to discuss this associated homology theory—a terrible crime against understanding and language, as the definitions of "bordism" and "cobordism" in the literature are equivalent.[a] We discuss the construction of this cohomology theory in Section 7.3, although we do not prove that its values are cobordism classes in the general case; we only prove it for spheres. A more careful analysis that proves this claim in general is found in [Tho54, Chapter IV].

In this chapter, all manifolds considered will be smooth, so we omit smoothness assumptions everywhere but in the statements of theorems. Although many of the theorems are true with finite degrees of smoothness (which is discussed in detail in [Tho54, Chapter I]) we disregard these in the interest of brevity and clarity.

Section 7.1 introduces the definition of cobordism groups and develops some first properties of cobordisms. Section 7.2 defines the main invariant of cobordism classes—Steifel–Whitney numbers—and states the main theorem of the chapter: that Steiel–Whitney numbers distinguish cobordism classes. Section 7.3 is an aside on cohomology theories and stability; it shows that cobordism represents a new cohomology theory. Section 7.4 introduces the main new tool for analyzing manifolds: transversality. Section 7.5 contains the main meat of the proof of the theorem: the Pontrjagin–Thom construction.

7.1 Introducing cobordisms

How is a torus diffeerent from a Klein bottle? The usual answer that one is orientable and one is not. But from a more basic perspective, the torus has an "inside" and the Klein bottle does not. In other words: the torus is the boundary of a 3-manifold, but the Klein bottle is not. More generally, we can ask the question: when is a closed n-manifold the boundary of an $n+1$-manifold? This is the basic question of *cobordism*.

Definition 7.1. Let M and N be two n-manifolds. M and N are *cobordant* if there exists an $(n+1)$-manifold W such that $\partial W = M \amalg N$. W is called the *cobordism* between M and N.

[a]The prefix 'co', often used in mathematics to mean "dual to" here means "together." When people say "bordism" to mean the associated homology theory they are either thinking that adding another 'co' would cancel out with the first one, or simply wishing "bordism" and "cobordism" to look parallel to "homology" and "cohomology." This is clearly silliness; the correct way to refer to the homology theory associated to cobordism is "cocobordism."

Example 7.2. Any manifold is cobordant to itself, since the boundary of $M \times I$ is $M \amalg M$.

Example 7.3. The reader may be wearing some nontrivial cobordisms: between S^1 and $S^1 \amalg S^1$ (a pair of pants), between S^1 and $\emptyset$ (a sock), or between S^1 and $S^1 \amalg S^1 \amalg S^1$ (a shirt or dress). This last cobordism is also interesting because it can also be interpreted as a cobordism between $S^1 \amalg S^1$ (armholes) and $S^1 \amalg S^1$ (head and body holes) which is not homeomorphic to $(S^1 \amalg S^1) \times I$.

Example 7.4. Cobordisms are not unique. $S^1 \times I$ is a cobordism between S^1 and S^1, but so is the torus with two ends cut off. In fact, the structure of cobordisms is quite interesting; it is addressed in the well-known cobordism theorems, such as the h-cobordism theorem and the s-cobordism theorem. See for example [Mil65].

In order to dive deeper into the theory of cobordisms, we begin by defining the cobordism group.

Definition 7.5. The *unoriented cobordism group* $\mathfrak{N}_n$ is defined as follows. As a set, $\mathfrak{N}_n$ consists of the equivalence classes of isomorphism classes of n-manifolds up to cobordism. Addition is defined to be $\amalg$. The empty manifold is considered to be an n-manifold for all n.

Lemma 7.6. *$\mathfrak{N}_n$ is an abelian group.*

Proof. First we check that the operation is well-defined. Suppose that $[M] = [M]'$ and $[N] = [N']$. We need to check that $[M \amalg N] = [M' \amalg N']$. Let W be a cobordism between M and M' and W' a cobordism between N and N'. Then $W \amalg W'$ is a cobordism between $M \amalg N$ and $M' \amalg N'$, so the operation is well-defined.

The identity is $[\emptyset]$.[b] Note that $2[M] = 0$ for all M, since $\partial(M \times I) = [M \amalg M]$. Thus inverses exist in the group.

$\mathfrak{N}_n$ is abelian because $M \amalg N \cong N \amalg M$. □

The goal of this chapter is to compute the group $\mathfrak{N}_n$ for all n. From the proof of the lemma we know that $\mathfrak{N}_n$ has exponent 2: twice any element is 0. Thus, if we knew that the group is finitely generated, we could conclude that it must be isomorphic to $(\mathbb{Z}/2)^k$ for some k. This turns out to be true, and to actually be a consequence of a far stronger result:

[b]If it is desitable to not consider $\emptyset$ as an n-manifold, we can write the identity to be $[S^n]$, since S^n is always a boundary.

Theorem 7.7 ([Tho54, Théorème IV.12])**.** *The group* $\mathfrak{N}_* = \bigoplus_{n\geq 0} \mathfrak{N}_k$ *is a graded* $\mathbb{Z}/2$*-algebra with the product given by* $[M][N] = [M \times N]$*. As a graded algebra,*

$$\mathfrak{N}_* \cong \mathbb{Z}/2\big[x_i \,\big|\, i \geq 1, i \neq 2^j - 1\big], \qquad |x_i| = i.$$

The fact that the algebra is well-defined is shown in Exercise 7.3. We will not cover the entire proof of the theorem, although we will discuss several important aspects of it.

7.2 Steifel–Whitney numbers

We now consider a structure which is coarser than Steifel–Whitney classes, but which will suffice to discriminate between cobordism classes. These are the *Steifel–Whitney numbers*.

Definition 7.8. Let M be an n-manifold, and let $[M] \in H_n(M;\mathbb{Z}/2)$ be its $\mathbb{Z}/2$-fundamental class. Let $r_1, \ldots, r_n$ be nonnegative integers such that

$$r_1 + 2r_2 + \cdots + nr_n = n.$$

Then the cohomology class $w_1(TM)^{r_1} \cdots w_n(TM)^{r_n}$ is in $H^n(M)$. The $(r_1, \ldots, r_n)$-*Steifel–Whitney number* is

$$(w_1(TM)^{r_1} w_2(TM)^{r_2} \cdots w_n(TM)^{r_n})[M] \in \mathbb{Z}/2.$$

This is denoted

$$w_1^{r_1} \cdots w_n^{r_n}[M].$$

As a quick exercise with the definition, we observe the following property:

Lemma 7.9. *Let* M *and* N *be two* n*-manifolds. For any* $r_1, \ldots, r_n$ *with* $r_1 + 2r_2 + \cdots + nr_n = n$*,*

$$w_i^{r_1} \cdots w_n^{r_n}[M \amalg N] = w_1^{r_1} \cdots w_n^{r_n}[M] + w_1^{r_1} \cdots w_n^{r_n}[N].$$

Proof. The inclusions $M \longrightarrow M \amalg N$ and $N \longrightarrow M \amalg N$ give isomorphisms

$$H^n(M \amalg N) \cong H^n(M) \times H^n(N) \qquad \text{and} \qquad H_n(M \amalg N) \cong H_n(M) \times H_n(N).$$

The tangent bundles on M and N can be obtained as pullbacks of the tangent bundle on $M \amalg N$ via pullbacks along these inclusions. The image of $[M \amalg N]$ is exactly $([M], [N])$. The statement of the lemma follows from the fact that Steifel–Whitney classes commute with pulling back vector bundles. □

We now compute an example.

Example 7.10. Let us compute the Stiefel–Whitney numbers for projective spaces. Recall that

$$H^*(\mathbf{R}P^n) \cong \mathbb{Z}/2[x]/x^{n+1} \qquad \text{and} \qquad w_i(T\mathbf{R}P^n) = \binom{n+1}{i} x^i.$$

First, suppose that n is even. Then $w_n(T\mathbf{R}P^n) = (n+1)x^n \neq 0$, and thus $w_n[M] \neq 0$. Similarly, since $w_1(TM) = (n+1)x$, $w_1^n[M] \neq 0$. In genereal,

$$w_1^{r_1} \cdots w_n^{r_n}[M] = \binom{n+1}{r_1}\binom{n+1}{r_2} \cdots \binom{n+1}{r_n} \bmod 2.$$

Depending on $n+1$, these vary. For example, when $n = 2^k - 2$ all of these are nonzero; on the other hand, when $n = 2^k$ the only ones which are nonzero are $w_1^n[M]$ and $w_n[M]$.

Now suppose that $n = 2k-1$ is odd. Note that $(1+x)^{2k} = (1+x^2)^k$, and thus $\binom{2k}{2i} = \binom{k}{i} \pmod 2$ and $\binom{2k}{2i+1} = 0 \pmod 2$. Thus in particular for all i, $w_{2i+1}(TM) = 0$. Since any $w_1^{r_1} \cdots w_n^{r_n}$ must have at least one $r_i \neq 0$ for an odd i, we see that all Steifel–Whitney numbers are 0.

From this example we conclude see that Steifel–Whitney numbers contain much less information than the classes themselves. It turns out that they contain just enough information to classify when a manifold is a boundary of another manifold.

Theorem 7.11 (Pontrjagin–Thom)**.** *Let M be a smooth closed n-manifold. There exists a smooth compact $(n+1)$-manifold B with boundary M if and only if all Steifel–Whitney numbers of M are* 0*.*

The "only if" direction was proved by Pontrjagin, and we are able to prove it now. The "if" direction is far more difficult; it was proved by Thom in [Tho54] and will require the development of some more theory before we are able to prove it.

Proof of "only if" direction. Let B be a $n+1$-manifold with boundary M. Consider the long exact sequence in homology for the pair (B, M):

$$H_{n+1}(M) \longrightarrow H_{n+1}(B) \longrightarrow H_{n+1}(B, M) \xrightarrow{\partial} H_n(M) \xrightarrow{i_*} H_n(B).$$

The first of these is 0. Note that $i_*[M]$ is the homology class in B represented by M; this is zero because M is a boundary: the boundary of B. Thus $[M]$ is the image of a class $[B, M] \in H_{n+1}(B, M)$. For any $v \in H^n(M)$,

$$v[M] = v(\partial[B, M]) = (\delta v)[B, M],$$

where $\delta\colon H^n(M) \longrightarrow H^{n+1}(B, M)$.

Let TB be the tangent bundle to B. Restricting it to M, we note that it has an everywhere-nonzero section, taking each point of M to the outward-facing vector. The orthogonal complement to this is the tangent bundle to M, so

$$TB|_M \cong TM \oplus \epsilon^1.$$

Thus the Steifel–Whitney classes of $TB|_M$ are the same as those of TM. If we thus consider the exact sequence in cohomology

$$H^n(B) \xrightarrow{i^*} H^n(M) \xrightarrow{\delta} H^{n+1}(B, M)$$

we see that $w_k TM = i^* w_k TB|_M$ for all k. Thus

$$\begin{aligned} w_1^{r_1} \cdots w_n^{r_n}[M] &= w_1^{r_1} \cdots w_n^{r_n}(\partial[B, M]) = (\delta(w_1^{r_1} \cdots w_n^{r_n}))[B, M] \\ &= (\delta i^*(w_1^{r_1} \cdots w_n^{r_n}))[B, M] = 0. \end{aligned}$$

Here in the last step we are using that $\delta i^* = 0$. Thus all Steifel–Whitney numbers of M are 0, as desired. □

The upshot of the Pontrjagin–Thom theorem is that the Steifel–Whitney numbers can identify exactly when a manifold is a boundary. By using the group structure on $\mathfrak{N}_n$ we can conclude that Stiefel–Whitney numbers can be used to detect cobordant manifolds:

Corollary 7.12. *M and N are cobordant if and only if their Stiefel–Whitney numbers are equal.*

It still remains to prove the "if" direction of the Pontrjagin–Thom theorem. In order to prove this, we must first develop more of the theory of Thom spaces.

7.3 Stability

An important notion in homotopy theory is that of *stability*: properties about invariants of spaces that are preserved by the functor taking a pointed space to its suspension: $X \longmapsto \Sigma X$. For example, the reduced homology and cohomology groups of a space are stable by definition (see Definition 4.4): the suspension homomorphism

$$\partial\colon h^i(X) \longrightarrow h^{i+1}(\Sigma X)$$

is an isomorphism. Homotopy, in general, is *not* stable. For example, $\pi_3(S^2) \cong \mathbb{Z}$ but $\pi_4(\Sigma S^2) \cong \pi_4(S^3) \cong \mathbb{Z}/2$. In homotopy there is no suspension isomorphism, but there is a convenient "stabilization" homomorphism between homotopy groups. For any map of pointed spaces $f\colon X \longrightarrow Y$ smashing with a circle induces a map $\Sigma f\colon \Sigma X \longrightarrow \Sigma Y$ inducing a function $[X, Y] \longrightarrow [\Sigma X, \Sigma Y]$. Although this function is not usually a bijection, for homotopy groups it is eventually an isomorphism:

Theorem 7.13 (Freudenthal Suspension Theorem, [Hat02, Theorem 4.23])**.** *For an n-connected pointed CW-complex X, the stabilization homomorphism*

$$\pi_k(X) \longrightarrow \pi_{k+1}(\Sigma X)$$

is an isomorphism if $k \leq 2n$. In particular, the homomorphism

$$\pi_k(S^n) \longrightarrow \pi_{k+1}(S^{n+1})$$

is an isomorphism for $k \leq 2n$.

The goal of this section is to prove an analogous result for Thom spaces. The key observation here is that suspending a Thom space of a bundle E corresponds to adding a trivial bundle. As discussed in Example 4.20, the Thom space of a trivial bundle over any compact base B is simple to understand:

$$\mathrm{Th}(\epsilon^k) \cong S^k \wedge B_+,$$

and more generally

$$\mathrm{Th}(E \times E') \cong \mathrm{Th}(E) \wedge \mathrm{Th}(E') \qquad \text{and} \qquad \mathrm{Th}(E \oplus \epsilon^k) \cong \mathrm{Th}(E) \wedge S^k.$$

As before, $\tau\colon G_k \longrightarrow G_{k+1}$ (Example 3.17) has $\tau^*\gamma_{k+1} \cong \gamma_k \oplus \epsilon^1$. By the above discussion, we therefore have a homomorphism

$$\pi_{n+k}(\mathrm{Th}(\gamma_k)) \longrightarrow \pi_{n+k+1}(\Sigma\mathrm{Th}(\gamma_k)) \cong \pi_{n+k+1}\mathrm{Th}(\gamma_k \oplus \epsilon^1) \xrightarrow{\tau_*} \pi_{n+k+1}\mathrm{Th}(\gamma_{n+k+1}).$$

The key observation is that for large enough k this is an isomorphism:

Lemma 7.14. *For $k > n$, the group $\pi_{n+k}(\mathrm{Th}(\gamma_k))$ is independent of k.*

We already know this is true for cohomology groups: By the Thom isomorphism theorem, $H^{n+k}(\mathrm{Th}(\gamma_k)) \cong H^n(G_k)$. Since $n < k$, this is going to be the group

$$\mathbb{Z}/2\{w_1^{i_1} \cdots w_n^{i_n} \mid i_1 + 2i_2 + \cdots + ni_n = n\}.$$

This is clearly independent of k. The point of this lemma is that this is also true for homotopy groups.

This is the first occurrence of *stable homotopy groups* in this course. The proof is in [Tho54, Section II.5, Theorem II.7] (keep in mind that Thom uses the notation $MO(k)$ for $\mathrm{Th}(G_k)$). We will not prove this lemma here, as it involves some technical computations of cohomology groups of Grassmannians with general $\mathbb{Z}/p$ coeffiecients; the interested reader is directed do the proof in Thom's paper. The main ingredients of the proof are the Freudenthal suspension theorem and a variant of the Hurewicz theorem.

Example 7.15. Consider the sequence of spaces

$$\mathrm{Th}(\gamma_0), \mathrm{Th}(\gamma_1,), \mathrm{Th}(\gamma_2), \ldots$$

As discussed above, the classifying map of the bundle $\gamma_n \oplus \epsilon^1$ induces a map

$$\Sigma\mathrm{Th}(\gamma_n) \cong \mathrm{Th}(\gamma_n \oplus \epsilon^1) \xrightarrow{\sigma} \mathrm{Th}(\gamma_{n+1}).$$

The adjoint of this is an inclusion $\sigma' \colon \mathrm{Th}(\gamma_n) \hookrightarrow \Omega\mathrm{Th}(\gamma_{n+1})$. If these were equivalences, we would have an Ω-spectrum (and thus a represented cohomology theory). To turn this into an Ω-spectrum, we "force" these to become equivalences using the following definition. Define

$$MO_n \overset{\text{def}}{=} \operatorname{colim} \big(\mathrm{Th}(\gamma_n) \longrightarrow \Omega\mathrm{Th}(\gamma_{n+1}) \longrightarrow \Omega^2\mathrm{Th}(\gamma_{n+2}) \longrightarrow \cdots\big),$$

where the maps are given by the adjoints. We claim that this is an Ω-spectrum, and that $\pi_{n+k}\mathrm{Th}(\gamma_k) \cong \operatorname{colim}_k \pi_{n+k} MO_k$ for k large enough.

To check that it is an Ω-spectrum, observe that the structure map is

$$\begin{aligned} MO_n \cong \operatorname{colim} \Omega^k \mathrm{Th}(\gamma_{n+k}) &\xrightarrow{\sigma'} \operatorname{colim} \Omega^k \Omega \mathrm{Th}(\gamma_{n+k+1}) \\ &\xrightarrow{\epsilon} \Omega \operatorname{colim} \Omega^k \mathrm{Th}(\gamma_{n+k+1}). \end{aligned}$$

To prove that this is an equivalence, it suffices to check that σ' and ϵ are weak equivalences. The map σ' is an isomorphism, as it is simply shifting by 1 in the colimit. The map ϵ is more complicated, as colimits do not, in general, commute with Ω. However, in this case it works, as each map in the colimit is a closed embedding.

This method did not rely on any properties of $\mathrm{Th}(\gamma_k)$; it works in general with any sequences of spaces $X_0, X_1, \ldots$ which are equipped with maps $\Sigma X_i \longrightarrow X_{i+1}$ (with suitable modifications if the adjoints are not closed embeddings). Such a sequence is called a *spectrum*; it turns out that spectra are more natural, and appear more often, than Ω-spectra. For more on this, see the "Further Reading" section.

Remark 7.16. Since we now have a second cohomology theory, it is possible to ask: what useful characteristic classes exist with values in MO? The existence of characteristic classes in a general cohomology theory turns out to rely on one extra property: the fact that the cohomology of $\mathbf{R}P^\infty$ (or $\mathbf{C}P^\infty$) is a polynomial ring with one generator. Once this is satisfied, characteristic classes can be defined, and they have similar properties to the classes we have already defined. For an in-depth discussion of this, see [Swi17, Chapter 16].

What is the cohomology theory represented by MO? We can compute its value on a point by

$$MO^n(S^0) \cong [S^0, \operatorname{colim} \Omega^k \mathrm{Th}(\gamma_{n+k})] \cong \operatorname{colim}[S^0, \Omega^k \mathrm{Th}(\gamma_{n+k})] \cong \operatorname{colim}[S^k, \mathrm{Th}(\gamma_{n+k})].$$

(Recall the notation in Definition 4.13.) By Lemma 7.14, the right-hand side of this eventually stabilizes, so it is a well-defined group. In fact, it turns out to be extremely geometric:

Theorem 7.17 ([Tho54, Théorème IV.8])**.** *When $k > n + 2$,*

$$\mathfrak{N}_n \cong \pi_{n+k} \mathrm{Th}(\gamma_k).$$

The rest of this chapter is concerned with the proof of this theorem and its relationship to the Pontrjagyn–Thom theorem.

7.4 Cobordism groups vs L-groups; an interlude on transversality

In order to relate cobordism groups to homotopy groups we will need to have a bit more data than a cobordism class provides. In particular, we will need to consider manifolds embedded into other manifolds, and to analyze their intersections.

An integral part of the proof is the following process: given a map $f\colon X \longrightarrow M$ and a submanifold $N \subseteq M$, we would like $f^{-1}(N)$ to be a submanifold of X of codimension equal to the codimension of N in M. This does not always hold, even in very simple cases. If we set $X = \mathbf{R}^2$ and $M = \mathbf{R}^3$, with the map being the inclusion of the xy-plane, and we let N be the zero set of $x^2 + y^2 - z$. Then N has codimension 1, but $f^{-1}(N)$ (a single point) has codimension 2.

For completeness we state the definition of transversality; however, we will not be using the definition, but instead will be using the properties stated in Theorem 7.20.

Definition 7.18. Let $f\colon X \longrightarrow M$ be a smooth map from an n-manifold to a p-manifold. Let $N \subseteq M$ be a submanifold of codimension q. For any point $y \in N$, let T_yM be the tangent space to M at y, and let T_yN be the subspace of T_yM which is the tangent space to N at y. Let $x \in f^{-1}(y)$. We then have a map $df_x\colon T_xX \longrightarrow T_yM \twoheadrightarrow T_yM/T_yN$. We say that f is *transverse to N at y* if this induced map is an epimorphism.

In general, f is transverse to N if it is transverse to N at every point of N. (If $f^{-1}(y) = \emptyset$ then transversality holds automatically.)

Transversality is dense and open, in the sense that any given f can be perturbed (in an arbitrarily small fashion) to a map which is transverse to N, and given any map which is transverse to N, so are all maps within some open neighborhood. To state this property formally, we need the following definition:

Definition 7.19. A homotopy $X \times I \longrightarrow Y$ is an *isotopy* if for each $t \in [0, 1]$ the restriction $X \times \{t\} \longrightarrow Y$ is smooth.

We are now ready to state the properties of transversality that we will need in this chapter:

Theorem 7.20 ([Tho54, Chapter 1])**.** *Let M be a p-manifold and N a compact submanifold of codimension q. Let T be a tubular neighborhood of N in M. We assume that X is a smooth n-manifold with boundary.*

(1) Let $f\colon X \longrightarrow M$ be any smooth map transverse to N, which is also transverse to N when restricted to the boundary. Then $f^{-1}(N)$ is a smooth $n - q$-submanifold of X.

(2) Let $f\colon X \longrightarrow M$ be any smooth map. There exists a smooth homeomorphism A of T, arbitrary close to the identity and equal to the identity on ∂T, such that $A \circ f$ is transverse to N, and such that $(A \circ f)|_{\partial X}$ is also transverse to N. In particular, $(A \circ f)^{-1}(N)$ is a smoothly embedded $n - q$-submanifold of X.

(3) Suppose f is transverse to N. Let A be an automorphism of T, equal to the identity on ∂T. If A is sufficiently close to the identity, the map $A\circ f$ is transverse to N and the submanifolds $f^{-1}(N)$ and $(A\circ f)^{-1}(N)$ are isotopic in V.

In particular, part (3) implies that $f^{-1}(N)$ and $(A \circ f)^{-1}(N)$ are isomorphic.

It's important to note that transversality is a local condition. In particular, note that for any bundle $E \longrightarrow M$ on a manifold, the space $\mathrm{Th}(E)$ is a manifold away from the basepoint. Thus all consequences of transversality apply to $\mathrm{Th}(E)$ as well, as long as the image of f does not contain the basepoint.

We can now proceed to relate homotopy and cobordism. Just as with the case of classification of vector bundles earlier, we're going to show that we can classify cobordism classes using homotopy classes of maps.

Theorem 7.21 ([Tho54, Théorème IV.4])**.** *Let $f, g\colon X \longrightarrow M$ be two smooth maps, where $m \geq n$, and suppose that both are transverse to N. Let $W = f^{-1}(N)$ and $W' = g^{-1}(N)$. If f and g are homotopic then W and W' are cobordant.*

Proof. [Tho54, Lemma IV.5] states that if there exists a homotopy between smooth maps then there exists a smooth homotopy between them. Thus we can assume that f and g are smoothly homotopic; let $h\colon X \times I \longrightarrow M$ be such a smooth homotopy. By Theorem 7.20(2), there exists an automorphism A such that $A \circ h$ is transverse to N. Applying Theorem 7.20(3) to this A, $(A \circ h)|_{X\times\{0\}}^{-1}(N)$ is isotopic (and therefore homeomorphic) to $f^{-1}(N)$ and $(A \circ h)|_{X\times\{0\}}^{-1}(N)$ is isotopic (and therefore homeomorphic) to $g^{-1}(N)$.

In particular, this analysis shows that it suffices to prove the theorem in the case where h is assumed to be smooth and transverse to N. Then $h^{-1}(N)$ is an $n - q + 1$-manifold with boundary $f^{-1}(N) \amalg g^{-1}(N)$, giving the desired cobordism. □

This theorem produces more than just a cobordism between $f^{-1}(N)$ and $g^{-1}(N)$; it produces a cobordism *with an embedding in* X. This motivates the following definition:

Definition 7.22. Let W_0, W_1 be two n-submanifolds of an $n + k$-manifold X. Then they are *L-equivalent in X* if there exists a $n + 1$-manifold Y with boundary $W_0 \amalg W_1$ with an embedding $Y \longrightarrow X \times I$ with $Y \cap X \times \{0\} = W_0$ and $Y \cap X \times \{1\} = W_1$.

By gluing these embeddings together we see that L-equivalence is an equivalence relation.

We denote by $L_n(X)$ the set of L-equivalence classes of n-submanifolds in X.

For example, if W_0 and W_1 are the boundary of a submanifold Z in X then they are L-equivalent: we can use Urysohn's Lemma to construct a function $\varphi: Z \longrightarrow [0,1]$ which is 0 on W_0 and 1 on W_1 and then take the embedding $f: Z \longrightarrow X \times I$ given by $f(z) = (i(z), \varphi(z))$. However, L-equivalence is in general a stronger condition than cobordism. For example, if X is not connected and W_0, W_1 lie in different connected components, then they are not L-equivalent in X even if they are isomorphic as abtract manifolds. However, in the case where we restrict X to be a sphere, L-equivalence is equivalent to cobordism:

Lemma 7.23. *For $k > n+2$, $L_k(S^{n+k})$ is an abelian group. The function $\varphi: L_k(S^{n+k}) \longrightarrow \mathfrak{N}_n$ taking the L-equivalence class of a submanifold to its cobordism class is an isomorphism.*

Proof. When $k > n+2$ we can always isotope two n-submanifolds of S^{n+k} to be disjoint. Then disjoint union gives a well-defined operation on $L_n(S^{n+k})$. As with cobordisms, the inverse of a submanifold W is itself (via a "horseshoe" cobordism)

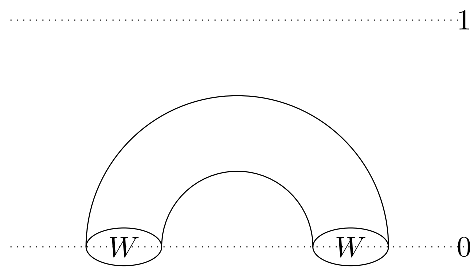

with the identity the empty submanifold. Thus $L_n(S^{n+k})$ is an abelian group. Since L-equivalence implies cobordism, φ is well-defined; with this operation it is clearly a homomorphism.

Since any manifold of dimension n can be embedded into S^{n+k} (as $n + k > 2n$) φ is surjective. Now suppose that a submanifold W of S^{n+k} is null-cobordant; thus there exists a manifold B of dimension $n+1$ with W as the boundary. Since $n+k > 2(n+1)$, there exists an embedding $f: B \longrightarrow S^{n+k}$. By Urysohn's lemma there exists a map $\psi: B \longrightarrow [0,1]$ such that $\psi^{-1}(0) = W$. Then the embedding $B \longrightarrow S^{n+k} \times [0,1]$ given by $x \longmapsto (f(x), \psi(x)/2)$ shows that W is L-equivalent to the empty set, as desired. Thus φ is injective, and thus an isomorphism. □

7.5 The Pontrjagin–Thom construction

Let $C = n+k+1$, let $N \overset{\text{def}}{=} G_k(\mathbf{R}^C)$, sitting as the 0-section inside $M = \mathrm{Th}(\gamma_{kC})$. Let $X = S^{n+k}$. Thus M is a manifold away from the basepoint, and X and N are both compact manifolds.[c] By Theorem 7.21 there is a well-defined function of *sets*

$$\begin{aligned}\pi_{n+k}(\mathrm{Th}(\gamma_{kC})) \cong [S^{n+k}, \mathrm{Th}(\gamma_{kC})] &\xrightarrow{G} L_n(S^{n+k}) \cong \mathfrak{N}_n \\ [f] &\longmapsto f^{-1}(G_k(\mathbf{R}^C)).\end{aligned}$$

This is actually a homomorphism, which follows from the observation that $G_k(\mathbf{R}^C)$ does not contain the basepoint in $\mathrm{Th}(\gamma_{kC})$. Indeed, given maps $f, g\colon S^{n+k} \longrightarrow \mathrm{Th}(\gamma_{kC})$, the class $[f]+[g]$ is given by a map $S^{n+k} \longrightarrow \mathrm{Th}(\gamma_{kC})$ which can be drawn as follows:

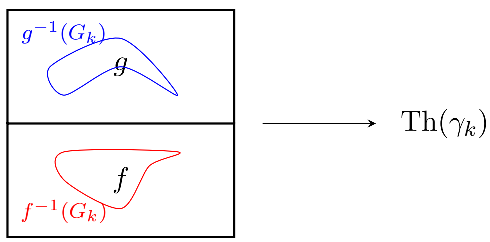

Here we are thinking of a map $S^{n+k} \longrightarrow \mathrm{Th}(\gamma_{kC})$ as a map $I^{n+k} \longrightarrow \mathrm{Th}(\gamma_{kC})$ which maps to the point at infinity on the boundary. In the diagram above, therefore, all solid lines in the domain map to the point at infinity; thus the red and blue lines do not intersect the black lines, and cannot intersect one another. It follows that $G([f]+[g]) = G[f] \amalg G[g]$, and G is a homomorphism.

To prove Theorem 7.17 it therefore remains to check that G is an isomorphism. To prove this we begin by constructing a right inverse J via the Pontrjagin–Thom construction.

Definition 7.24. We consider S^{n+k} to be sitting inside $\mathbf{R}^C$ in the standard way. Let $\alpha \in L_n(S^{n+k})$ be an L-equivalence class. Pick a representative embedding $W \hookrightarrow S^{n+k}$; let ν_W be the orthogonal complement of TW inside TS^{n+k}. This is a k-bundle, so it is classified by a map $f\colon W \longrightarrow G_k(\mathbf{R}^C)$. Pick $\epsilon > 0$, and let T_ϵ be the space of points within ϵ of W. If ϵ is small enough, every point in T_ϵ has a unique closest (via a geodesic in S^{n+k}) point in W, and the induced projection $T_\epsilon \longrightarrow W$ is a disk bundle. The boundary ∂T_ϵ is a sphere bundle. Moreover, [Tho54, Section I.3] T_ϵ inherits a canonical isomorphism with the unit disk bundle $D(\nu_W)$. Thus f induces

[c] This can also be done without a finite C by extending Theorem 7.20 to well-partitionable spaces.

a map $\widetilde{f}: T_\epsilon \longrightarrow \gamma_{kC}$ which takes ∂T_ϵ to the unit sphere bundle. We can then define a map $PT_W: S^{n+k} \longrightarrow \mathrm{Th}(\gamma_{kC})$ by

$$PT_W(x) = \begin{cases} * & \text{if } x \notin T \\ \widetilde{f}(x) & \text{if } x \in T. \end{cases}$$

This is smooth away from the preimage of the basepoint. Note, also, that the preimage of $G_k(\mathbf{R}^C) \subseteq \mathrm{Th}(\gamma_{kC})$ is exactly W.

This is the *Pontrjagin–Thom construction.*

We define a function

$$J: L_n(S^{n+k}) \longrightarrow \pi_{n+k}\mathrm{Th}(\gamma_{kC}) \qquad \text{by} \qquad [W] \longmapsto [PT_W].$$

Lemma 7.25. *The homotopy class of PT_W depends only on α, not the choice of W, f, or ϵ. Consequently, J is well-defined.*

Proof. First, notice that the homotopy class of PT_W does not depend on choice of f. Indeed, suppose an alternate representative f' is chosen. By the classification theorem of bundles, f is homotopic to f'. Since $\nu_W \cong f^*\gamma_{kC} \cong (f')^*\gamma_{kC}$ the induced maps $\widetilde{f}$ and $\widetilde{f}'$ must also be homotopic. Thus the induced maps $S^{n+k} \longrightarrow \mathrm{Th}(\gamma_{kC})$ must be homotopic, as desired.

Next, consider two different choices of ϵ. The difference between these involves "flowing" some of portion of S^{n+k} into the basepoint along geodesics orthogonal to W, which does not change the homotopy class of the map.

Lastly, consider an L-equivalence between two n-submanifolds in S^{n+k}. This is given by an embedding of $Y \hookrightarrow S^{n+k} \times I$, where Y is an $n+1$-manifold with boundary; the embedding takes the boundary of Y to $S^{n+k} \times \{0, 1\}$; write $W_i = Y \cap S^{n+k} \cap \{i\}$ for $i = 0, 1$. Points on the boundary within ϵ of Y are still within ϵ on the boundary; thus T_ϵ for Y restricts to a T_ϵ for each W_i. Then $PT_Y: S^{n+k} \times I \longrightarrow \mathrm{Th}(\gamma_{kC})$ is a homotopy between PT_{W_0} and PT_{W_1}. □

Lemma 7.26. *J is a group homomorphism.*

Proof. Let $\alpha, \alpha' \in L_n(S^{n+k})$ be two L-equivalence classes. We can choose representatives W, W' of these equivalence classes so that they lie inside disjoint hemispheres, and choose tubular neighborhoods T, T' of W and W' (respectively) so that they are disjoint, and each lie in their own respective hemisphere. Then $T \amalg T'$ is a tubular neighborhood of $W \amalg W'$. Using these choices, the diagram

$$\begin{array}{ccc} S^{n+k} & \xrightarrow{PT_{W \amalg W'}} & \mathrm{Th}(\gamma_{kC}) \\ & \searrow^{\text{collapse equator}} & \nearrow_{PT_W \vee PT_{W'}} \\ & S^{n+k} \vee S^{n+k} & \end{array}$$

commutes. The composition around the bottom represents the sum in $\pi_{n+k}(\mathrm{Th}(\gamma_{kC}))$ of $[PT_W]$ and $[PT_{W'}]$. The map across the top represents $[PT_{W \amalg W'}]$. Thus $J([W \amalg W']) = J([W]) + J([W'])$, as desired. □

We can increase C without affecting the definition of J, as S^{n+k} embeds into $\mathbf{R}^{C'}$ for any $C' \geq C$ without affecting any of the definitions in the Pontrjagyn–Thom construction. Increasing C does not affect the definition of G, either. In light of this, we drop C from our notation and disregard it in the future analysis.

Consider the function $G \circ J$. This takes an L-equivalence class $[W]$ to a map $f\colon S^{n+k} \longrightarrow \mathrm{Th}(\gamma_k)$ such that the preimage of $G_k \subseteq \mathrm{Th}(\gamma_k)$ is exactly W; thus $G \circ J$ is the identity. In particular, this implies that J is injective and G is surjective.

To prove that G is an isomorphism it suffices to prove that J is an isomorphism, or in other words that J is surjective. As every map is homotopic to one which is transverse to a fixed submanifold, the key lemma to prove is the following:

Lemma 7.27. *Let $f\colon S^{n+k} \longrightarrow \mathrm{Th}(\gamma_k)$ be a map which is transverse to G_k. Then f is homotopic to a map obtained from a classifying map of the normal bundle to $f^{-1}(G_k)$. Consequently, J is surjective.*

Proof. Let T_ϵ be the set of points within ϵ of $W \overset{\text{def}}{=} f^{-1}(G_k)$. If ϵ is sufficiently small, T_ϵ is isomorphic to the disk bundle in ν_W. Up to homotopy (and making ϵ smaller if necessary), we may assume that f takes the interior of T_ϵ to $\mathrm{Th}(\gamma_k) \smallsetminus \{*\}$, and that the preimage of $*$ is exactly $S^{n+k} \smallsetminus T_\epsilon$.

With this picture, there are two maps $T_\epsilon \longrightarrow W$: the map given by the canonical projection (which exhibits T_ϵ as the disk bundle of the normal bundle to W) and the composition

$$T_\epsilon \xrightarrow{f} \mathrm{Th}(\gamma_k) \smallsetminus * \longrightarrow G_k \xrightarrow{f^{-1}} W.$$

If we can show that these two maps are homotopic via homotopies that extend to one-point compactifications, we will be done.

□

We have thus shown that

$$L_n(S^{n+k}) \cong \pi_{n+k}(\mathrm{Th}(\gamma_k)).$$

By Lemma 7.23, $L_n(S^{n+k}) \cong \mathfrak{N}_n$. Putting these together we get

$$\mathfrak{N}_n \cong L_n(S^{n+k}) \cong \pi_{n+k}\mathrm{Th}(\gamma_k) \qquad \text{when } k > n,$$

which is eactly the statement of Theorem 7.17. □

To finish this section we use Theorem 7.17 to prove the "if" direction of the Pontrjagin–Thom Theorem (Theorem 7.11). We begin with a simple corollary which rephrases Theorem 7.17 in way most useful for the proof of Pontrjagin–Thom.

Corollary 7.28. *Suppose $k > n + 2$. The n-manifold M is a boundary if and only if the map $PT_M \colon S^{n+k} \longrightarrow \mathrm{Th}(\gamma_k)$ is null-homotopic.*

We will also need the following technical lemma, whose proof we omit as it relies on a close homotopical analysis of $\mathrm{Th}(\gamma_k)$; for the interested reader, one explanation can be found in [Tho54, discussion between theorems IV.9 and IV.10].

Lemma 7.29. *When $k > n+2$, a map $f \colon S^{n+k} \longrightarrow \mathrm{Th}(\gamma_k)$ is null-homotopic if and only if the induced homomorphism on cohomology $f^* \colon \widetilde{H}^{n+k}(\mathrm{Th}(\gamma_k); \mathbb{Z}/2) \longrightarrow \widetilde{H}^{n+k}(S^{n+k}; \mathbb{Z}$ is zero.*

We are now ready to finish the proof of the Pontrjagin–Thom theorem:

Proof of "if" direction of Theorem 7.11. Suppose that a manifold M has all Steifel–Whitney numbers 0. The Pontrjagin–Thom construction produces the following diagram:

$$\begin{array}{ccccc} N & \hookrightarrow & S^{n+k} & \xrightarrow{PT_M} & \mathrm{Th}(\gamma_k) \\ \uparrow & \nearrow_{\iota} & & & \uparrow \\ M & & \xrightarrow{\widetilde{f}} & & G_k \end{array}$$

where N is a tubular neighborhood of M and f is the classifying map of a normal bundle ν_M, modeled using the embedding $\iota; M \hookrightarrow S^{n+k}$.

For any tuple $(r_1, \ldots, r_n)$ with $\sum ir_i = n$ we can define the *normal Stiefel–Whitney numbers* to be the terms $w_1^{r_1}(\nu_M) \cdots w_n^{r_n}(\nu_M)[M] \in \mathbb{Z}/2$. By the Whitney Duality Theorem, the Stiefel–Whitney numbers of M are zero if and only if the normal Stiefel–Whitney numbers of M are 0.

Since $\widetilde{H}^{n+k}\mathrm{Th}(\gamma_k) \cong H^n(G_k)$ (by the Thom isomorphism theorem),

$$\widetilde{H}^{n+k}(\mathrm{Th}(\gamma_k); \mathbb{Z}/2) \cong \mathbb{Z}/2 \left\{ (r_1, \ldots, r_n) \in \mathbb{Z}^n_{\geq 0} \;\middle|\; \sum i r_i = n \right\}.$$

By definition, $PT^*_M(r_1, \ldots, r_n)$ is the $(r_1, \ldots, r_n)$-th normal Steifel–Whitney number. By Lemma 7.29 these are all 0 exactly when PT_M is null-homotopic, which by Corollary 7.28 implies that M is a boundary, as desired. □

Further reading

Thom's paper [Tho54] contains a *lot* of very cool material. Readers should be aware that, given that this was the original paper for much of this material, the notation is very nonstandard. The original paper is in French, but is well-worth the read. There is an English translation in [NT07]. Smoothness considerations are discussed carefully there, especially in [Tho54, Chapter I]. The sections relevant to the discussion in this chapter are I and IV.

When we proved that cobordism groups are isomorphic to homotopy groups, we went through L-equivalence classes inside a sphere. In fact, the proof works in more generality: for any $n+k$-manifold X, the group $L_n(X)$ is isomorphic to $[X, \mathrm{Th}(\gamma_k)]$ (which will also naturally be a group). The proof works analogously to the proof we gave, although requires more machinery; for a complete exploration of this topic, see [Tho54, Chapter IV].

For more computations of Steifel–Whitney numbers, see [MS74, Section 4]. For a discussion of Steifel—Whitney classes inside other cohomology theories, including MO, see [Swi17, Chapter 16].

There are many good resources on spectra. A place to start is [Sch], whose first couple of sections give a good overview of symmetric spectra.

There are many similar examples of groups of equivalence classes manifolds equipped with structure relative to cobordisms being isomorphic to certain homotopy groups. A basic example is *oriented cobordism*, also discussed in [Tho54]. Other examples abound, however, with one of the more interesting ones being *framed cobordism*: cobordisms with a choice of embedding and trivialization of the *normal bundle*. The associated homotopy groups turn out to be the stable homotopy groups of spheres. To read about this, see for example [Ran02, Chapters 2,6].

Exercises and Extensions

7.1 Compute the Stiefel–Whitney numbers of the Klein bottle and use them to prove that the Klein bottle is not the boundary of a 3-

manifold.

7.2 Prove directly that $\mathbf{R}P^{2k-1}$ is a boundary of a $2k$-manifold. (Hint: $\mathbf{R}P^{2k-1}$ double-covers a lens space, and is thus the boundary of an I-bundle on the lens space.)

7.3 Prove that the disjoint union and cartesian product make $\mathfrak{N}_*$ into a well-defined algebra over $\mathbb{Z}/2$.

7.4 Complete the proof of Lemma 7.27.

7.5 Suppose that

$$X_0 \hookrightarrow X_1 \hookrightarrow X_2 \hookrightarrow \cdots$$

is a sequence of closed embeddings of manifolds. Prove that the induced map

$$\operatorname{colim} \Omega X_n \longrightarrow \Omega \operatorname{colim} X_n$$

is a homeomorphism.

7.6 (Pontrjagin) A *framing* of a manifold M is an embedding $M \hookrightarrow \mathbf{R}^N$ together with a choice of trivialization of the normal bundle. A framing of a cobordism W between M and N is a framing on W which restricts to a framing of M and N (using a collar neighborhood of the boundary of W for the extra dimension). Let Ω_n^{fr} be the group of framed cobordism classes of n-manifolds. Use the technique in this chapter to prove that $\Omega_n^{fr} \cong \operatorname{colim}_k \pi_{n+k} S^k$.

Chapter 8

Bott Periodicity

With the addition of MO we have now seen two examples of cohomology theories. In the next chapter we will construct a third: topological K-theory. Topological K-theory turns out to be significantly simpler than either of the two previously-mentioned cohomology theories. However, in the spirit of this book, we will be introducing it homotopically, instead of algebraically. Therefore before we can classify the cohomology theory we must construct its representing spectrum.

The spectrum representing this cohomology theory, usually called KO or KU, is constructed out of unitary groups and their classifying spaces. Because of this algebraic underpinning it turns out to be surprisingly simple to work with—so simple, in fact, that Atiyah wrote a book[a] [Ati89] on topological K-theory which does not mention singular cohomology at all. Part of the goal of this book, in fact, was to try and convince people that topological K-theory is a better introduction to cohomology theory than the classical example. For those readers interested in computations, or in a purely algebraic introduction to the subject, the book is highly recommended.

Remark 8.1. In this chapter we shift our attention from real vector bundles to complex vector bundles. As in previous chapters, we will state the theorems for both cases but only prove it for one. Previously we focused on the real case, as its proof displayed some interesting internal structure of Grassmannians. In this chapter we take the opposite tack: we focus on the complex case, as the substancially simpler one, in order to display the

[a] Atiyah also wrote a paper [Ati66] on Real K-theory, which is (unfortunately for current and future generations) not the same thing as the real topological K-theory discussed in this book. The paper is also very interesting, but is aimed at the seasoned topologist.

machinery of the proof more clearly.

In this chapter, Section 8.1 introduces Bott periodicity and gives the statement of the main result. The rest of the chapter is focused on the proof. Section 8.2 explains how to construct the Bott map. Section 8.3 gives some necessary background on H-spaces, and Section 8.4 constructs the cell structure on SU. Section 8.5 completes the proof. A brief discussion of Bott periodicity in the real case is given in Section 8.6.

8.1 The Statement of Bott Periodicity

We have spent some time studying Grassmannians. We know their cohomology (and thus their homology), and we have mentioned that, in general, their homotopy groups are not known. However, it turns out that these spaces satisfy a nice stability condition which allows us to calculate some of their homotopy groups. Recall the definition of the classifying space of a group given in Definition 2.31. We proved in Theorem 2.32 (and Exercise 2.1) that $GU_n \simeq BU(n)$; thus to study complex Grassmannians it suffices to fully understand $U(n)$.

The group $U(n)$ acts on $S^{2n-1} \subseteq \mathbf{C}^n$, with the stabilizer of a point isomorphic to $U(n-1)$. Since this action is continuous, the orbit-stabilizer theorem applies and we get that $S^{2n-1} \cong U(n)/U(n-1)$. More importantly, this produces a fiber sequence

$$U(n-1) \longrightarrow U(n) \longrightarrow S^{2n-1}.$$

The long exact sequence of this fibration is the following:

$$\cdots \longrightarrow \pi_{i+1}S^{2n-1} \longrightarrow \pi_i U(n-1) \longrightarrow \pi_i U(n) \longrightarrow \pi_i S^{2n-1} \longrightarrow \cdots$$

Since $\pi_i S^{2n-1} = 0$ for $i < 2n-1$, the sequence gives an isomorphism $\pi_i U(n-1) \cong \pi_i U(n)$ for $i < 2n-2$. Thus the low-dimensional homotopy groups do not depend on the dimension of the ambient space. It is reasonable, therefore, to guess that calculating the homotopy groups of $U = \operatorname{colim} U(n)$ might be simpler than calculating the homotopy groups of $U(n)$.

Theorem 8.2 (Bott Periodicity)**.** *For $k \geq 0$,*

$$\pi_k U \cong \pi_{k+2} U.$$

The statement of real Bott periodicity is given in Section 8.6.

We will be following the proof from [DL61]. There are many different proofs of this theorem, from Bott's original proof using Morse Theory to a spectral sequence argument of Moore's, to new proofs using quasifibrations of Behrens and Aguilar–Prieto; a list of references is given in the "Further Reading" section. The approach that we follow has the advantage that it does not require a lot of theory, relying mostly on an understanding of algebra and some topological techniques.

The idea of the proof is to construct the *Bott map*

$$\Phi\colon BU \longrightarrow \Omega SU,$$

where $SU = \operatorname{colim} SU(n)$.

Theorem 8.3 (Existence of Bott map)**.** *The Bott map exists and is a weak equivalence.*

Assuming Theorem 8.3 we can prove Bott periodicity.

Proof of Bott periodicity. Write $\Omega_0 X$ for the connected component of ΩX containing the constant loop. Then $\Omega_0 U \simeq \Omega SU$. Thus there is a weak equivalence

$$U \simeq \Omega BU \xrightarrow{\Omega\Phi} \Omega^2 SU \simeq \Omega\Omega_0 U \simeq \Omega^2 U,$$

The first step is true for any group G, and is left as an exercise for the reader. The last step follows by definition, since for any space X, $\Omega X = \Omega X_x$, where X_x is the connected component of X containing the basepoint. Since $\pi_k \Omega X \cong \pi_{k+1} X$, the theorem follows. □

We can use this to compute the homotopy groups of U. By Theorem 8.2 it suffices to compute π_0 and π_1. Analogously to the observation about $O(n)$ at the beginning of this section, $\pi_i U(n-1) \cong \pi_i U(n)$ for $i < 2n-2$. Thus to find all of the homotopy groups up to π_1 it suffices to consider $n = 2$. But $U(2) \cong S^1 \times SU(2) \cong S^1 \times S^3$, which has $\pi_0 = 0$ and $\pi_1 \cong \mathbb{Z}$. Thus

$$\pi_{\text{even}} U \cong 0 \qquad \text{and} \qquad \pi_{\text{odd}} U \cong \mathbb{Z}.$$

We now turn our attention to proving Theorem 8.3.

8.2 Constructing the Bott map Φ

Definition 8.4. For $k < k'$, write $\iota_{kk'}\colon \mathbf{C}^k \longrightarrow \mathbf{C}^{k'}$ for the inclusion of $\mathbf{C}^k$ into the first k coordinates of $\mathbf{C}^{k'}$. Write

$$\iota_{kk'|nn'} \overset{\text{def}}{=} \iota_{kk'} \times \iota_{nn'}\colon \mathbf{C}^k \times \mathbf{C}^n \longrightarrow \mathbf{C}^{k'} \times \mathbf{C}^{n'}.$$

For fixed k and n and $\theta \in [0, 2\pi]$ we define a continuous family of linear maps, $\alpha_\theta^{k,n}$, by

$$\begin{aligned} \alpha_{k,n}^\theta \colon \mathbf{C}^k \times \mathbf{C}^n &\longrightarrow \mathbf{C}^k \times \mathbf{C}^n \\ (z_1, z_2) &\longmapsto (z_1 e^{i\theta}, z_2 e^{-i\theta}) \end{aligned}$$

For each θ, $\alpha_{k,n}^\theta \in U(k+n)$; this data can therefore be used to define a map $\alpha_{k,n} \colon S^1 \longrightarrow U(k+n)$, or in other words a point in $\Omega U(k+n)$. Moreover, this is *natural* in k and n, in the following sense: the diagram

$$\begin{array}{ccc} \mathbf{C}^k \times \mathbf{C}^n & \xrightarrow{\alpha_{k.n}^\theta} & \mathbf{C}^k \times \mathbf{C}^n \\ {\scriptstyle \iota_{kk'|nn'}} \downarrow & & \downarrow {\scriptstyle \iota_{kk'|nn'}} \\ \mathbf{C}^{k'} \times \mathbf{C}^{n'} & \xrightarrow{\alpha_{k',n'}^\theta} & \mathbf{C}^{k'} \times \mathbf{C}^{n'} \end{array}$$

commutes.

Consider the map $\widetilde{\Phi}_{k,n} \colon U(k+n) \longrightarrow \Omega SU(k+n)$, defined by

$$T \longmapsto \left(\theta \longmapsto T \circ \alpha_{k,n}^\theta \circ T^{-1} \circ (\alpha_{k,n}^\theta)^{-1} \right).$$

Suppose that $T = T_k \times T_n \in U(k) \times U(n)$. Then

$$\Phi_{k,n}(T_k \times T_n) = \begin{pmatrix} T_k e^{i\theta} T_k^{-1} e^{-i\theta} & \\ & T_n e^{-i\theta} T_n^{i-1} e^{i\theta} \end{pmatrix} = I.$$

Thus $\widetilde{\Phi}$ takes any $T \in U(k) \times U(n)$ to the trivial loop and induces a map

$$\Phi_{k,n} \colon U(k+n)/U(k) \times U(n) \longrightarrow \Omega SU(k+n).$$

Warning: The quotient above is as a space with a group action, *not* as groups, since $U(k) \times U(n)$ is not a normal subgroup of $U(k+n)$.

The map $\iota_{kk'|nn'}$ induces a map $U(k+n) \hookrightarrow U(k'+n')$ which in turn induces maps

$$\iota_{kk'|nn'} \colon U(k+n)/U(k) \times U(n) \longrightarrow U(k'+n')/U(k') \times U(n').$$

These fit into a commutative square of the form

$$\begin{array}{ccc} U(k+n)/U(k) \times U(n) & \xrightarrow{\Phi_{k,n}} & \Omega SU(k+n) \\ {\scriptstyle \iota_{kk'|nn'}} \downarrow & & \downarrow {\scriptstyle \Omega(\iota_{kk'|nn'})} \\ U(k'+n')/U(k') \times U(n') & \xrightarrow{\Phi_{k',n'}} & \Omega SU(k'+n'). \end{array} \tag{8.5}$$

Lemma 8.6.

$$\operatorname*{colim}_{k\to\infty} U(k+n)/U(k)\times U(n) \cong Gr_n(\mathbf{C}^\infty).$$

Proof. Let $V_n(\mathbf{C}^{k+n})$ be the Stiefel manifold of n-frames in $\mathbf{C}^{k+n}$. Then, analogously to Exercise 2.4,

$$V_n(\mathbf{C}^{k+n}) \cong U(k+n)/U(k).$$

Thus

$$\begin{aligned}\operatorname*{colim}_{k\to\infty} U(k+n)/U(k)\times U(n) &\cong \operatorname*{colim}_{k\to\infty} V_n(\mathbf{C}^{k+n})/U(n)\\ &\cong \operatorname*{colim}_{k\to\infty} Gr_n(\mathbf{C}^{k+n}) = Gr_n(\mathbf{C}^\infty).\end{aligned}$$

□

Taking vertical colimits as k and n go to infinity in (8.5) induces a map

$$\Phi\colon \operatorname*{colim}_{k,n\to\infty} U(k+n)/U(k)\times U(n) \longrightarrow \operatorname*{colim}_{k,n\to\infty} \Omega SU(k+n).$$

Since

$$\operatorname*{colim}_{n\to\infty}\operatorname*{colim}_{k\to\infty} U(k+n)/U(k)\times U(n) \cong \operatorname*{colim}_{n\to\infty} Gr_n(\mathbf{C}^\infty) \cong \operatorname*{colim}_{n\to\infty} BU(n) \cong BU,$$

Φ is a map

$$BU \longrightarrow \Omega SU,$$

as desired. This is the Bott map.

8.3 H-spaces

Recall[b] the definition of an H-space.

Definition 8.7. An *H-space* is a space X equipped with a *multiplication map* $\mu\colon X\times X \longrightarrow X$ and an *identity* $e\in X$ such that $\mu(\cdot,e)$ and $\mu(e,\cdot)$ are homotopic to the identity map on X. An H-space is called *homotopy associative* if the multiplication is associative up to homotopy (i.e. if there is a homotopy between the maps $\mu(\cdot,\mu(\cdot,\cdot))$ and $\mu(\mu(\cdot,\cdot),\cdot)$).

An *H-map* $f\colon X\longrightarrow Y$ between H-spaces X and Y is a map such that the square

[b] As this book only assumes one semester of alagebraic topology, the word "recall" may be somewhat disingenuous here.

$$
\begin{array}{ccc}
X \times X & \xrightarrow{\mu} & X \\
{\scriptstyle f \times f}\downarrow & & \downarrow{\scriptstyle f} \\
Y \times Y & \xrightarrow{\mu} & Y
\end{array}
$$

commutes up to homotopy.

Thus an H-space is a space with unital multiplicaiton "up to homotopy"; note, however, that this operation is not required to be associative or to have an inverse.

Remark 8.8. There are three definitions of an H-space in the literature: the definition above, a definition requiring that $\mu(\cdot, e)$ and $\mu(e, \cdot)$ are *equal* to the identity (rather than homotopic), and one that allows homotopies but requires them to fix e. Luckily, these three definitions are equivalent for CW-complexes.

Example 8.9. Any topological group is a homotopy associative H-space.

Example 8.10. Suppose that we are given a skew field structure on $\mathbf{R}^n$: a unital bilinear multiplication $p\colon \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}^n$ with no zero divisors. This gives a homotopy associative H-space structure on S^{n-1} in the following manner. Define $H\colon S^{n-1} \times S^{n-1} \longrightarrow S^{n-1}$ via

$$H(x, y) = \frac{p(x, y)}{|p(x, y)|}.$$

This is well-defined because $p(x, y)$ is never equal to 0 for $x, y \in S^{n-1}$, since there are no zero divisors. Let e be the unit of p. Then H will be unital because

$$H\left(\frac{e}{|e|}, y\right) = \frac{p\left(\frac{e}{|e|}, y\right)}{\left|p\left(\frac{e}{|e|}, y\right)\right|} = \frac{\frac{1}{|e|}p(e, y)}{\left|\frac{1}{|e|}p(e, y)\right|} = y.$$

and analogously for $H(x, e)$. The proof for associativity is analogous.

Example 8.11. For any pointed space Y, ΩY is a homotopy associative H-space, with operation given by concatenation of loops. If a strictly unital model is desired, one can instead take

$$\Omega^t Y \stackrel{\text{def}}{=} \{f\colon [0, t] \longrightarrow Y \mid t \in \mathbf{R}_{\geq 0},\ f(0) = f(t) = *\}.$$

The identity element in this case is the unique pointed map $[0, 0] \longrightarrow Y$.

Example 8.12. When G is an abelian discrete group, BG is a homotopy associative H-space. To prove this it is important to note that B is a functor $\mathbf{Gp} \longrightarrow \mathbf{Top}$ which commutes with products. When G is abelian the multiplication $\mu: G \times G \longrightarrow G$ is a group homomorphism, and we can define the multiplication to be

$$BG \times BG \cong B(G \times G) \xrightarrow{B\mu} BG.$$

The fact that μ is unital follows from the fact that G has an identity element.

Example 8.13. The space BU is a homotopy associative H-space via the *block diagonal* map. Consider the map $\mathbf{C}^\infty \oplus \mathbf{C}^\infty \longrightarrow \mathbf{C}^\infty$ embedding the coordinates of the first copy of $\mathbf{C}^\infty$ as the odd coordinates, and the second copy as the even coordinates. This gives a group homomorphism $U \times U \longrightarrow U$ which, by the reasoning of Example 8.12, produces a map of spaces $BU \times BU \longrightarrow BU$, as desired.

Example 8.14. The spheres S^0, S^1, S^3 and S^7 are the only spheres which are H-spaces, with the multiplication given by real multiplication, complex multiplication, quaternion multiplication and octonion multiplication, respectively. (This will be proved later; it is another consequence of Theorem 10.5.) Since octonion multiplication is not associative, S^7 is not an associative H-space. (In fact, no homotopy-associative H-space structure on S^7 exists; see [Jam57, Theorem 1.4].)

The extra structure of the H-space multiplication also endows homology with a product structure.

Definition 8.15. Let X be an H-space. By the Eilenberg–Zilber theorem there is a quasi-isomorphism $C_*(X) \otimes C_*(X) \simeq C_*(X \times X)$. Using this quasi-isomorphism we define the bilinear *cross-product*

$$\times: H_i(X) \oplus H_j(X) \longrightarrow H_{i+j}(X \times X)$$

by sending $[a] \oplus [b]$ to $[a \otimes b]$. When X is an H-space, we can therefore define a multiplication

$$H_*(X) \otimes H_*(X) \xrightarrow{\times} H_*(X \times X) \xrightarrow{\mu_*} H_*(X).$$

When the H-space structure on X is homotopy associative, this multiplication will also be associative. This is called the *Pontrjagin ring* of X.

In particular, both BU and ΩSU are H-spaces.

Lemma 8.16. *The map Φ is an H-map of H-spaces.*

Proof. The following diagram commutes for all k, k', n, n':

$$\begin{array}{ccc} U(k+n) \times U(k'+n') & \xrightarrow{\widetilde{\Phi}_{k,n} \times \widetilde{\Phi}_{k',n'}} & \Omega SU(k+n) \times \Omega SU(k'+n') \\ \Big\downarrow \text{diag} & & \Big\downarrow \Omega\text{diag} \\ U(k+k'+n+n') & \xrightarrow{\widetilde{\Phi}_{k+k',n+n'}} & \Omega SU(k+k'+n+n') \end{array}$$

The map Ωdiag is homotopic to the loop concatenation map. The analogous diagram to the above with $\Phi_{k,n}$ instead of $\widetilde{\Phi}_{k,n}$ also commutes; if we take the colimit as k, n go to infinity we get that Φ is an H-map of H-spaces, where the H-space structure on ΩU is the loop concatentation map, and the H-space structure on BU is the block diagonal map. □

The following theorem is a refinement of Whitehead's theorem:

Theorem 8.17 ([DL61, Theorem 1.6]). *Let $f\colon X \longrightarrow Y$ be a map of connected spaces. If f is an H-map of H-spaces and $f_*\colon H_i(X;\mathbb{Z}) \longrightarrow H_i(Y,\mathbb{Z})$ is an isomorphism for all i then f is a weak equivalence.*

Thus in order to prove that Φ is a weak equivalence it suffices to check that it is an isomorphism on homology.

8.4 Cell structure on SU

Recall the definition of an exterior algebra:

Definition 8.18. The exterior algebra generated by x over a ring R is defined to be

$$\Lambda_R[x] \overset{\text{def}}{=} R[x]/x^2.$$

For a set of generators S,

$$\Lambda_R[x \mid x \in S] \overset{\text{def}}{=} \bigotimes_{x \in S} \Lambda_R[x],$$

where the tensor product is taken over R.

There is an explicit description of the Pontrjagin ring structures of U and SU:

Theorem 8.19 ([Yok57, Theorem 8.1(7)])**.** *The Pontrjagin ring structures of $U(n)$ and $SU(n)$ are given by*

$$H_*(U(n)) \cong \Lambda_{\mathbb{Z}}[e_1, e_3, \dots, e_{2n-1}]$$
$$H_*(SU(n)) \cong \Lambda_{\mathbb{Z}}[e_3, \dots, e_{2n-1}],$$

where $|e_i| = i$. The inclusion map $U(n) \longrightarrow U(n+1)$ takes e_i to e_i. Thus the Pontrjagin rings of U and SU are given by

$$H_*U \cong \Lambda_{\mathbb{Z}}[e_{2i-1} \,|\, i \geq 1] \qquad \textit{and} \qquad H_*SU \cong \Lambda_{\mathbb{Z}}[e_{2i-1} \,|\, i \geq 2].$$

For a discussion and calculation of this structure (as well as analogous structures for other classical Lie groups), see [Yok57, Theorem 8.1, p.111]. We will not provide all of the details of the proof of the theorem, but we will show how the cell structure is constructed, as it will be useful for the rest of the section. For a more in-depth discussion of the cell structure, see [Yok56, Section 7].

Proof sketch. We will give a reduction of the cell structure for $U(n)$ to one of $SU(n)$. We will then describe the cell structure on $SU(n)$. From this description it will follow that this structure is compatible with the inclusion $SU(n) \hookrightarrow SU(n+1)$—and thus the inclusion into SU.

We begin by reducing the case of $U(n)$ to $SU(n)$. Consider the map

$$\begin{aligned} U &\longrightarrow S^1 \times SU(n) \\ A &\longmapsto (\det A, \operatorname{diag}((\det A)^{-1}, 1, \dots, 1) \cdot A). \end{aligned}$$

This is a homeomorphism, and by the Kunneth Formula[c]

$$H_*U(n) \cong H_*S^1 \otimes H_*SU(n) \cong \Lambda_{\mathbb{Z}}[e_1] \otimes H_*SU(n).$$

Thus in particular, to construct the cell structure on $U(n)$ it suffices to construct the cell structure on $SU(n)$.

The definition of cells in $SU(n)$ will be somewhat strange, as instead of using spheres to define the maps we will use $\Sigma\mathbf{C}P^{n-1}$ instead. The key point here is that $\Sigma\mathbf{C}P^{n-1} \backslash \Sigma\mathbf{C}P^{n-2}$ is homeomorphic to a disk, and therefore such a map can be considered the characteristic map of a cell.

[c] It is important to be careful here: the product structure here *is* the product of Pontrjagin rings, but this does *not* follow directly from the above homeomorphism; it requires some more work to prove this claim.

The characteristic maps are defined as follows. For a point $x \in \mathbf{C}P^{n-1}$, pick a representative $(x_1, \ldots, x_n)$ with $\sum_{i=1}^{n} |x_i|^2 = 1$. Define the matrix

$$A_x \stackrel{\text{def}}{=} (x_i \overline{x}_j)_{i,j=1}^{n} ;$$

note that this is independent of the choice of representative because any two representatives differ by multiplication by $\lambda \in \mathbf{C}$ with $\lambda\overline{\lambda} = 1$. This matrix also has the properties that

$$A_x = A_x^* \qquad \text{and} \qquad A_x A_x^* = A_x,$$

where $\cdot^*$ is the conjugate transpose of A_x. Consider points of $\Sigma\mathbf{C}P^{n-1}$ to be pairs of points (θ, x) with $\theta \in [-\pi/2, \pi/2]$ and $x \in \mathbf{C}P^{n-1}$. We define $f_n \colon \Sigma\mathbf{C}P^{n-1} \longrightarrow SU(n)$ by

$$f_n(\theta, x) = \Big(I_n - 2e^{-i\theta}\cos\theta A_x\Big)\begin{pmatrix} -e^{-2i\theta} & \\ & I_{n-1} \end{pmatrix}.$$

Each of the two matrices in the product is unitary, and their product has determinant 1. When $\theta = -\pi/2$ or $\pi/2$ this is equal to I_nThis map is also injective on the interior of the top cell in $\Sigma\mathbf{C}P^{n-1}$; this is the *characteristic map of the cell* e_{2n+1}. This map is stable, in the sense that the following diagram commutes for $k < n$:

$$\begin{array}{ccc} \Sigma\mathbf{C}P^{k-1} & \xrightarrow{f_k} & SU(k) \\ \downarrow & & \downarrow \\ \Sigma\mathbf{C}P^{n-1} & \xrightarrow{f_n} & SU(n), \end{array}$$

where the vertical arrows are the natural inclusions as the "first k coordinates." Directly from the definition we can show that f_k maps $\Sigma\mathbf{C}P^{k-1} \backslash \Sigma\mathbf{C}P^{k-2}$ homeomorphically into $SU(k) \subseteq SU(n)$.

For $n \geq k_1 > k_2 > \cdots > k_j \geq 2$ we define a map

$$f_{k_1,\ldots,k_j} \colon \Sigma\mathbf{C}P^{k_1-1} \times \cdots \times \Sigma\mathbf{C}P^{k_2-1} \longrightarrow SU(n)$$

by defining $f_{k_1,\ldots,k_j}(z_1, \ldots, z_j) = f(z_1)\cdots f(z_j)$. This is the *characteristic map of the cell* $e_{2k_1+1}e_{2k_2+1}\cdots e_{2k_j+1}$.

These are all of the cells of the CW structure on $SU(n)$. The multiplication structure in the Pontrjagin ring is via "multiplication on cell names", as the H-space structure on $U(n)$ is exactly the group structure. □

See also Exercise 5.5 for another construction of this form.

8.5 Proof that Φ is a weak equivalence

By Theorem 8.17, in order to prove that Φ is a weak equivalence it suffices to check that it induces an isomorphism on all homology groups. In order to prove this, we require some extra structure:

Definition 8.20. Consider (8.5) again. Setting $n = 1$ and $k' = k$ produces the square

$$\begin{array}{ccc} \mathbf{C}P^k & \xrightarrow{\Phi_{k,1}} & \Omega SU(k+1) \\ \Big\downarrow{\scriptstyle \iota_{kk|1n'}} & & \Big\downarrow{\scriptstyle \Omega(\iota_{kk|1n'})} \\ Gr_{n'}(\mathbf{R}^k) & \xrightarrow{\Phi_{0,n'}} & \Omega SU(k+n'). \end{array}$$

Taking $k, n' \longrightarrow \infty$ we get a square

$$\begin{array}{ccc} \mathbf{C}P^\infty & \xrightarrow{\Phi'} & \Omega SU \\ \Big\downarrow{\scriptstyle J} & & \Big\downarrow{\scriptstyle \Omega J'} \\ BU & \xrightarrow{\Phi} & \Omega SU \end{array} \tag{8.21}$$

defining the maps Φ', J and J'.

Tracing through the definitions gives following two observations:

Lemma 8.22. *The map $\Phi_{k,1}: \mathbf{C}P^k \longrightarrow \Omega SU(k+1)$ is the adjoint to the map $f_{k+1}: \Sigma\mathbf{C}P^k \longrightarrow \Omega SU(k+1)$. Thus the map $\Sigma\mathbf{C}P^\infty \longrightarrow SU(k+1)$ adjoint to Φ', takes the generator in degree $2k+1$ to f_k.*

This lemma is left as an exercise for the reader.

Lemma 8.23. *The map J' is a homology isomorphism.*

Proof. From the cell structure we constructed on SU, $\iota_{kk|1n'}: SU(k+1) \longrightarrow SU(k+n')$ is a homology isomorphism up to degree $2k$. Thus as $k \longrightarrow \infty$ this becomes a homology isomorphism. □

Since Φ is an H-map of H-spaces, by Theorem 8.17 it suffices to check that it takes the generators of the Pontrjagin ring structure on H_*BU to the generators of the Pontrjagin ring structure on $H_*\Omega SU$, and that these structures are isomorphic. To begin, it turns out that we can deduce the Pontrjagin ring structure on $H_*\Omega SU$ from what we know of the Pontrjagin ring structure on H_*SU:

Theorem 8.24 ([DL61, Theorem 2.7]). *Let X be an H-space such that $H_*(X)$ is a transgressively generated exterior algebra on odd generators of degree at least 3. Then $H_*(\Omega X)$ is a polynomial algebra generated by the adjoints of the generators.*

We will not prove this theorem; it is a relatively straightforward use of spectral sequences; the interested reader is encouraged to try it on their own, with reference to the "Further Reading" section at the end of this chapter.

Theorem 8.24 applies to SU. Thus we know that $H_*(\Omega SU)$ is a polynomial algebra generated by the adjoint maps to the generators of $H_*(SU)$—which, as mentioned before, are exactly the maps $\Phi_{k,1}$. Thus $H_*(\Omega SU)$ is a polynomial algebra on generators of even degrees. From this theorem and the previous two lemmas it also follows that ΩJ is a homology isomorphism.

To complete the proof it suffices to prove the following:

Proposition 8.25. *The Pontrjagin ring structure on H_*BU is a polynomial algebra generated by the images of the generators of $H_*\mathbf{C}P^\infty$.*

Proof. First, consider the H-space structure on BU. By analogy to the result about $H^*(BO(k))$,

$$H^*BU \cong \mathbb{Z}\left[c_{2i} \,\middle|\, |c_{2i}| = 2i\right].$$

The H-space multiplication map is given by the map which takes two matrices to their block diagonal sum. This map is exactly the map characterizing the Whitney sum of two bundles; thus on cohomology it takes the generator c_{2i} to $\sum_{j+k=i} c_{2j}c_{2k}$. The Pontrjagin product is (directly from the definition) the Poincare dual of this comultiplication, and the algebra structure on $H_*(BU)$ is given by the dual of this, we can conclude that $H_*(BU)$ is also a polynomial algebra on even generators z_{2i}.

It remains to check that z_{2i} is the image of the generator $b_{2i} \in H_{2i}(\mathbf{C}P^\infty)$ under J_*. By definition, J is the inclusion $\mathbf{C}P^\infty \cong BU(1) \longrightarrow BU$. By the definition of (complex) characteristic classes, $J^*(c_2)$ is a generator of $H^2(\mathbf{C}P^\infty)$, which is dual to b_2; on the other hand, $J^*(c_{2i}) = 0$ for $i > 1$. Thus $J_*(b_{2i}) = z_{2i}$ (since otherwise pushing forward, dualizing and then pulling back would give the wrong result), and the proof is complete. □

8.6 Real Bott Periodicity

The statement of real Bott periodicity is similar to the statement of complex Bott periodicity:

Theorem 8.26 (Bott Periodicity)**.** *Let* $O = \operatorname{colim} O(n)$*. Then*

$$\pi_k O \cong \pi_{k+8} O.$$

The proof of this is similar to that in the complex case, although it requires many more maps:

$$\begin{array}{ll} \Phi_1 \colon BSp \longrightarrow \Omega(U/Sp) & \Phi_2 \colon BO \longrightarrow \Omega(U/O) \\ \Phi_3 \colon U/Sp \longrightarrow \Omega(SO/U) & \Phi_4 \colon U/O \longrightarrow \Omega(Sp/U) \\ \Phi_5 \colon SO/U \longrightarrow \Omega SO & \Phi_6 \colon Sp/U \longrightarrow \Omega Sp. \end{array}$$

Here, Sp is the infinite symplectic group, $\operatorname{colim} Sp(n)$.

The proof is similar, although more complicated than the complex case. The homology computations become more involved, and the proof must be done in two stages: first showing that the maps analogous to Φ and J are isomorphisms in mod p homology for all primes p, and then lifting this to imply that they are isomorphisms in homology with $\mathbb{Z}$ coefficients. The case when $p = 2$ must be handled separately. For the interested reader these are explained in detail in [DL61]. The chain of necessary equivalences is the following:

$$\begin{aligned} O \simeq \Omega BO &\longrightarrow \Omega^2(U/O) \longrightarrow \Omega^3(Sp/U) \longrightarrow \Omega^4 Sp \\ &\simeq \Omega^5 BSp \longrightarrow \Omega^6(U/Sp) \longrightarrow \Omega^7(SO/U) \longrightarrow \Omega^8 SO \\ &\simeq \Omega^8 O. \end{aligned}$$

This proof also proves that the homotopy groups of Sp are shifts of the homotopy groups of O.

Further reading

For an in-depth discussion of the CW structure on U and SU, as well as on other classical Lie groups, see [Yok56, Yok57].

There are many other proofs of Bott periodicity in the literature. Many use tools outside the scope of this book, or else use material from later chapters. For the interested reader we list some here:

- The proof presented in this chapter is Dyer and Lashof's proof [DL61].
- Milnor's *Morse Theory* [Mil63] contains an expanded version of Bott's original proof.

- Atiyah's *K-theory* [Ati89] includes a proof constructing a different version of the Bott map which has an explicit homotopy inverse.

- Atiyah and Bott [AB64] found a proof using clutching functions and some clever uses of linear agebra.

- Aguilar and Prieto [AP99] prove Bott periodicity using quasifibrations; Behrens has a simplification [Beh02].

There are many other proofs in the literature, using many different aspects of topology and geometry.

For a primer on spectral sequences see [HatA]. A more comprehensive and advanced reference is [McC01].

Exercises and Extensions

8.1 Prove that for any topological group G,

$$G \simeq \Omega BG.$$

8.2 This exercise explores some basic properties of H-spaces.

(a) Prove that the fundamental group of any H-space is abelian.

(b) Suppose that $f\colon X \xrightarrow{\sim} Y$ is a weak equivalence of pointed spaces. Prove that if there exists an H-space structure on X there also exists an H-space structure on Y.

(c) Suppose that $f\colon X \longrightarrow Y$ is a map of pointed spaces. Prove that the induced map $\Omega f\colon \Omega X \longrightarrow \Omega Y$ is an H-map.

8.3 Prove Lemma 8.22.

8.4 What is the Pontrjagin ring structure on $H_*(S^1)$? What about $H_*(S^3)$? (The H-space structures are discussed in Example 8.14.)

8.5 What goes wrong if we try to do real Bott periodicity in the same way as the first part of the chapter? In other words, if we naïvely replace BU with BO everywhere, where do things start to go wrong?

Chapter 9

Topological K-theory

In this chapter we introduce topological K-theory, which will be the third cohomology theory discussed in this book. This chapter also brings us full-circle, as topological K-theory is an invariant of vector bundles. In Chapter 3 we showed that vector bundles are classified by homotopy classes of maps into representing spaces. However, in a disappointing twist these were not computable invariants. In this chapter we show how it is possible to make a slightly weaker invariant of vector bundles computable, and discuss many of its properties.

The key observation is that *moving up to higher dimensions makes computations simpleer*. This is the fundamental observation behind various stability properties in algebraic topology. The computation of the cohomology of Grassmannians is complicated for finite-dimensional Grassmannians, but once we move to the infinite-dimensional case the cohomology is much simpler. In the previous chapter, we saw a similar phenomenon arise when increasing the dimension of the subspace, as well, as the colimit over $BU(n)$ was much simpler in structure than each individual space. It is this colimit that exhibits Bott periodicity, and allows for the construction of topological K-theory.

In this chapter we also wrap up the thread on skew field constructions on $\mathbf{R}^n$ started in Chapter 1 and prove that skew field structures exist only when $n = 1$, 2, 4, or 8.

Section 9.1 sets up topological K-theory as a cohomology theory and gives its algebraic definition. Section 9.2 discusses the relationship of topological K-theory and characteristic classes, defines the Chern character, and proves that it is an isomorphism between rationalized K-theory and rational cohomology. Section 9.3 discusses vector bundles on spheres and computes the K-theory of all spheres. Section 9.4 introduces and proves the Splitting Principle, which is very useful for computations, and will be used to prove that Adams Operations exist in the next chapter..

9.1 The definition of topological K-theory

As mentioned in the previous chapter, as topological spaces $U \cong S^1 \times SU$. Since Ω is a right adjoint,

$$\Omega U \cong \Omega(S^1 \times SU) \cong \Omega S^1 \times \Omega SU \simeq \mathbb{Z} \times \Omega SU.$$

Thus the map

$$\mathbb{Z} \times \Phi \colon \mathbb{Z} \times BU \xrightarrow{\sim} \mathbb{Z} \times \Omega SU \simeq \Omega U$$

is a weak equivalence. On the other hand, there is a map

$$\phi \colon U \xrightarrow{\sim} \Omega BU \cong \Omega(\mathbb{Z} \times BU),$$

since for any space X with basepoint component X_0, $\Omega X = \Omega X_0$. This implies that there is an Ω-spectrum

$$K := \mathbb{Z} \times BU, U, \mathbb{Z} \times BU, U, \ldots$$

Here the first map is given by $\mathbb{Z} \times \Phi$, the second by ϕ, the third by $\mathbb{Z} \times \Phi$, and so on.[a]

The cohomology theory represented by K is *topological K-theory.* For any compact space X, $K^0(X)$ is given by formal differences of vector bundles on X. (Recall the notation for cohomology groups in Definition 4.13.) This is the abelian group which is closest to the monoid of vector bundles on X.

Theorem 9.1. *For any compact connected space X,*

$$K^0(X) = \left\{ \begin{array}{c} \textit{free ab. gp. on} \\ \textit{iso classes of vector bundles on } X \end{array} \right\} \Big/ [E \oplus E'] = [E] + [E'].$$

For general i, $K^i(X) = \widetilde{K}^0(\Sigma^i X_+)$.

[a] This spectrum is also often called KU, to differentiate it from KO, the analogously-constructed real K-theory spectrum.

Proof. We prove the second part first. For $i \geq 0$ define

$$K_{-i} \stackrel{\text{def}}{=} \begin{cases} \mathbb{Z} \times BU & \text{if } i \text{ odd,} \\ U & \text{if } i \text{ even,} \end{cases}$$

thus extending the spectrum structure "to the left." Then

$$\begin{aligned} K^i(X) = [X_+, K_i] &\cong [X_+, K_{-|i|}] \cong [X_+, \Omega^i(\mathbb{Z} \times BU)] \\ &\cong [\Sigma^i X_+, \mathbb{Z} \times BU] = \widetilde{K}^0(\Sigma^i X_+), \end{aligned}$$

as desired.

Let G be the group given by the right-hand side of the expression in the theorem. Define a homomorphism $\varphi\colon K^0(X) \longrightarrow G$ in the following manner. Consider a class $\alpha \in [X_+, \mathbb{Z} \times BU]$. Since X is connected, the image of X lies in a single component $\{i\} \times BU$. Since X is compact and $BU = \operatorname{colim}_n BU(n)$, α is represented by a map $f\colon X \longrightarrow \{i\} \times BU$ which factors through $BU(n)$ for some n. Define

$$\varphi[f, i] \stackrel{\text{def}}{=} [f^*\gamma_n] - [\epsilon^{n-i}],$$

where, if $j < 0$, $[\epsilon^j] \stackrel{\text{def}}{=} -[\epsilon^j]$.

To check that φ is well-defined it is necessary to check that it is independent of the choices of f and n. Changing $f\colon X \longrightarrow BU(n)$ up to homotopy changes $f^*\gamma_n$ by an isomorphism, so that the value of φ is unchanged. To show that φ is independent of n it suffices to show that increasing n does not affect φ. Increasing n by 1 corresponds to composing the factorization of f through $BU(n)$ with the map $BU(n) \longrightarrow BU(n+1)$ induced by the map $U(n) \longrightarrow U(n+1)$ adding a 1 in the lower-right corner of the matrix. The pullback of a vector bundle along this map takes γ_{n+1} to $\gamma_n \oplus \epsilon^1$. Thus if we replace n by $n+1 \geq n$ we replace $f^*\gamma_n$ by $f^*\gamma_n \oplus \epsilon^1$. But then, inside G,

$$\begin{aligned} [f^*\gamma_n \oplus \epsilon^1] - [\epsilon^{(n+1)-i}] &= [f^*\gamma_n] \oplus [\epsilon^1] - [\epsilon^{n-i}] - [\epsilon^1] \\ &= [f^*\gamma_n] - [\epsilon^{n-i}]. \end{aligned}$$

Thus $\varphi[f, i]$ is well-defined.

The addition in G is induced by the H-space structure on $\mathbb{Z} \times BU$. As discussed in the previous chapter, this H-space structure is given by the sum on the $\mathbb{Z}$-components and the Whitney sum of representing bundles on the BU-component. Therefore two elements of $[X, \mathbb{Z} \times BU]$ represented by $[f, i]$ and $[f', j]$ have sum represented by $[f \oplus f', i + j]$, and

$$\begin{aligned}\varphi[f \oplus f', i + j] &= [(f \oplus f')^* \gamma_{n+m}] - [\epsilon^{n+m-i-j}] \\ &= [f^* \gamma_n \oplus f'^* \gamma_m] - [\epsilon^{n-i} \oplus \epsilon^{m-j}] \\ &= \varphi[f, i] + \varphi[f', j].\end{aligned}$$

It remains to show that φ is an isomorphism. First, consider surjectivity. Let

$$\alpha = \sum_{i=1}^{n} a_i [E_i]$$

be any element of G. By using the relation in G, we can rewrite this element as $[E] - [E']$, where E is the sum of all E_i with positive coefficient (taken sufficiently many times) and E' is the sum of all E_i with negative coefficient. Since X is compact, by Proposition 6.18, there exists an E'' such that $E' \oplus E'' \cong \epsilon^m$ for some m. Then

$$[E] - [E'] = [E \oplus E''] - [E' \oplus E''] = [E \oplus E''] - [\epsilon^m].$$

Let $n = \dim E \oplus E''$, and let $f\colon X \longrightarrow BU(n)$ be the classifying map of $E \oplus E''$. Then the map

$$X \xrightarrow{f} \{n - m\} \times BU(n) \longrightarrow \{n - m\} \times BU$$

is mapped to α by φ, as desired.

Now consider injectivity. Let $f\colon X \longrightarrow \{i\} \times BU$ be a map such that $\varphi[f, i] = 0$; assume that f factors through $\{i\} \times BU(n)$. Since $\varphi[f, i] = 0$, there exists a bundle E such that $f^* \gamma_n \oplus E \cong \epsilon^{n-i} \oplus E$. The dimension of the left-hand side is $n + \dim E$ and the dimension of the right-hand side is $n - i + \dim E$, so it must be that $i = 0$. In addition, since X is compact there exists a bundle E' such that $E \oplus E' \cong \epsilon^m$; hence

$$f^* \gamma_n \oplus \epsilon^m \cong \epsilon^{n+m}.$$

This implies that the map

$$X \xrightarrow{f} BU(n) \longrightarrow BU(n + m)$$

classifies a trivial bundle, and therefore is null-homotopic. But then

$$X \xrightarrow{f} BU(n) \longrightarrow BU$$

is null-homotopic, and $[f]$ was 0, as desired. Thus φ is injective. □

Corollary 9.2. *Two vector bundles E, E' over a compact base X have $[E] = [E']$ in $K^0(X)$ if and only if there exists $n \geq 0$ such that $E \oplus \epsilon^n \cong E' \oplus \epsilon^n$.*

Proof. Suppose that $E \oplus \epsilon^n \cong E' \oplus \epsilon^n$. Then

$$[E] + [\epsilon^n] = [E'] + [\epsilon^n]$$

in $K^0(X)$, and thus $[E] = [E']$, as desired.

Now suppose that $[E] = [E']$ in $K^0(X)$. This means that they are related using a finite number of the relations in Theorem 9.1. In particular, this implies that there is some vector bundle F such that $E \oplus F \cong E' \oplus F$. Let F' be a complementary bundle to F, so that $F \oplus F' \cong \epsilon^n$ for some n. Then

$$E \oplus \epsilon^n \cong E \oplus F \oplus F' \cong E' \oplus F \oplus F' \cong E' \oplus \epsilon^n,$$

as desired. □

There are two (isomorphic) presentations for $\widetilde{K}^0(X)$, when X is a pointed space. The first of these is that $\widetilde{K}^0(X)$ is the subgroup of $K(X)$ of the classes of maps which factor through the $\{0\} \times BU$ component. In other words, this is the kernel of the map $K^0(X) \longrightarrow \mathbb{Z}$ defined by $[E] \longmapsto \dim E$. Using this presentation, the elements of the form $[E] - [\epsilon^{\dim E}]$ are the generators of $\widetilde{K}^0(X)$.

An alternate presentation for $\widetilde{K}^0(X)$ is that it is the cokernel of the map $K^0(*) \longrightarrow K^0(X)$: i.e., it is the quotient of $K^0(X)$ by the submodule of the trivial bundles. In this case, the generators are still classes vector bundles $[E]$, but there is an additional relation given by stating that $[E] = [E']$ if there exist k, k' such that $E \oplus \epsilon^k \cong E' \oplus \epsilon^{k'}$. In fact, analogously to Corollary 9.2, this is the condition for $[E] = [E']$. In addition, if E and E' are complementary vector bundles then

$$[E] + [E'] = [E \oplus E'] = [\epsilon^{\dim E + \dim E'}] = 0,$$

so $[E'] = -[E]$. In other words, the complementary bundle represents the inverse of a bundle.

Theorem 9.3 (Bott periodicity for K-theory)**.** *Let X be a pointed space. Then for all i,*

$$\widetilde{K}^i(X) \cong \widetilde{K}^{i-2}(X).$$

Proof. For all i,

$$\widetilde{K}^i(X) \overset{\text{def}}{=} [X, K_i] \cong [X, \Omega^2 K_i] \cong [\Sigma^2 X, K_i] = \widetilde{K}^{i-2}(X).$$

Thus $\widetilde{K}^*(X)$ is 2-periodic for all X. □

This theorem can also be proven directly from the presentation in Theorem 9.1. An elementary proof, just from this definition and geometric information about vector bundles, is given in [AB64].

Since K represents a cohomology theory, it has an associated long exact sequence. However, in light of Bott periodicity, this long exact sequence takes on an usually compact form:

Proposition 9.4. *For any cofiber sequence $A \hookrightarrow X \longrightarrow X/A$ of compact spaces there is an exact rectangle*

$$\begin{array}{ccccc} \widetilde{K}^0(X/A) & \longrightarrow & \widetilde{K}^0(X) & \longrightarrow & \widetilde{K}^0(A) \\ \uparrow & & & & \downarrow \\ \widetilde{K}^1(A) & \longleftarrow & \widetilde{K}^1(X) & \longleftarrow & \widetilde{K}^1(X/A). \end{array}$$

Proof. Since $\widetilde{K}$ is a cohomology theory, it follows that for any cofiber sequence $A \hookrightarrow X \longrightarrow X/A$ there exists an exact sequence

$$K^{-1}(A) \longrightarrow \widetilde{K}^0(X/A) \longrightarrow \widetilde{K}^0(X) \longrightarrow \widetilde{K}^0(A) \longrightarrow \widetilde{K}^1(X/A) \longrightarrow K^1(X) \longrightarrow K^1(A).$$

The natural isomorphism $\widetilde{K}^{-1}(X) \cong \widetilde{K}^1(X)$ turns this into the desired rectangle, where the left-hand vertical map is the boundary composed with this isomorphism. □

In light of Theorem 9.1 for the rest of this chapter we will focus exclusively on $K^0(X)$ and $\widetilde{K}^0(X)$.

As is the case in singular cohomology, $K^0(X)$ is a ring. The multiplicative structure arises from the tensor product of vector bundles (see also Exercise 3.7), which can be described on the representing spaces via the following:

$$(\mathbb{Z} \times BU) \times (\mathbb{Z} \times BU) \longrightarrow (\mathbb{Z} \times \mathbb{Z}) \times (BU \times BU) \xrightarrow{\mu_{\mathbb{Z}} \times \mu_{BU}} \mathbb{Z} \times BU,$$

where $\mu_{\mathbb{Z}}$ is the multiplication on the integers, and μ_{BU} is the tensor product of matrices. On $K^0(X)$ this can be represented by

$$[E] \cdot [E'] \stackrel{\text{def}}{=} [E \otimes E'].$$

This structure is also well-defined on $\widetilde{K}^0(X)$, since the total dimension of the right-hand side in

$$([E] - [\epsilon^{\dim E}])([E'] - [\epsilon^{\dim E}]) = [E \otimes E] - [\epsilon^{\dim E} \otimes E'] - [E \otimes \epsilon^{\dim E}] + [\epsilon^{\dim E \dim E'}].$$

is 0.

Just as in ordinary cohomology, maps between spaces induce ring homomorphisms on K-theory.

Lemma 9.5. *Let $f\colon X \longrightarrow Y$ be a map of compact spaces. Then then induced homomorphism $f^*\colon K^0(Y) \longrightarrow K^0(X)$ is a ring homomorphism.*

Proof. The multiplication on $K^0(X)$ is represeented by a structure map on $\mathbb{Z}\times BU$. Since the function f^* is given by precomposing $[Y, \mathbb{Z}\times BU] \longrightarrow [X, \mathbb{Z}\times BU]$ and the multiplication is induced by a map $BU \times BU \longrightarrow BU$ (which is given by postcomposition), f^* is a ring homomorphism. □

9.2 An aside on characteristic classes

Analogously to the Steifel–Whitney classes, for complex bundles we have Chern classes:

Definition 9.6. The n-th *Chern class* the characteristic class determined by the pullback of the generator $c_n \in H^{2n}(G_n)$. For a bundle E over X, we write this as $c_n(E) \in H^{2n}(X)$.

All of the results from Section 6.1 hold in an analogous fashion for complex bundles. Properties (SW2)-(SW4) hold as stated; property (SW1) requires only a shift in the degree of the cohomology group to hold.

From the second presentation of $\widetilde{K}^0(X)$ we get the following observation:

Lemma 9.7. *Let X be a compact pointed space. The function $c_n\colon \widetilde{K}^0(X) \longrightarrow H^{2n}(X)$ given by $[E] \longmapsto c_n(E)$ is well-defined. When $n = 1$ this function is a group homomorphism.*

Proof. It suffices to check that if $[E] = [E']$ in $K^0(X)$ then $c_n(E) = c_n(E')$. By Corollary 9.2, if $[E] = [E']$ then there exists an m such that $E \oplus \epsilon^m \cong E' \oplus \epsilon^m$. But then

$$c_n(E) = c_n(E \oplus \epsilon^m) = c_n(E' \oplus \epsilon^m) = c_n(E'),$$

as desired.

Now suppose that $n = 1$. To check that this is a group homomorphism it suffices to check that $c_1(E \oplus E') = c_1(E) + c_1(E')$. This is exactly the Whitney sum formula for $n = 1$. □

For $n \geq 2$, c_n is *not* a group homomorhpism, since $c_n(E \oplus E') \neq c_n(E) + c_n(E')$ for $n > 1$. However, as functions between groups which are not group homomorphisms are unusual and unnatural, it is very tempting to extend the definition of Chern classes in order to see it as a group homomorphism. Here the total Chern class is tempting. Recall that we define the total Chern class as

$$c(E) \stackrel{\text{def}}{=} 1 + c_1(E) + c_2(E) + \cdots \in H^*(X).$$

Then, by the Whitney sum formula,

$$c(E \oplus E') = c(E)c(E'),$$

which gives the total Chern class the appearance of a group homomorphism. Unfortunately, it is *not* a group homomorphism $K^0(X) \longrightarrow H^*(X)$, since under multiplication $H^*(X)$ is not a group.

The idea of the Chern character is to use Chern classes to define a ring homomorphism

$$\text{ch}\colon K^0(X) \otimes \mathbb{Q} \longrightarrow H^*(X;\mathbb{Q}).$$

From the above lemma it may be tempting to just use c_1. However, this runs into a problem: for line bundles L and L',

$$c_1(L \oplus L') = c_1(L) + c_1(L') = c_1(L \otimes L').$$

(See Exercise 6.5.) This implies that, at least for line bundles, trying to use c_1 in a naïve way is not going to work.

The first observation towards fixing this is to recall that exponentiation takes addition to multiplication. So a simple way of fixing the problem with multiplication above is to define ch using exponentiation. For a line bundle L, define

$$\text{ch}([L]) = 1 + c_1(L) + \frac{1}{2}c_1(L)^2 + \frac{1}{3!}c_1(L)^3 + \cdots = e^{c_1(L)}.$$

This is well-defined in $H^*(X;\mathbb{Q})$, since X is compact and therefore finite-dimensional. For line bundles L and L', $L \otimes L'$ is also a line bundle, and thus we can check that

$$\text{ch}([L \otimes L']) = e^{c_1(L \otimes L')} = e^{c_1(L)}e^{c_1(L')} = \text{ch}(L)\text{ch}(L').$$

Thus, at least on the multiplicative subgroup[b] of vector bundles, ch is a group homomorphism to the multiplicative monoid in $H^*(X;\mathbb{Q})$.

[b] See Exercise 3.8.

To extend this to all bundles, consider first those bundles E which are sums of line bundles: $E \cong L_1 \oplus \cdots \oplus L_n$. Since we would like ch to be a ring homomorphism, we want

$$\mathrm{ch}(E) = \mathrm{ch}(L_1) + \cdots + \mathrm{ch}(L_n).$$

Expanding the formula,

$$\mathrm{ch}(E) = \sum_{m \geq 0} \frac{1}{m!}(c_1(L_1)^m + \cdots + c_1(L_n)^m).$$

Definition 9.8. The *elementary symmetric polynomial* of degree m in n variables $x_1, \ldots, x_n$ is

$$\sigma_m(x_1, \ldots, x_n) \stackrel{\text{def}}{=} \sum_{\{j_1,\ldots,j_m\} \subseteq \{1,\ldots,n\}} x_{j_1} \cdots x_{j_m}.$$

In other words, it is the unique polynomial of degree m which is of degree 1 when any $n-1$ of the variables are fixed and which is invariant under permutations of the variables.

Directly from the definition we see that if $m > n$ then $\sigma_m(x_1, \ldots, x_n) = 0$.

Proposition 9.9 ([Mac15, Section I.2])**.** *For every k there exists a polynomial $q_k(y_1, \ldots, y_k)$, such that for all n*

$$q_k(\sigma_1, \ldots, \sigma_k) = x_1^k + \cdots + x_n^k.$$

Here we write $\sigma_i \stackrel{\text{def}}{=} \sigma_i(x_1, \ldots, x_n)$ for conciseness.

Applying this to the formula for $\mathrm{ch}(E)$ gives

$$\mathrm{ch}(E) = \sum_{m \geq 0} \frac{1}{m!} q_m(\sigma_1(E), \ldots, \sigma_m(E)),$$

where we write $\sigma_i(E) \stackrel{\text{def}}{=} \sigma_i(c_1(L_1), \ldots, c_1(L_n))$.

Lemma 9.10. *For all m, and all vector bundles $E \cong L_1 \oplus \cdots \oplus L_n$,*

$$\sigma_m(E) = c_m(E).$$

Proof. We prove this by induction on m and n. First, notice that if $m > n$ then both sides are 0. If $n = 1$ the lemma holds for all m; if $m = 1$ then it holds for all n. In other words, the lemma holds for all pairs (m, n) where $m = 1$ or $n = 1$.

Fix (m, n). Suppose that the lemma holds for all (i, j) with $i < m$ or $j < n$. (In other words, it holds for all lattice points to the left and below (m, n).) Write $E = E' \oplus L_n$. Then

$$\begin{aligned}\sigma_m(E) &= \sigma_m(E') + \sigma_{m-1}(E')c_1(L_n) = c_m(E')c_0(L_n) + c_{m-1}(E')c_1(L_n) \\ &= c_m(E' \oplus L_n) = c_m(E),\end{aligned}$$

as desired. Thus it also holds for (m, n).

Together with the base cases, by induction the lemma holds for all m and n. ☐

We can now give a definition of the Chern character that is well-defined for all vector bundles.

Definition 9.11. Define

$$\mathrm{ch}\colon K^0(X) \otimes \mathbb{Q} \longrightarrow H^*(X; \mathbb{Q})$$

by

$$\mathrm{ch}(E) = \sum_{m \geq 0} \frac{1}{m!} q_m(c_1(E), \ldots, c_m(E)).$$

Proposition 9.12. ch *is a ring homomorphism.*

Proof. We must show that for all vector bundles E, E', $\mathrm{ch}(E \oplus E') = \mathrm{ch}(E) + \mathrm{ch}(E')$ and $\mathrm{ch}(E \otimes E') = \mathrm{ch}(E)\mathrm{ch}(E')$. First, suppose that both E and E' decompose as sums of line bundles, $E = L_1 \oplus \cdots \oplus L_n$ and $E' = L'_1 \oplus \cdots \oplus L'_m$. The first relation then holds by definition, since the sum also decomposes as a sum of line bundles. For the second,

$$\begin{aligned}\mathrm{ch}(E \otimes E') &= \mathrm{ch}\left(\bigoplus_{i,j} L_i \otimes L'_j\right) = \sum_{i,j} \mathrm{ch}(L_i \otimes L_j) \\ &= \sum_{i,j} \mathrm{ch}(L_i)\mathrm{ch}(L_j) = \mathrm{ch}(E)\mathrm{ch}(E'),\end{aligned}$$

as desired.

Now suppose that E splits as a direct sum of line bundles and E' is arbitrary. Write $f\colon P(E') \longrightarrow X$ for the map from the projectivization of E' to X. Then we have the following commutative diagram:

$$\begin{array}{ccc} K^0(X)\otimes\mathbb{Q} & \xrightarrow{\text{ch}} & \bigoplus_i H^{2i}(X;\mathbb{Q}) \\ {\scriptstyle f^*}\big\downarrow & & \big\downarrow{\scriptstyle f^*} \\ K^0(P(E'))\otimes\mathbb{Q} & \xrightarrow{\text{ch}} & \bigoplus_i H^{2i}(P(E');\mathbb{Q}) \end{array}$$

The two vertical morphisms are both injective by the Leray–Hirsch theorem. Moreover, both f^*E and f^*E' split as a sum of line bundles. Thus

$$\begin{aligned} f^*(\text{ch}(E\otimes E')) &= \text{ch}(f^*E\otimes f^*E') = \text{ch}(f^*(E))\text{ch}(f^*(E')) \\ &= f^*(\text{ch}(E)\text{ch}(E')). \end{aligned}$$

Since f^* is injective, it follows that $\text{ch}(E\otimes E') = \text{ch}(E)\text{ch}(E')$, as desired.

Lastly, suppose that both bundles are arbitrary. An analogous argument to the above shows that pulling back to $P(E)$ (which again induces an injective homomorphism on both K-theory and cohomology) reduces to the previous case. □

The homomorphism ch would have been equally well-defined as a homomorphism $K^0(X)\longrightarrow H^*(X;\mathbb{Q})$, since the $\mathbb{Q}$-coefficients were only necessary after applying Chern classes. We stated it in the given form so that we could write down the following:

Proposition 9.13. *Let X be a finite chain complex with only even-degree cells. Then* ch *is a ring isomorphism.*

To prove this proposition we will need some computations of the topological K-theory of spheres, so we postpone it until the end of the next section.

In fact, Proposition 9.13 can be expressed in a much stronger form:

Theorem 9.14. *Let X be a finite CW-complex. The Chern character induces an isomorphism*

$$\text{ch}\colon K^*(X)\otimes\mathbb{Q}\longrightarrow H^*(X;\mathbb{Q})$$

We leave this theorem to the exercises. In fact, the proof of this theorem also implies that ch induces a ring isomorphism $K^0(X)\otimes\mathbb{Q}\longrightarrow\bigoplus H^{2i}(X;\mathbb{Q})$.

Remark 9.15. Unlike the total Chern class, it is possible to model the Chern character using a map of spectra, modeling it as a natural transformation of cohomology theories.

9.3 Classifying vector bundles on spheres

Consider complex vector bundles on S^n, for $n > 1$. According to the classification theorem, Theorem 3.29,

$$\mathbf{Vect}_m(S^n) = [S^n, G_m] = \pi_n G_m.$$

As discussed earlier, these homotopy groups are not simple to compute. However, they have two advantages over the general classification given in Theorem 3.29: they are abelian groups, rather than sets, and spheres have a particularly simple cover by contractible spaces.

The sphere S^n can be written as $D^n \cup_{S^{n-1}} D^n$; in other words, it is two disks glued along their boundaries.[c] Since each D^n is contractible, every vector bundle on D^n is trivial; thus any vector bundle on S^n can be written in terms of local trivializations with respect to these two sets. Moreover, since there are only two sets, there is exactly one transition function, which can be taken to be a function $S^{n-1} \longrightarrow GL_n(\mathbf{C})$.

Definition 9.16. Let E be a vector bundle on S^n, presented as two trivial bundles $E_+ = D^n \times \mathbf{C}^n$ (E restricted to the northern hemisphere), $E_- = D^n \times \mathbf{C}^n$ (E restricted to the southern hemisphere), and a gluing function $g_{+-}\colon S^{n-1} \times \mathbf{C}^n \longrightarrow S^{n-1} \times \mathbf{C}^n$ which is linear on each fiber. This data can be encoded as a function $g_E\colon S^{n-1} \longrightarrow GL_m(\mathbf{C})$ taking $x \in S^{n-1}$ to $g_{+-}|_{E_x}$.

The clutching function g_E is continuous, and thus represents a class in $\pi_{n-1}GL_m(\mathbf{C})$. This class is actually an invariant of the vector bundle:

Lemma 9.17. *Let E, E' be two vector bundles on S^n. If $g_E \sim g_{E'}$ then $E \cong E'$.*

Proof. Suppose that $g_E \sim g_{E'}$. Let $h\colon S^{n-1} \times I \longrightarrow GL_m(\mathbf{C})$ be the homotopy between them. We can then construct a bundle $\widetilde{E} \longrightarrow S^n \times I$ by using $h(\cdot, t)$ as the clutching function for the bundle above $S^n \times \{t\}$. The restriction of $\widetilde{E}$ to $S^n \times \{0\}$ is then E and the restriction to $S^n \times \{1\}$ is E'. By Lemma 3.19, these are therefore isomorphic, as desired. □

Corollary 9.18. *Since $GL_m(\mathbf{C}) \simeq U(m)$, clutching functions can be assumed to take values in $U(m)$.*

[c] Technically, for this analysis to work it is necessary to cover S^n by open balls, rather than by disks, and consider the intersection. However, due to the classification theorem we know that the sets of vector bundles on homotopy equivalent spaces are in bijection, and we therefore ignore this detail.

Moreover, clutching functions have very simple interactions with sums and tensor products of bundles:

Lemma 9.19. *Let E, E' be vector bundles over S^n with clutching functions g_E and $g_{E'}$, respectively. Then*

$$g_{E\oplus E'} = g_E \oplus g_{E'} \qquad \text{and} \qquad g_{E\otimes E'} = g_E \otimes g_{E'}.$$

Here, the $\oplus$ and $\otimes$ on the right-hand sides are the maps $\oplus\colon U(m)\times U(m') \longrightarrow U(m+m')$ (the block sum of matrices) and $\otimes\colon U(m)\times U(m') \longrightarrow U(mm')$ (the tensor product of matrices).

The proof of this lemma is left as an exercise for the reader.

This lemma is useful for exhibiting relations between vector bundles. Moreover, because $GL_m(\mathbf{C})$ is a Lie group, the group structure often makes such analyses relatively simple.

Example 9.20. In this example we compute $K^0(S^2)$. First, observe that

$$K^0(S^2) \cong [S^2_+, \mathbb{Z}\times BU] \cong \mathbb{Z}\oplus [S^2, BU] \cong \mathbb{Z}\times\mathbb{Z}.$$

Thus the additive structure is $\mathbb{Z}\times\mathbb{Z}$. The first $\mathbb{Z}$ is generated by the trivial bundles; the second $\mathbb{Z}$ is generated by the canonical line bundle on $\mathbf{C}P^1$, which we call H.[d] We can check that these are different using c_1: all of the characteristic classes of trivial bundles are 0, while H has a nontrivial first Chern class.

We claim that H satisfies the relation

$$(H\otimes H)\oplus 1 \cong H\oplus H.$$

We begin by computing the clutching function of H. Consider a point in $\mathbf{C}P^1$ as a pair $[z_0\colon z_1]$. The line above this point is the line consisting of the points $(\lambda z_0, \lambda z_1)$. Thus inside the unit disk we can assume that $z_1 = 1$ and write the trivialization of the bundle as $[z_0/z_1\colon 1]\times\mathcal{C} \longrightarrow (\lambda z_0/z_1, \lambda)$. Outside the unit disk we can assume that $z_0 = 1$ and write the trivialization of the bundle as $[1\colon z_1/z_0]\times\mathcal{C} \longrightarrow (\lambda, \lambda z_1/z_0)$. When $|z_1/z_0| = 1$, the transition function takes $([z_0/z_1\colon 1], \lambda)$ to $([1\colon z_1/z_0], \lambda)$, so this function takes $z\in S^1$ to the function $(z)\in U(1)$.

Using this, we see that the clutching functions for $(H\otimes H)\oplus 1$ and $H\oplus H$ are

$$\begin{pmatrix} z^2 & \\ & 1 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} z & \\ & z \end{pmatrix}.$$

[d]We could also call it γ_{11}, but this could cause confusion with the 11-dimensional canonical bundle. We also prefer H as it is consistent with both [Ati89] and [HatB].

To show that these are isomorphic we just need to show that these clutching functions are homotopic. However, the first one is the pointwise multiplication in ΩU of (z) with itself. The second one is loop addition. For any topological group, these two functions are homotopic (see Exercise 9.2).

From this discussion it follows that $H^2 + 1 = 2H$, or in other words that $(H-1)^2 = 0$. We thus have a natural homomorphism $\mathbb{Z}[H]/(H-1)^2 \longrightarrow K(S^2)$. Since this homomorphism is onto and since the rank of the source is equal to the rank of the target, it must be an isomorphism.

If we think of $K(S^2) \cong \mathbb{Z} \times \mathbb{Z}$ as generated by $[H] - [\epsilon^1]$ and $[\epsilon^1]$ then $\widetilde{K}(S^2)$ is generated by $H-1$, with the relation $(H-1)^2 = 0$.

Remark 9.21. By tracing carefully through this analysis and the definition of the Bott map we can show explicitly that the homomorphism

$$\widetilde{K}^0(X) \longrightarrow \widetilde{K}^0(\Sigma^2 X)$$

can be defined to be multiplication by $H-1$. Thus we can extend the ring structure on $\widetilde{K}^0(X)$ to $\widetilde{K}^*(X)$.

Proposition 9.22. *For all $n \geq 0$,*

$$\widetilde{K}^0(S^n) \cong \begin{cases} 0 & \text{if } n \text{ is odd} \\ \mathbb{Z} & \text{if } n \text{ is even.} \end{cases}$$

For all of these cases, the multiplicative structure is trivial if $n > 0$.

Proof. To prove the odd case it suffices to compute $\widetilde{K}^0(S^1)$. But

$$\widetilde{K}^0(S^1) \cong [S^1, BU] \cong [S^0, \Omega BU] \cong [S^0, U].$$

Since U is connected, this last group is trivial, as desired.

The even case follows from Bott periodicity. The multiplicative structure is trivial because $\widetilde{K}^0(S^{2m}) \cong \widetilde{K}^{-2m}(S^0)$, and the product of two classes in these naturally lands in $\widetilde{K}^{-4m}(S^0)$—which is not the group in question. □

We are now ready to prove Proposition 9.13.

Proof of Proposition 9.13. We prove a stronger statement by induction on the number of cells in X. We claim that not only does the given statement hold, but $K^1(X) = 0$. When X is a single 0-cell, the statement is trivially true. Now suppose that the desired statement holds for X; we wish to show that it holds for $Y = X \cup_f D^{2n}$. There is a cofiber sequence $X_+ \hookrightarrow Y_+ \longrightarrow S^{2n}$ which induces the exact rectangle

$$\begin{array}{ccccc} \widetilde{K}^0(S^{2n}) & \longrightarrow & K^0(Y) & \longrightarrow & K^0(X) \\ \uparrow & & & & \downarrow \\ K^1(X) & \longleftarrow & K^1(Y) & \longleftarrow & K^1(S^{2n}) \end{array}$$

$K^1(X) = 0$ by the inductive hypothesis, and $K^1(S^{2n}) = 0$ by Proposition 9.22. Thus $K^1(Y) = 0$, as desired. Moreover, the top row of the exact square is a short exact sequence.

Since the cohomologies of X,Y, and S^{2n} are concentrated in even degrees, the long exact sequence in cohomology for the given cofiber sequence produces an exact sequence

$$0 \longrightarrow \mathbb{Q} \longrightarrow \bigoplus_i H^{2i}(Y;\mathbb{Q}) \longrightarrow \bigoplus_i H^{2i}(X;\mathbb{Q}) \longrightarrow 0.$$

The two above two exact sequences give a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \widetilde{K}^0(S^{2n}) \otimes \mathbb{Q} & \longrightarrow & K^0(Y) \otimes \mathbb{Q} & \longrightarrow & K^0(X) \otimes \mathbb{Q} & \longrightarrow & 0 \\ & & \downarrow \text{ch} & & \downarrow \text{ch} & & \downarrow \text{ch} & & \\ 0 & \longrightarrow & H^{2n}(S^{2n}) & \longrightarrow & \bigoplus_i H^{2i}(Y;\mathbb{Q}) & \longrightarrow & \bigoplus_i H^{2i}(X;\mathbb{Q}) & \longrightarrow & 0 \end{array}$$

in which every row is exact. The right-hand homomorphism is an isomorphism by the induction hypothesis. Thus to prove the proposition it remains to check that the Chern character is an isomorphism for S^{2n}.

To prove this, it suffices to show that it is nonzero, since both groups are isomorphic to $\mathbb{Q}$. Since $c_1(H-1) \neq 0$, the Chern character is an isomorphism when $n = 1$. For higher n, notice that the Chern character is natural with respect to suspension, in the sense that the diagram

$$\begin{array}{ccc} \widetilde{K}^0(X) & \xrightarrow{\cong} & \widetilde{K}^0(\Sigma^2 X) \\ \downarrow \text{ch} & & \downarrow \cong \\ \bigoplus_i H^{2i}(X;\mathbb{Q}) & \xrightarrow{\text{ch}} & \bigoplus_i H^{2i}(\Sigma^2 X;\mathbb{Q}) \end{array}$$

commutes. When $X = S^2$ we know that the left-hand homomorphism is nonzero; thus, by induction, it will be nonzero for all even speres, as desired. □

We can now use the Chern character to compute the K-theory of general complex projective spaces:

Proposition 9.23. *As a ring, $K^0(\mathbf{C}P^n) \cong \mathbb{Z}[L]/(L-1)^{n+1}$, where L is the canonical line bundle over $\mathbf{C}P^n$.*

Proof. Since $\mathbf{C}P^n$ has only even-dimensional cells, the exact rectangle in K-theory and the computations of $\widetilde{K}^0(S^{2i})$ imply that as a group $K^0(\mathbf{C}P^n) \cong \mathbb{Z}^{n+1}$. The ring structure follows from the fact that the Chern character is a ring homomorphism. □

9.4 The Splitting Principle

As discussed in Example 6.23, some vector bundles can be decomposed as a sum of line bundles, but some cannot. The Splitting Principle is the next-best-thing to being able to split a vector bundle as a sum of line bundles: it states that any bundle can be pulled back to another base in such a way that over the new base it splits. It is important to remember that splitting into a sum of line bundles is *not* the same as being isomorphic to a trivial bundle, as nontrivial line bundles exist.

Theorem 9.24 (Splitting Principle)**.** *For any bundle $p\colon E \longrightarrow X$ with X compact, there exists a compact X' with a map $f\colon X' \longrightarrow X$ such that f^*E splits as a sum of line bundles, and the induced homomorphism $f^*\colon K^0(X) \longrightarrow K^0(X')$ is injective.*

The idea behind the splitting principle lies again in the familiar map $\tau\colon BU(n) \longrightarrow BU(n+1)$. Pulling back the canonical $n+1$-plane bundle along τ produces the canonical n-bundle plus a trivial bundle (Example 3.17). The idea of the Splitting Principle is to generalize this idea to any bundle.

We begin by generalizing the notion of a canonical line bundle from considering only lines in $\mathbf{C}^n$ to more generally considering lines in a vector bundle.

Definition 9.25. For any vector space V, write $P(V)$ for the projectivization of V: this is the space $V/\mathbf{C}^\times$. The points of $P(V)$ are therefore the lines through the origin in V; this is homeomorphic to $\mathbf{C}P^{\dim V - 1}$.

Let $p\colon E \longrightarrow X$ be a rank-n vector bundle. The fiber bundle $g\colon P(E) \longrightarrow X$ has as its total space

$$P(E) = \{(x, v) \in X \times E \mid v \in E_x\}/\langle (x, v) \sim (x, \lambda v), \lambda \in \mathbf{C}^\times \rangle.$$

In other words, $P(E)$ is the space of lines which are contained inside fibers of E. We write points in $P(E)$ as (x, ℓ). The induced projection $p'\colon P(E) \longrightarrow X$ has the structure of a fiber bundle with fiber $\mathbf{C}P^{n-1}$.

There is a canonical line bundle $L \longrightarrow P(E)$, in which the fiber over a point is exactly the line represented by that point. More concretely,

$$L = \{(x, \ell, w) \in P(E) \times E \mid w \in \ell\},$$

with the structure map given by projection to $P(E)$.

Since each fiber of $P(E)$ is compact, and X is compact, so is $P(E)$. It turns out that once we pull E back to $P(E)$ we can split off a line bundle component.

Lemma 9.26. *Let $p\colon E \longrightarrow B$ be a bundle of rank n, and let $p'\colon P(E) \longrightarrow X$ be the associated projectivization. Then L is a subbundle of $(p')^*E \longrightarrow P(E)$, so $(p')^*E$ splits as a sum of a line bundle and a bundle of rank $n-1$.*

Proof. A point in $P(E)$ is a pair (x, ℓ), where $\ell \subseteq E_x$. The fiber over this point in $(p')^*E$ is isomorphic to E_x. Thus the map $L \longrightarrow (p')^*E$ given by $(x, \ell, w) \longmapsto (w \in E_{(x,\ell)})$ is a morphism of vector bundles, proving the lemma. ☐

The first part of the splitting principle follows directly by induction:

Corollary 9.27. *For any compact space X and any bundle $p\colon E \longrightarrow X$ there exists a compact space X' and a map $f\colon X' \longrightarrow X$ such that the bundle $f^*E \longrightarrow X'$ splits as a sum of line bundles.*

However, the splitting principle would not be useful without its second part. After all, a point is compact, and the pullback along the inclusion of a point into any vector bundle is a trivial bundle, which obviously splits into lines. The key point is that the induced homomorphism on K-theory is injective, which means that computations can be done inside $K^0(X')$ instead of $K^0(X)$ without losing any information.

To prove the algebraic portion of the Splitting Principle we need one more ingredient:

Theorem 9.28 (Leray–Hirsch)**.** *Let $p\colon E \longrightarrow B$ be a fiber bundle with E and B compact Hausdorff and with fiber F such that K^*F is free. Suppose that there exist classes $c_1, \dots, c_k \in K^*(E)$ which restrict to a basis in each fiber F of p. If F is a finite cell complex with cells of only even dimensions[e] then as a module, $K^*(E)$ is a free module over $K^*(B)$ with basis $\{c_1, \dots, c_k\}$.*

We omit the proof of the Leray–Hirsch Theorem. The Leray–Hirsch Theorem holds in other cohomology theories as well, and was originally proved for ordinary cohomology. In that form it is a mild generalization of the Thom isomorphism. For a proof of this, see for example [KT06, Theorem 3.1].

Lemma 9.29 (Inductive step of the algebraic portion of the Splitting Principle)**.** *Let $p\colon E \longrightarrow B$ be a bundle of rank n, and let $p'\colon P(E) \longrightarrow X$ be the associated projectivization. Let $L \longrightarrow P(E)$ be the canonical line bundle. Then $K^0(X) \hookrightarrow K^0(P(E))$.*

Proof. Consider $[1], [L], [L^2], \dots, [L^n] \in K^*(P(E))$. Over each fiber these restrict to the powers of the canonical line bundle in $\mathbf{C}P^n$; by Proposition 9.23 this is a basis for $K^*(\mathbf{C}P^n)$. Thus the Leray–Hirsch Theorem applies, and $K^*(P(E))$ is a free $K^*(X)$-module with basis $1, [L], \dots, [L^n]$, and the homomorphism is the inclusion of $K^*(X)$ as $K^*(X)\cdot 1$. In particular the homomorphism $K^*(X) \longrightarrow K^*(P(E))$ is injective. Specializing to K^0 gives the desired statement of the lemma. □

The Splitting principle now follows directly by induction.

Let

$$F(E) = \left\{ (x, \ell_1, \dots, \ell_n) \in X \times P(E)^n \;\middle|\; \begin{array}{c} p(\ell_i) = x \ \forall i \\ \ell_i \perp \ell_j = 0 \ \forall i \neq j \end{array} \right\}.$$

More informally, $F(E)$ is the space whose points are n-tuples of orthogonal lines in a fiber of p. There is a natural projection $g\colon F(E) \longrightarrow X$. This space is called the *flag bundle*. (It is important to keep in mind that we're using the orthogonality structures that we have in order to keep track of each of the lines; in general, each point would be a *flag* of spaces: $V_1 \subseteq V_2 \subseteq \cdots \subseteq V_n$. However, as we have orthogonal complements and each space goes up a single dimension, we obtain an n-tuple of lines instead.) Flag bundles arise in many other situations, especially in algebraic geometry.

Remark 9.30. Since the Leray–Hirsch theorem works for singular cohomology as well, the same proof shows that the inclusion $H^*(X) \longrightarrow H^*(P(E))$ is injective.

[e] This can also be replaced by a condition on B.

Further Reading

For a more classical description of topological K-theory, which starts with the definition and proves that K-theory is a cohomology theory and Bott periodicity directly from it, see [Ati89]. For the reader interested in a more computational and more detailed approach, one place to start is [Kar08B].

Exercises and Extensions

9.1 Let E, E' be vector bundles on S^n, with clutching functions g_E and $g_{E'}$, respectively.

(a) Prove that the clutching function $g_{E\oplus E'}$ of the Whitney sum is given by the composition

$$S^{n-1} \xrightarrow{g_E, g_{E'}} GL_m(E) \times GL_{m'}(E) \xrightarrow{\oplus} GL_{m+m'}(E),$$

where $\oplus$ takes the block diagonal sum of matrices.

(b) Prove that the clutching function $g_{E\otimes E'}$ of the tensor product of the bundles is given by the composition

$$S^{n-1} \xrightarrow{g_E, g_{E'}} GL_m(E) \times GL_{m'}(E) \xrightarrow{\otimes} GL_{mm'}(E),$$

where $\otimes$ takes the tensor product of matrices.

(c) Prove that if $E \cong E'$ then $g_E \sim g_{E'}$, showing that the homotopy class of the clutching function is a *complete* invariant of the bundle.

9.2 Let G be a topological group. There are two different operations on $\pi_1 G$: the first one is the usual concatenation of loops, whereas the second one takes two loops $f, g\colon [0,1] \longrightarrow G$ and defines $(f \cdot g)(t) = f(t)g(t)$. Prove that these are homotopic, so that they are the same operation on $\pi_1 G$.

9.3 Let $p\colon E \longrightarrow B$ be a vector bundle. Let $g\colon F(E) \longrightarrow X$ be the flag bundle of E. Prove directly (i.e. not by induction) that the bundle g^*E splits as a direct sum of line bundles.

9.4 Let $p\colon E \longrightarrow B$ be a vector bundle of rank n. Prove that $P(E) \longrightarrow X$ is a fiber bundle with fiber $\mathbf{C}P^{n-1}$.

9.5 Use the Chern character to prove that $\widetilde{K}(S^{2n})$ is generated by the bundle $(H-1)^n$.

9.6 Use the Leray–Hirsch theorem to compute $K^0(F(E))$ in terms of $K^0(X)$. (Hint: look carefully at the proof of Lemma 9.29.)

9.7 Write

$$H^{even}(X;\mathbb{Q}) = \bigoplus_{i\in\mathbb{Z}} H^{2i}(X;\mathbb{Q}) \qquad \text{and} \qquad H^{odd}(X;\mathbb{Q}) = \bigoplus_{i\in\mathbb{Z}} H^{2i+1}(X;\mathbb{Q}).$$

(a) Prove that for any cofiber sequence $A \hookrightarrow X \longrightarrow X/A$ there is an exact rectangle

$$\begin{array}{ccccc} H^{even}(X/A) & \longrightarrow & H^{even}(X) & \longrightarrow & H^{even}(A) \\ \uparrow & & & & \downarrow \\ H^{odd}(A) & \longleftarrow & H^{odd}(X) & \longleftarrow & H^{odd}(X/A). \end{array}$$

(b) Use this, together with the definition of $K^1(X)$, to prove Theorem 9.14.

Chapter 10

Adams Operations

For many of the applications of topological K-theory, we do not need the entire structure of $K^*(X)$; the group $K^0(X)$ suffices. One of the main reason for this is the *Adams operations*, which create a very rigid structure on the groups $K^0(X)$. These do not exist in ordinary cohomology, and thus some computations that are more difficult in ordinary cohomology are simpler in K^0.

In this chapter we introduce Adams operations abd use them to resolve a classical question about maps of Hopf invariant 1 (which also control the existence of spaces with polynomial cohomology, H-space structures on $\mathbf{R}^n$, and the parallelizability of spheres).

We begin the chapter with a short section containing the main properties of the Adams operations. In practice, this is often all that is necessary in order to use them for computations. In Section 10.2 we use these properties to prove the nonexistence of maps of Hopf invariant 1 above dimension 8, and use this result to wrap up the geometric applications that we have been working with. The rest of the chapter is concerned with proving that these operations exist . The construction of the Adams operations is in Section 10.3, with the proofs of the definition on line bundles and of naturality in Lemma 10.16 and Proposition 10.17, respectively. The proof that the Adams operations are ring homomorphisms is in Theorem 10.18. Uniqueness is proved in Lemma 10.19.

10.1 An overview of the properties of Adams operations

The main theorem about Adams operations is the following:

Theorem 10.1. *For every compact space X there exist unique ring homomorphisms*

$$\psi^k\colon K^0(X) \longrightarrow K^0(X)$$

which are

powers on line bundles *for every line bundle L over X,*

$$\psi^k[L] = [L^{\otimes k}], \text{ and}$$

natural *for every homomorphism $f\colon Y \longrightarrow X$ the square*

$$\begin{array}{ccc} K^0(X) & \xrightarrow{\psi^k} & K^0(X) \\ {\scriptstyle f^*}\downarrow & & \downarrow{\scriptstyle f^*} \\ K^0(Y) & \xrightarrow{\psi^k} & K^0(Y) \end{array}$$

commutes.

The Adams operations have several other nice properties:

Theorem 10.2. *For all compact X, the Adams operations satisfy the following extra properties:*

(i) *For all $k, \ell \geq 1$, $\psi^k \circ \psi^\ell = \psi^{k\ell}$.*

(ii) *For any prime p, $\psi^p(\alpha) = \alpha^p \pmod p$, in the sense that for all α there exists a $\beta \in K^0(X)$ such that $\psi^p(\alpha) = \alpha^p + p\beta$.*

(iii) *When $X \cong S^{2n}$, for any class α with total dimension 0, $\psi^k(\alpha) = k^n\alpha$.*

10.2 Maps of Hopf Invariant 1

The most famous application of Adams operations relates to the nonexistence of maps of Hopf invariant one above dimension 8.

Definition 10.3. The attaching map of a $2n$-cell is a map $f\colon S^{2n-1} \longrightarrow S^n$. We define the *Hopf invariant* of f to be the integer h such that $\alpha \smile \alpha = h\beta$ (for α the cohomology class represented by S^n and β the cohomology class represented by the $2n$-cell attached by f).

When n is odd the Hopf invariant is always 0, so we restrict attention to the case when n is even. For $n = 2, 4, 8$ there is a map of Hopf invariant 1 by considering the attaching map of the $2n$-cell in $\mathbf{C}P^2$, $\mathbf{H}P^2$, or $\mathbf{O}P^2$, respectively.

Let $X = S^n \cup_f D^{2n}$. Then there exists a cofiber sequence

$$S^n \hookrightarrow X \longrightarrow S^{2n}.$$

Since $\widetilde{K}$ is an evenly graded cohomology theory on spheres by Proposition 9.22, we get the following short exact sequence:

$$0 \longrightarrow \widetilde{K}^0(S^{2n}) \longrightarrow \widetilde{K}^0(X) \longrightarrow \widetilde{K}^0(S^n) \longrightarrow 0.$$

Let α' be the generator of $\widetilde{K}^0(S^{2n})$ and let β' be the generator of $\widetilde{K}^0(S^n)$. (Both of these groups are $\mathbb{Z}$ by Proposition 9.22.) Write α for the image of α' in $\widetilde{K}^0(X)$ and β for any preimage of β' in $\widetilde{K}^0(X)$. Since multiplication is trivial in $\widetilde{K}^0(S^n)$, $\beta'^2 = 0$; this implies that $\beta^2 = k\alpha$ for some integer k.

Proposition 10.4. *k is the Hopf invariant of f.*

Proof. The Chern character is a ring homomorphism $\mathrm{ch}\colon \widetilde{K}^0(X)\otimes\mathbb{Q} \longrightarrow \widetilde{H}^*(X;\mathbb{Q})$. Thus

$$k\mathrm{ch}(\alpha) = \mathrm{ch}(\beta^2) = \mathrm{ch}(\beta)^2 = h\mathrm{ch}(\alpha).$$

Since both are $\mathbb{Q}$-vector spaces it follows that $k = h$, as desired. □

The Hopf invariant is partially interesting because it controls how "polynomial" a simple cohomology ring can be. Indeed, suppose that we want to construct a cell complex X with 3 cells such that the cohomology is $\mathbb{Z}[x]/x^3$, where $|x| = n$ for some n? When $n = 2$ we can do this by setting $X = \mathbf{C}P^2$. When $n = 4$ we can do this by setting $H = \mathbf{H}P^2$, and when $n = 8$ we can do this by setting $X = \mathbf{O}P^2$.

We now assume $n > 1$.

Theorem 10.5. *k can equal ± 1 only if $n = 2, 4, 8$.*

Proof. Write $n = 2m$. It suffices to check that for $m \neq 1, 2, 4$, $\beta^2 \equiv 0 \pmod 2$. Using Theorem 10.2(ii), we note that it is equivalently sufficient to show that $\psi^2(\beta) = 0 \pmod 2$.

Write $n = 2m$. By Theorem 10.2(iii),

$$\psi^2(\alpha) = 2^m\alpha \qquad \text{and} \qquad \psi^3(\alpha) = 3^m\alpha$$

and

$$\psi^2(\beta) = 2^m\beta + \mu\alpha \qquad \text{and} \qquad \psi^3(\beta) = 3^m\beta + \lambda\alpha$$

for some integers μ and λ. It therefore suffices to show that μ is even. By Theorem 10.2(i) $\psi^2\psi^3 = \psi^3\psi^2$, and therefore

$$2^m(3^m\beta + \lambda\alpha) + \mu 3^{2m}\alpha = 3^m(2^m\beta + \mu\alpha) + \lambda 2^{2m}\alpha.$$

Rearranging this we get

$$3^m(3^m - 1)\mu = 2^m(2^m - 1)\lambda.$$

We'll need the following two observations:

(O1) When m is odd,

$$3^m - 1 \equiv 2 \pmod 8 \qquad \text{and} \qquad 3^m + 1 \equiv 3 + 1 = 4 \pmod 4$$

(O2) When m is even,

$$3^m + 1 \equiv 1 + 1 = 2 \pmod 8.$$

Let $\nu(m)$ be the largest power of 2 dividing $3^m - 1$. To show that μ is even it suffices to check that $\nu(m) < m$. When m is odd, by (O1), $\nu(m) = 1$ and thus $2|\mu$ unless $m = 1$.

Now suppose that $m = 2^\ell j$, with j odd. Then we can write

$$3^{2^\ell j} - 1 = (3^{2^{\ell-1}j} - 1)(3^{2^{\ell-1}j} + 1) = (3^{2^{\ell-2}j} - 1)(3^{2^{\ell-2}j} + 1)(3^{2^{\ell-1}j} + 1)$$
$$= \cdots = (3^j - 1)\prod_{L=0}^{\ell}(3^{2^L j} + 1).$$

By (O2) each of the terms in the product is divisible by exactly one power of 2 except when $L = 0$, which gives 2; the first term in the product is also divisible by exactly one power of 2. Thus

$$\nu(2^\ell j) = \ell + 2.$$

In particular, $\nu(m) \geq m$ exactly when $2^\ell j \leq \ell + 2$. $2^\ell \leq \ell + 2$ when $\ell = 0, 1, 2$; in all of these cases we cannot have $j > 1$, so these ar the only solutions. Thus the only possible values for m are $1, 2, 4$, as desired. □

To connect this back to the question of parallelizability of projective spaces recall from Figure 1.1 that an H-space structure on S^{n-1} produces a map of Hopf invariant 1. Since we have now limited the possible dimensions of maps of Hopf invariant 1, this automatically limits the dimensions of H-space structures on S^{n-1}, and thus also the paralellizability of $\mathbf{R}P^n$ and S^n.

10.3 Constructing Adams Operations

We have claimed some strong implications of the existence of Adams operations, but we have not yet actually constructed them. The rest of this chapter will be taken up with the theory necessary to construct Adams operations. Although the theory may seem opaque and specialized, it turns out to be surprisingly versatile, with applications in representation theory, algebraic geometry, and knot theory.

The idea of Adams operations is to construct an operation that behaves just like a "pure power," but also works correctly with respect to addition. Thus, for example, we cannot consider the function $x \longmapsto x^2$ to be a "pure power," because it is not compatible with addition:

$$(a+b)^2 = a^2 + 2ab + b^2 \neq a^2 + b^2.$$

If, instead, we say that the "degree-2" operation on k variables $a_1, \ldots, a_k$ is

$$\sigma_1^2 - 2\sigma_2 = a_1^2 + \cdots + a_k^2,$$

where σ_1 and σ_2 are the symmetric polynomials on k variables of degree 1 and 2, respectively, then the formula works *independently of* k. This is the idea of Adams operations: to take a formalization of the notion of symmetric polynomial and define the "pure powers" in terms of these.

Definition 10.6. Let R be a commutative ring. A *pre-λ-ring structure*[a] on R is operations $\lambda^n \colon R \longrightarrow R$ for integers $n \geq 0$ satisfying the following relations for all $r, s \in R$:

(L1) $\lambda^0 r = 1$.

(L2) $\lambda^1 = 1_R$.

(L3) $\lambda^n(r+s) = \sum_{k=0}^{n} \lambda^k(r)\lambda^{n-k}(s)$.

A *morphism of pre-λ-rings* is a ring homomorphism which commutes with the λ-operations.

The operations λ^* are designed to model operations that behave akin to exterior powers of vector spaces. Thus Axiom (L1) states that a 0-fold skew-symmetric product is the unit, and Axiom (L2) states that a 1-fold skew-symmetric product is the identity. The final axiom states that this skew-symmetic product distributes across addition in an appropriately "graded" manner.

[a]Sometimes these are called *λ-rings* and what we call "λ-rings" are called *special λ-rings*. This does indeed occasionally lead to confusion.

Example 10.7. Let $R = \mathbb{Z}$, and let $\lambda^m(n) \stackrel{\text{def}}{=} \binom{n}{m}$. Axioms (L1) and (L2) hold automatically; Axiom (L3) holds because

$$\binom{n+n'}{m} = \sum_{j+k=m} \binom{n}{j}\binom{n'}{k}.$$

Example 10.8. There exists a pre-λ-ring structure on $K^0(X)$ using the exterior product of bundles. More concretely, define $\lambda^n[E] = \bigwedge^n E$. This also satisfies the additional property that

$$\lambda^n(E) = 0 \qquad \text{if } n > \dim E,$$

although this is not necessary for a general pre-λ-ring structure.

Axioms (L1) and (L2) hold by definition. Axiom (L3) is satisfied because for any vector spaces V and W, the vector space $\bigwedge^n(V \oplus W)$ is spanned by vectors of the form $v_1 \wedge \cdots \wedge v_k \wedge w_{k+1} \wedge \cdots \wedge w_n$, with $v_1, \ldots, v_k \in V$ and $w_{k+1}, \ldots, w_n \in W$,

A pre-λ-ring structure is a generalization of the way that symmetric polynomials behave. This can be more rigorously explained by constructing a pre-λ-ring structure on a ring whose underlying abelian group is the *multiplicative* group of formal power series with constant term 1. The multiplication in this ring will be "pairwise multiplication of roots"; the pre-λ-ring structure will be "elementary symmetric polynomials in the roots."

Example 10.9. Let R be a commutative ring. Let $\Lambda(R)$ be the abelian group of power series with coefficients in R and constant term equal to 1, with operation being multiplication. $\Lambda(R)$ has an associated multiplication, constructed in the following manner.

Consider the following formal products:

$$\alpha(t) = \prod_{n\geq 1}(1+\xi_n t) \qquad \text{and} \qquad \beta(t) = \prod_{n\geq 1}(1+\eta_n t).$$

These can be formally expanded to series

$$\alpha(t) = 1 + \sum_{n\geq 1} a_n t^n \qquad \text{and} \qquad \beta(t) = 1 + \sum_{n\geq 1} b_n t^n,$$

where the coefficients are allowed to be infinite formal sums in R. Similarly consider the formal product

$$\prod_{m,n}(1+\xi_m\eta_n t) = 1 + \sum_{n\geq 1} P_n t^n.$$

Note that P_n is a symmetric series in the ξ_i, η_i, and thus it can be written as a series in terms of a_i's and b_i's (whose coefficients are the elementary symmetric series in the ξ_i and η_j, respectively). In fact, by degree considerations, P_n will only depend on $a_1, \ldots, a_n, b_1, \ldots, b_n$. For example,

$$a_1 = \sum_{n \geq 1} \xi_n \qquad b_1 = \sum_{n \geq 1} \eta_n$$

and

$$P_1 = \sum_{n,m \geq 1} \xi_n \eta_m = a_1 b_1.$$

For P_2,

$$\begin{aligned} P_2 &= \sum_{\substack{m_1, m_2 \\ n_1, n_2 \\ (m_1,n_1) \neq (m_2,n_2)}} \xi_{m_1} \xi_{m_2} \eta_{n_1} \eta_{n_2} \\ &= 2 \sum_{\substack{m_1 < m_2 \\ n_1 < n_2}} \xi_{m_1} \xi_{m_2} \eta_{n_1} \eta_{n_2} + \sum_{m, n_1 < n_2} \xi_m^2 \eta_{n_1} \eta_{n_2} + \sum_{m_1 < m_2, n} \xi_{m_1} \xi_{m_2} \eta_n^2 \\ &= 2a_2 b_2 + b_2 \sum_m \xi_m^2 + a_2 \sum_n \eta_n^2 = a_2 b_2 + b_2(a_1^2 - 2a_2) + a_2(b_1^2 - 2b_2) \\ &= b_2 a_1^2 + a_2 b_1^2 - 2a_2 b_2. \end{aligned}$$

By the general theory of symmetric polynomials, it will always be possible to express P_n in terms of $a_1, \ldots, b_n$. In fact, in order to compute P_n it is not necessray to assume that there are infinitely many variables η_i and ξ_j; if there are at least n of each then the calculation of P_n in terms of the a_i and b_j will be correct.

Define a multiplication on $\Lambda(R)$ by

$$\alpha(t) * \beta(t) \stackrel{\text{def}}{=} 1 + \sum_{n \geq 1} P_n(a_1, \ldots, b_n) t^n.$$

To obtain a pre-λ-structure on $\Lambda(R)$, we define

$$\prod_{i_1 < \cdots < i_n} (1 + \xi_{i_1} \cdots \xi_{i_n}) = 1 + \sum_{m \geq 1} L_{n,m} t^m.$$

By a similar argument to the above, these will only depend on the a_i and b_j. For example,

$$L_{2,1} = \sum_{n < m} \xi_n \xi_m = a_2.$$

Define

$$\lambda^n\alpha(t) \stackrel{\text{def}}{=} 1 + \sum_{m\geq 1} L_{n,m}t^m.$$

A pre-λ-ring structure on R defines a homomorphism of abelian groups $\lambda\colon R \longrightarrow \Lambda(R)$ given by

$$r \longmapsto \sum_{n\geq 0} \lambda^n(r)t^n.$$

Definition 10.10. R is a *λ-ring* if $\lambda\colon R \longrightarrow \Lambda(R)$ is a morphism of pre-λ-rings.

An analysis of this definition produces following alternate formulation:

Definition 10.11. A *λ-ring* R is a pre-λ-ring satisfying the following extra conditions:

(L4) $\lambda^n(1) = 0$ for $n > 1$.

(L5) $\lambda^n(rs) = P_n(\lambda^1(r), \ldots, \lambda^n(r), \lambda^1(s), \ldots, \lambda^n(s))$.

(L6) $\lambda^n(\lambda^m(r)) = L_{n,m}(\lambda^1(r), \ldots, \lambda^{mn}(r))$.

Example 10.12. Continuing on from Example 10.8, $K^0(X)$ is a λ-ring.

(L4) holds because for any vector space V of dimension n, $\bigwedge^m V = 0$ for $m > n$. For any vector space V, we can write $V \cong V_1 \oplus \cdots \oplus V_m$ for linear subspaces V_i; for any other space W write $W \cong W_1 \oplus \cdots \oplus W_n$. Then

$$V \otimes W \cong \sum_{\substack{1\leq i\leq m\\ 1\leq j\leq n}} V_i \otimes W_j.$$

In other words, the tensor product $V \otimes W$ is a "symmetric polynomial" in the variables V_i and W_j. The polynomial P_ℓ tells us exactly how to get this polynomial in terms of the symmetric polynomials in V_i and the symmetric polynomials in W_j; thus $\lambda^\ell(V \otimes W) = P_\ell(\lambda^1 V, \ldots, \lambda^\ell W)$.

A similar explanation works for $\lambda^m\lambda^n(V)$ and is left to the reader.

Example 10.13. Continuing on from Example 10.7, $\mathbb{Z}$ is a λ-ring. Axiom (L4) holds (because $\binom{1}{n} = 0$ for $n > 1$). Axiom (L5) holds because the computations of the polynomials P_n correspond to expressing the method for picking n objects out of an $r \times s$ grid. Similarly, Axiom (L6) will hold because the computation for $L_{n,m}$ will correspond exactly to a method for picking n subsets out of the m-element subsets of $\{1, \ldots, r\}$. Alternately, using Example 10.12 one notes that $\lambda^n(r)$ is exactly the dimension of $\lambda^n(V)$ in Example 10.12 when $\dim V = r$.

Thus a λ-ring is one where symmetric polynomials "work correctly."

Once we have a theory of symmetric polynomials we can use Newton's identities to define good "power operations."

Definition 10.14. Recall the polynomial q_k from Proposition 9.9. We define the *k-th Adams operation* $\psi^k\colon K^0(X) \longrightarrow K^0(X)$ to be

$$\psi^k([E]) = q_k(\lambda^1([E]), \dots, \lambda^k([E])).$$

As a final remark, we note that using the theory of generating functions there is a somewhat-more-compact definition of Adams operations.

Definition 10.15. Suppose that R is a λ-ring. Write

$$\lambda_t(r) \stackrel{\text{def}}{=} 1 + \sum_{n>0} \lambda^n(r)t^n \in R[\![t]\!].$$

Define

$$\psi_t(r) \stackrel{\text{def}}{=} -t\frac{d}{dt}\log \lambda_{-t}(r) = \sum_{k=1}^{\infty} \psi^k(r)t^k.$$

10.4 Properties of Adams operations

In this section we prove several important properties of the Adams operations. First we give an explicit description of the Adams operations on line bundles.

Lemma 10.16. *Let L be a line bundle over a compact X. Then*

$$\psi^k([L]) = [L]^k = [L^{\otimes k}].$$

Proof. By definition, $\lambda^n([L]) = 0$ for $n > 1$, and

$$\psi^k([L]) = s_k([L], 0, \dots, 0).$$

The polynomials s_k always start with the k-th power of the first variable, and all other terms always have at least one higher-indexed term. Thus

$$\psi^k([L]) = [L]^k = [L^{\otimes k}].$$

□

The Adams operations are also natural:

Proposition 10.17. *Let $f\colon X \longrightarrow Y$ be a map of compact spaces. Then the induced map $f^*\colon K^0(Y) \longrightarrow K^0(X)$ commutes with the Adams operations; i.e., for all $k > 0$ the following diagram commutes:*

$$\begin{array}{ccc} K^0(Y) & \xrightarrow{\psi^k} & K^0(Y) \\ {\scriptstyle f^*}\big\downarrow & & \big\downarrow{\scriptstyle f^*} \\ K^0(X) & \xrightarrow{\psi^k} & K^0(X). \end{array}$$

Proof. Since f^* is a ring homomorphism, it suffices to check that λ^n commutes with f^* for all n. Since $\lambda^n[E] = [\bigwedge^n E]$ it suffices to check that $\wedge$ is natural—which it is, because it is classified by a map on Grassmannians. □

In fact, the Adams operations are not only natural, they are *ring homomorphisms*. This statement is not at all obvious; it is not even direct to show that they are homomorphisms of abelian groups.

Theorem 10.18. *For each $k \geq 1$, $\psi^k\colon K^0(X) \longrightarrow K^0(X)$ is a ring homomorphism.*

Proof. Suppose we are given two vector bundles E and E', and we wish to compute $\psi^k(E \oplus E')$. Let $f\colon X' \longrightarrow X$ be a map that splits E; then the pullback of $E \oplus E'$ is $L_1 \oplus \cdots \oplus L_n \oplus f^*E'$. Then let $f'\colon X'' \longrightarrow X'$ be a map that splits f^*E'; this pulls back to $L_1 \oplus \cdots \oplus L_n \oplus L'_1 \oplus \cdots \oplus L'_{n'}$. By construction we know that

$$\psi^k(L_1 \oplus \cdots \oplus L'_{n'}) = \psi^k(L_1) + \cdots + \psi^k(L'_{n'})$$

holds; thus $\psi^k(E \oplus E') = \psi^k(E) + \psi^k(E')$, since the following diagram commutes:

$$\begin{array}{ccc} K^0(X) & \xrightarrow{\psi^k} & K^0(X) \\ {\scriptstyle (f'f)^*}\big\downarrow & & \big\downarrow{\scriptstyle (f'f)^*} \\ K^0(X'') & \xrightarrow{\psi^k} & K^0(X'') \end{array}$$

Thus on bundles ψ^k is additive. Since $K^0(X)$ is a group completion and ψ^k is additive on generators, it is also additive on all of $K^0(X)$.

Now consider $E \otimes E'$. It pulls back along $f'f$ to $\bigoplus L_i \otimes L'_j$, which is a sum of line bundles; thus

$$\psi^k(\bigoplus L_i \otimes L'_j) = \bigoplus L_i^k \otimes {L'_j}^k = \left(\bigoplus L_i^k\right) \otimes \left(\bigoplus {L'_j}^k\right) = \psi^k(E)\psi^k(E').$$

Thus ψ^k is a ring homomorphism. □

It is now possible to complete the proof of Theorem 10.1 by showing that the Adams operations are unique.

Lemma 10.19. *The Adams operations are uniquely determined by the properties (other than uniqueness) listed in Theorem 10.1.*

Proof. Let $\widetilde{\psi}^k \colon K^0(X) \longrightarrow K^0(X)$ be a ring homomorphism satisfying naturality and being a power on line bundles. We claim that $\widetilde{\psi}^k$ is the k-th Adams operation. For any class $[E] \in K^0(X)$, let $f \colon X' \longrightarrow X$ be such that $f^*E \cong \bigoplus_{i=1}^n L_i$ splits as a sum of line bundles L_i and the homomorphism $K^0(X) \longrightarrow K^0(X')$ is injective. Then

$$\widetilde{\psi}^k[E] = (f^*)^{-1}(\widetilde{\psi}^k[f^*E]) = (f^*)^{-1}\sum_{i=1}^n [L_i^k] = (f^*)^{-1}\psi^k[f^*E] = \psi^k[E].$$

Thus $\widetilde{\psi}^k$ is an Adams operation, as desired. □

We can also prove the extra properties claimed in Theorem 10.2. Just like all of the proofs above, the theorem is proved by applying the Splitting Principle and the fact that Adams operations are powers on line bundles.

Proof of Theorem 10.2. Let E be any bundle over X, and suppose that $f \colon X' \longrightarrow X$ is a map such that $f^*E \cong \bigoplus_{i=1}^n L_i$ with each L_i a line bundle and $f^* \colon K^0(X) \longrightarrow K^0(X')$ is injective. For any $k, \ell \geq 1$ diagram

$$\begin{array}{ccccc} K^0(X) & \xrightarrow{\psi^k} & K^0(X) & \xrightarrow{\psi^\ell} & K^0(X) \\ {\scriptstyle f^*}\downarrow & & \downarrow{\scriptstyle f^*} & & \downarrow{\scriptstyle f^*} \\ K^0(X') & \xrightarrow{\psi^k} & K^0(X') & \xrightarrow{\psi^\ell} & K^0(X') \end{array}$$

commutes. Then

$$\psi^\ell\psi^k[f^*E] = \psi^\ell\psi^k\sum_{i=1}^n [L_i] = \sum_{i=1}^n [L_i^{k\ell}] = f^*(\psi^{k\ell}E).$$

Since f^* is injective and E was arbitrary, $\psi^\ell \psi^k = \psi^{k\ell}$ as desired.

For Property (ii), suppose that $k = p$ is prime. Then

$$\psi^p([f^*E] - [\epsilon^n]) = \sum_{i=1}^{n} [L_i^p] - [\epsilon^n] = \left(\sum_{i=1}^{n} [L_i] - [\epsilon^n] \right)^p - p([E'] - [\epsilon^{n'}])$$

for a bundle E' and integer n'.

Now consdider (iii). First suppose that $n = 1$. Since $\widetilde{K}^0(S^2)$ is generated by $L - 1$, it is the case that

$$\psi^k(L-1) = L^k - 1 = (1 + (L-1))^k - 1 = 1 + k(L-1) - 1 = k(L-1),$$

since $(L-1)^i = 0$ for $i > 1$. Property (iii) then follows from the fact that $\widetilde{K}^0(S^{2n}) \cong \widetilde{K}^0(S^2) \otimes \cdots \otimes K^0(S^2)$, with the Adams operations acting termwise, as proved in Proposition 9.22. □

Further Reading

For a more classical description of topological K-theory, which starts with the definition and proves that K-theory is a cohomology theory and Bott periodicity directly from it, see [Ati89].

The paper from which the solution of the Hopf Invariant 1 problem is taken is [AA66], which also contains an interesting discussion of p-th powers in cohomology.

Exercises and Extensions

10.1 Show that the λ-ring structure on $\mathbb{Z}$ given in Example 10.7 is the only possible λ-ring structure on $\mathbb{Z}$.

10.2 Complete the argument in Example 10.12 to show that $K^0(X)$ is a λ-ring.

10.3 Prove that $\Lambda(R)$ is a λ-ring.

10.4 Compute P_3 and $L_{m,n}$ for $m, n \leq 2$.

10.5 Verify that the definition of Adams operations in Definition 10.15 agrees with the more explicitly-given definition.

10.6 Compute the Adams operations on $\widetilde{K}^0(S^{2n})$. (Hint: Use the result of Exercise 9.5.) Use this to prove that $S^2 \vee S^4 \not\simeq \mathbf{C}P^2$.

Chapter 11

Next Directions

I do not claim that this book is a comprehensive introduction to anything at all. I sincerely hope that something in this book has sparked the reader's interest and curiosity.[a] In this chapter I include various references that I used, sources for deeper dives into topics that I glossed over, and next topics that the reader may find interesting. This chapter should be read as my personal opinion on what is interesting, rather than as a survey of the field.

All books and papers mentioned have complete entries in the bibliography.

11.1 Classic Textbooks and Topics

This textbook could not have been written without the classic textbooks in the field:

- John Milnor and Jim Shasheff's *Characteristic Classes*.
- Allen Hatcher's *Vector Bundles and K-theory*.
- Michael Atiyah's *K-theory*.
- Dale Husemoller's *Fibre Bundles*.

These books contain much more material than it is possible to package into this tiny book. Almost every topic in this book is discussed in more detail and depth in (at least) one of these books. For the reader interested in learning more about this material, these are highly recommended.

[a] Or, let's face it, annoyance at my flippant treatment of a deep topic.

Haynes Miller also has a set of notes called “Vector Fields on Spheres” at, at least at the time of writing, it was still possible to find on the internet. They are highly recommended as further exploration of vector fields on spheres, Clifford algebras, and other related topics. The classic reference is Adams’s paper “Vector fields on spheres.”

For the reader interested in these things called “spectral sequences” that I kept referring to, I recommend Allen Hatcher’s *Spectral Sequences*.

For an introduction to spectra, I recommend the beginning of Stefan Schwede’s unpublished book *Symmetric Spectra*. (It is not necessary to understand all of the examples; I suggest skipping examples that get too technical unless they are of specific interest.) The classic book remains Adams’s *Stable Homotopy and Generalized Homology*, with the standard warning that Part III is meant to be read first.

The proof presented of the nonexistence of maps of Hopf invariant 1 is not the first proof published. The first was published in “On the non-existence of elements of Hopf invariant one” [Ada60], and relied heavily on the structure of the Steenrod algebra. The Steenrod algebra is a set of cohomology operations that has been instrumental in much of the modern computational and analytic progress in homotopy theory. References abound, but I recommend starting with Robert Mosher and Martin Tangora’s “Cohomology Operations and Applications in Homotopy Theory.”

11.2 Category Theory

Category theory is a deep rabbit hole that I am not certain I want to recommend people jump down with no preparation. (I say this as someone already thirty miles down and still gleefully falling.) The classic textbook is Saunders Mac Lane’s *Categories for the Working Mathematician*, although I heartily do not recommend this book. It is a classic for a reason, but it is also outdated in notation and terminology, and is much harder to read than more contemporary sources. My favorite reference is Emily Riehl’s *Category Theory in Context*.

However, I believe strongly that category theory is learned significantly better by *doing*, rather than reading. From this perspective, I have two recommendations. Appendix A is a crash course in category theory which largely consists of exercises that I recommend to the students taking my classes. In addition, Tom Leinster has a great set of four exercise sheets from a course he taught in Cambridge in 2004. If it is still possible to find these on the internet at the moment of reading this sentence, I strongly recommend them.

11.3 Algebraic K-theory

Algebraic K-theory is not at all like topological K-theory, except when it is. It is not, in the sense that it is not computable, it is not simple to define, and is it extremely difficult to work with in every possibly conceivable sense of the word. It is, in the sense that it works with modules over a ring analogously to the way that we work with vector bundles over a base: by classifying them up to addition, and seeing what happens.

For a ring R, define $K_0(R)$ to be the free abelian group generated by finitely generated projective R-modules, modulo the relation that for every short exact sequence $A \hookrightarrow B \twoheadrightarrow C$, $[B] = [A] + [C]$. Since surjections between projective R-modules split, this is equivalent to the relation that $[A \oplus B] = [A] + [B]$. Thus the definition of $K_0(R)$ is explicitly analogous to that of $K^0(X)$ for a compact space X.

There are two important connections between algebraic K-theory and topological K-theory. The first of these is the Serre–Swan Theorem, which states that for a compact space X, the category of bundles on X is equivalent to the category of projective modules over the ring $C(X)$ of continuous real-valued functions on X. (For a complex version, replace "real-valued" with "complex-valued.") In particular, it means that $K^0(X)$ is isomorphic to $K_0(C(X))$. The second connection states that if X is an algebraic variety and we restrict to algebraic vector bundles, then algebraic K-theory of X (thought of as the correct "union" of the rings defining it) is isomorphic to topological K-theory of X. For discussions of these results (and far more!) see for example Max Karoubi's book *K-theory: An Introduction*, Charles Weibel's book *The K-book*, or the *Handbook of K-theory*.[b]

11.4 Chromatic homotopy theory

Topological K-theory is the lowest level of "chromatic homotopy theory," which splits the homotopy groups of a spectrum into different "colors" and analyzes each color independently. The development of chromatic homotopy theory started with an analysis of the image of the "J-homomorphism",

[b]Be warned: these are not small books, and they do not contain simple material.

which is a family of homomorphisms $J_i\colon \pi_i(O(n)) \longrightarrow \pi_{n+i}(S^n)$. The analysis of the image was done by Adams in a series of four papers (conveniently all called "On the groups $J(X)$"). It turns out to be a small part of a much larger story, involving formal group laws, algebraic geometry, and interesting periodic elements in the stable homotopy groups of spheres. The classic references are two books by Douglas Ravenel, *Nilpotence and Periodicity in Stable Homotopy Theory* and *Complex Cobordism and Stable Homotopy Groups of Spheres*.

Appendix A

A Crash Course on Category Theory

This chapter contains a crash course in category theory, designed to get a reader to the level of understanding necessary for the main book. As category theory is best understood through doing, rather than reading, the main content of the chapter is a series of exercises with the definitions. This is *not* intended as a comprehensive, thorough, or complete[a] introduction to the subject.

A.1 Categories and Functors

Definition A.1. A *category* $\mathcal{C}$ consists of a collection[b] $\mathrm{ob}\,\mathcal{C}$ of *objects* of $\mathcal{C}$, together with a set $\mathcal{C}(A, B)$ for all $A, B \in \mathrm{ob}\,\mathcal{C}$. (When it is clear from context we often write $A \in \mathcal{C}$ instead of $A \in \mathrm{ob}\,\mathcal{C}$.) The set $\mathcal{C}(A, B)$ is the *set of morphisms*, or Hom-*set* from A to B. It is sometimes denoted $\mathrm{Hom}(A, B)$ or $\mathrm{Hom}_{\mathcal{C}}(A, B)$. A morphism $f \in \mathcal{C}(A, B)$ is often denoted $f \colon A \longrightarrow B$.

For every object $A \in \mathrm{ob}\,\mathcal{C}$ there is a distinguished morphism $1_A \in \mathcal{C}(A, A)$ called the *identity*. For every triple of objects $A, B, C \in \mathrm{ob}\,\mathcal{C}$ there is a *composition*

$$\circ_{\mathcal{C}} \colon \mathcal{C}(B, C) \times \mathcal{C}(A, B) \longrightarrow \mathcal{C}(A, C).$$

When $\mathcal{C}$ is clear from context it is omitted from the notation. This compo-

[a] or any other word suggesting sufficient information

[b] A collection is a set, or larger than a set; for example, there is a collection of all sets, although not a set of all sets.

sition is *associative*, in the sense that given

$$f \in \mathcal{C}(A, B) \qquad g \in \mathcal{C}(B, C) \qquad h \in \mathcal{C}(C, D)$$

it is the case that

$$h \circ (g \circ f) = (h \circ g) \circ f.$$

It is also *unital* in the sense that for any $A, B \in \operatorname{ob} \mathcal{C}$ and any $f \in \mathcal{C}(A, B)$ it is the case that

$$1_B \circ f = f = f \circ 1_A.$$

A morphism $f\colon A \longrightarrow B$ is called an *isomorphism* if there exists a morphism $g\colon B \longrightarrow A$ such that $g \circ f = 1_A$ and $f \circ g = 1_B$. In this case g is called the *inverse to* f.

A fairly important example of constructions with categories is that of the opposite category.

Definition A.2. For a category $\mathcal{C}$, the *opposite category*, denoted $\mathcal{C}^{\mathrm{op}}$, is the category with $\operatorname{ob} \mathcal{C}^{\mathrm{op}} \stackrel{\mathrm{def}}{=} \operatorname{ob} \mathcal{C}$ and $\mathcal{C}^{\mathrm{op}}(A, B) \stackrel{\mathrm{def}}{=} \mathcal{C}(B, A)$. Composition is defined by

$$g \circ_{\mathcal{C}^{\mathrm{op}}} f \stackrel{\mathrm{def}}{=} f \circ_{\mathcal{C}} g,$$

and identities are the same ones as the ones in $\mathcal{C}$.

Morphisms between categories are called functors; they act on both objects and morphisms, in a compatible fashion.

Definition A.3. Let $\mathcal{C}$ and $\mathcal{D}$ be categories. A *functor* (often just referred to as a *functor*) $F\colon \mathcal{C} \longrightarrow \mathcal{D}$ is a function (also denoted F by an abuse of notation) $F\colon \operatorname{ob} \mathcal{C} \longrightarrow \operatorname{ob} \mathcal{D}$ and a collection of functions (also denoted F) $F\colon \mathcal{C}(A, B) \longrightarrow \mathcal{D}(F(A), F(B))$. These functions *respect composition*, in the sense that given $A, B, C \in \operatorname{ob} \mathcal{C}$ and $f \in \mathcal{C}(A, B)$ and $g \in \mathcal{C}(B, C)$ it is the case that

$$F(g \circ_{\mathcal{C}} f) = F(g) \circ_{\mathcal{D}} F(f).$$

In addition, it must be the case that for all $A \in \operatorname{ob} \mathcal{C}$,

$$F(1_A) = 1_{F(A)}.$$

A.1 (a) Consider the category $\mathcal{C}$ where

objects are nonnegative integers,

morphisms $\mathcal{C}(m, n) = \{n \times m$ matrices with $\mathbb{Q}$-entries$\}$, and

composition defined by matrix multiplication.

Check that this is, in fact, a category.

(b) Let $\mathbf{Vect}_{\mathbb{Q}}$ be the category of finite-dimensional $\mathbb{Q}$-vector spaces and linear maps. Write down functors $F\colon \mathcal{C} \longrightarrow \mathbf{Vect}_{\mathbb{Q}}$ and $G\colon \mathbf{Vect}_{\mathbb{Q}} \longrightarrow \mathcal{C}$ such that $G \circ F = 1_{\mathcal{C}}$. (Hint: you will need the axiom of choice to construct G.)

A.2 Prove that $\cdot^{\mathrm{op}}\colon \mathbf{Cat} \longrightarrow \mathbf{Cat}$ is a functor.

A.3 (a) Let P be a partially ordered set: a set equipped with a relation $\leq$ which is reflexive and transitive. Let NP be the set consisting of those finite subsets $\{p_0, \ldots, p_n\}$ of P where $p_0 < p_1 < \cdots < p_n$. Check that NP is a simplicial complex.

(b) Prove that N is a functor $\mathbf{PoSet} \longrightarrow \mathbf{SimpCplx}$. Here, $\mathbf{PoSet}$ has as morphisms functions of elements which preserve order, and $\mathbf{SimpCplx}$ has as morphisms the functions on vertices which take a simplex to a simplex.

(c) Let P^{op} be P with the opposite relation: if $x \leq y$ in P then $y \leq x$ in P^{op}. What is the relationship between $|NP|$ and $|NP^{op}|$?

(d) Let L be an ordered simplicial complex. Let cL be the poset whose elements are the simplices of L, and where $\sigma \leq \tau$ if $\sigma \subseteq \tau$. Prove that c is a functor $\mathbf{SimpCplx} \longrightarrow \mathbf{PoSet}$.

(e) Prove that there is a simplicial morphism $NcL \longrightarrow L$ given on vertices by $\sigma \longmapsto \max \sigma$.

(f) What is the relationship between L and NcL?

Instad of writing formulas it is often more useful to draw diagrams in which morphisms are represented by arrows. Suppose that we are given morphisms $f\colon A \longrightarrow B$, $g\colon B \longrightarrow D$, $h\colon A \longrightarrow C$ and $j\colon C \longrightarrow D$. We can represent these in a diagram

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

The diagram is said to *commute* if $gf = jh$.

Definition A.4. Let $F, G\colon \mathcal{C} \longrightarrow \mathcal{D}$ be two functors. A *natural transformation from F to G*, denoted $\alpha\colon F \Longrightarrow G$, is a choice of morphism $\alpha_A\colon F(A) \longrightarrow G(A)$ for every $A \in \operatorname{ob}\mathcal{C}$, satisfying the condition that for all pairs $A, B \in \operatorname{ob}\mathcal{C}$ and every $f \in \mathcal{C}(A, B)$, the square

$$\begin{array}{ccc} F(A) & \xrightarrow{F(f)} & F(B) \\ {\scriptstyle \alpha_A}\big\downarrow & & \big\downarrow{\scriptstyle \alpha_B} \\ G(A) & \xrightarrow{G(f)} & G(B) \end{array}$$

commutes.

A natural transformation α is a *natural isomorphism* if for all $A \in \operatorname{ob}\mathcal{C}$, α_A is an isomorphism.

Although defining isomorphisms of categories is fairly straightforward, the more natural notion turns out to be that of equivalence.

Definition A.5. A functor $F\colon \mathcal{C} \longrightarrow \mathcal{D}$ is an *equivalence of categories* if there exists a functor $G\colon \mathcal{D} \longrightarrow \mathcal{C}$ and natural isomorphisms $1_{\mathcal{C}} \Longrightarrow G \circ F$ and $F \circ G \Longrightarrow 1_{\mathcal{D}}$.

A.4 (a) Let G be a group. We can consider G to be a category by considering a category $\mathcal{C}$ with one object $*$ and with $\mathcal{C}(*, *) = G$. Composition is defined through the group operation. If G and H are two groups considered as categories, what are functors $F\colon G \longrightarrow H$? Given two functors $F, F'\colon G \longrightarrow H$, what is a natural transformation $\alpha\colon F \Longrightarrow F'$?

(b) Write down two functors $F, G\colon \mathbf{Sets} \longrightarrow \mathbf{Sets}$ such that there is a natural transformation $\alpha\colon F \Longrightarrow G$ but there is no natural transformation $\beta\colon G \Longrightarrow F$.

(c) Let $\mathcal{I}$ be the category $0 \longrightarrow 1$. Suppose that we are given two functors $F, G\colon \mathcal{C} \longrightarrow \mathcal{D}$. Prove that there exists a natural transformation $\alpha\colon F \Longrightarrow G$ if and only if there exists a functor $H\colon \mathcal{C} \times \mathcal{I} \longrightarrow \mathcal{D}$ such that $H|_{\mathcal{C}\times 0} = F$ and $H|_{\mathcal{D}\times 1} = G$. What definition from class is the definition of equivalent categories analogous to?

(d) A category is called *small* if its collection of objects is a set.[c] For any small category $\mathcal{C}$, define the space $X_{\mathcal{C}}$ by the following

[c] Yes, sets can get very large, but the point is that categories can get even larger.

method. Take a point for every object in $\mathcal{C}$, and for each morphism $f\colon A \longrightarrow B$ in $\mathcal{C}$ attach an arc to A at one end and B at the other. Prove that if $\mathcal{C}$ and $\mathcal{D}$ are equivalent then the connected components of $X_{\mathcal{C}}$ are in bijection with the connected components of $X_{\mathcal{D}}$.

A.5 Let $F\colon \mathcal{C} \longrightarrow \mathcal{D}$ be a functor. We say that F is

full if for all A and B, the function $\mathrm{Hom}_{\mathcal{C}}(A, B) \longrightarrow \mathrm{Hom}_{\mathcal{D}}(F(A), F(B))$ is surjective,

faithful if for all A and B, the function $\mathrm{Hom}_{\mathcal{C}}(A, B) \longrightarrow \mathrm{Hom}_{\mathcal{D}}(F(A), F(B))$ is injective, and

essentially surjective if for every object $D \in \mathcal{D}$ there exists an object $C \in \mathcal{C}$ and an isomorphism $F(C) \longrightarrow D$.

Prove that F is an equivalence of categories if and only if it is full, faithful and essentially surjective.

Definition A.6. Let $F\colon \mathcal{C} \longrightarrow \mathcal{D}$ and $G\colon \mathcal{D} \longrightarrow \mathcal{C}$. We say that F is *left adjoint to G* or G is *right adjoint to F* if for all $A \in \mathcal{C}$ and $B \in \mathcal{D}$,

$$\mathrm{Hom}_{\mathcal{D}}(FA, B) \cong \mathrm{Hom}_{\mathcal{C}}(A, GB).$$

This isomorphism needs to be natural in both A and B. In other words, this means that given any commutative diagram

$$\begin{array}{ccc} FA & \xrightarrow{\alpha} & B \\ {\scriptstyle Ff}\downarrow & & \downarrow{\scriptstyle g} \\ FA' & \xrightarrow{\beta} & B' \end{array}$$

there exists a commutative diagram

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

and vice versa.

A.6 (a) Let $\mathcal{C} = \mathcal{D} = \mathbf{Set}$, and fix a set S. Let $F(A) = A \times S$ and $G(B) = B^S$. Prove that F is left adjoint to G.

(b) Suppose that $\mathcal{C}$ is such that for all objects A and B, $|\mathrm{Hom}_{\mathcal{C}}(A, B)| \leq 1$. (Such a $\mathcal{C}$ is sometimes called a *preorder*.) Let $\mathcal{D}$ be the category with one object and no non-identity morphisms. Prove that $F\colon \mathcal{C} \longrightarrow \mathcal{D}$ has a left adjoint if and only if $\mathcal{C}$ has a minimal object and a right adjoint if and only if $\mathcal{C}$ has a maximal object.

(c) Suppose that F is left adjoint to G. Prove that there exists a natural transformation $\eta\colon 1_{\mathcal{C}} \Longrightarrow GF$ and a natural transformation $\epsilon\colon FG \Longrightarrow 1_{\mathcal{D}}$ such that for every $A \in \mathcal{C}$ and $B \in \mathcal{D}$ the following commute:

$$
\begin{array}{ccc}
GB & \xrightarrow{\eta_{GB}} & GFGB \\
& \searrow^{=} & \downarrow{\scriptstyle G(\epsilon_B)} \\
& & GB
\end{array}
\qquad\qquad
\begin{array}{ccc}
FA & \xrightarrow{F(\eta_A)} & FGFA \\
& \searrow^{=} & \downarrow{\scriptstyle \epsilon_{FA}} \\
& & FA
\end{array}
$$

(d) Now prove the converse of the above: if η and ϵ exist, show that F and G are adjoints.

(e) Recall how to construct the nerve of a partial order. We can consider a partial order P to be a category $\mathcal{C}_P$ with the objects the elements of the partial order and a unique morphism $x \longrightarrow y$ if $x \leq y$. Suppose that $f\colon S \longrightarrow T$ and $g\colon T \longrightarrow S$ are two order-preserving functions of partial orders such that the induced functors $F\colon \mathcal{C}_S \longrightarrow \mathcal{C}_T$ and $G\colon \mathcal{C}_T \longrightarrow \mathcal{C}_S$ are adjoints. Prove that NS and NT are homotopy equivalent. (Hint: See problem 6 on homework 2.)

Definition A.7. An *initial object* in a category $\mathcal{C}$ is an object $\varnothing$ such that for every other object $A \in \mathcal{C}$, $|\mathrm{Hom}(\varnothing, A)| = 1$.

A *terminal object* in $\mathcal{C}$ is an initial object in $\mathcal{C}^{\mathrm{op}}$.

Definition A.8. Let I be a category. A *diagram of shape* I in $\mathcal{C}$ is a functor $X\colon I \longrightarrow \mathcal{C}$. If X is a diagram of shape I, a *cone over* X is a natural transformation from a constant functor $I \longrightarrow \mathcal{C}$ to X. The *limit of* X, denoted $\lim X$, is the terminal object in the category of cones over X.

A *colimit* of a functor $F\colon \mathcal{C} \longrightarrow \mathcal{D}$ is the limit of the induced functor $F\colon \mathcal{C}^{\mathrm{op}} \longrightarrow \mathcal{D}^{\mathrm{op}}$.

A.7 (a) Give an example of a category with an initial object. Give an example of a category without an initial object.

(b) Give an example of a category with an object which is both initial and terminal. Give an example of a category with an an object I which is initial but such that $|\mathrm{Hom}(A, I)| > 1$ for all $A \ncong I$.

(c) Let $\mathcal{C}$ be a category, and consider the diagram

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

in $\mathcal{C}$. Let $\mathcal{C}_{//ACB}$ be the category with

objects commutative diagrams

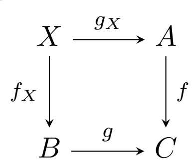

which we denote for simplicity as (X, f_X, g_X).

morphisms $(X, f_X, g_X) \longrightarrow (Y, f_Y, g_Y)$ are morphisms $\varphi \colon X \longrightarrow Y$ such that the diagram

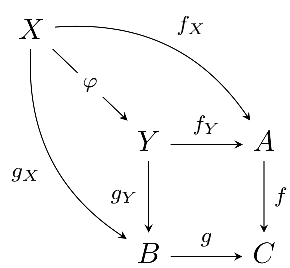

commutes.

Prove that for a fixed diagram $A \longrightarrow C \longleftarrow B$ in **Set**, the category $\mathbf{Set}_{//ACB}$ has a terminal object. This object is called the *pullback* of the diagram.

(d) Let $\mathcal{D}$ be the category

$$\bullet \longrightarrow \bullet \longleftarrow \bullet .$$

Let $[\mathcal{D}, \mathcal{C}]$ be the category of functors $\mathcal{D} \longrightarrow \mathcal{C}$ and natural transformations between them. There is a functor $c \colon \mathcal{C} \longrightarrow [\mathcal{D}, \mathcal{C}]$ given by sending each object of $\mathcal{C}$ to a constant diagram. Suppose that $\mathcal{C}$ has all pullbacks. Prove that c has a right adjoint.

A.8 (a) Let I be the category $\bullet \longrightarrow \bullet \longleftarrow \bullet$. Check that the limit of X is the pullback of X.

(b) Let $I = \bullet \rightrightarrows \bullet$ and let $X: I \longrightarrow \mathbf{Set}$ be a diagram. What is the limit of X?

(c) Let I be the empty diagram. Check that the limit of I is the terminal object.

(d) Let $F: \mathcal{C} \rightleftarrows \mathcal{D} :G$ be an adjoint pair, and let I be arbitrary. Let $X: I \longrightarrow \mathcal{D}$ be a diagram in $\mathcal{D}$. Prove that

$$G(\lim X) \cong \lim G \circ X.$$

(e) Use this to show that the underlying set of a pullback in **Top** or **Gp** is the pullback of the underlying sets.

A.9 A category $\mathcal{C}$ is *filtered* if it satisfies the following two conditions:

- For every pair of objects $A, B \in \mathcal{C}$, there exists an object C and morphisms $A \longrightarrow C$ and $B \longrightarrow C$.
- For every pair of morphisms $f, g: A \longrightarrow B \in \mathcal{C}$, there exists a morphism $h: B \longrightarrow C$ such that $hf = hg$.

(a) Prove that a product of filtered categories is filtered.

(b) Let $\mathcal{C}$ be a filtered category, and let $F: \mathcal{C} \longrightarrow \mathbf{Gp}$ be a diagram in $\mathcal{C}$ in which the image of every morphism in $\mathcal{C}$ is an injection. Prove that $\operatorname{colim} F = \bigcup_{A \in \mathcal{C}} F(A)$.[d]

(c) Prove that a filtered colimit of exact sequences is exact.

(d) Suppose that $\mathcal{D}$ is a full filtered subcategory of $\mathcal{C}$ such that for every object $A \in \mathcal{C}$ there exists an object $B \in \mathcal{D}$ and a morphism $A \longrightarrow B$ in $\mathcal{C}$. If we write $i: \mathcal{D} \longrightarrow \mathcal{C}$ for the inclusion and we let $F: \mathcal{C} \longrightarrow \mathbf{Set}$ be any functor, prove that the induced morphism on colimits

$$\operatorname*{colim}_{\mathcal{D}} Fi \longrightarrow \operatorname*{colim}_{\mathcal{C}} F$$

is an isomorphism.

[d] Such a colimit is called a *filtered colimit*.

Bibliography

[AA66] J. F. Adams and M. F. Atiyah. K-theory and the Hopf invariant. *Quart. J. Math. Oxford Ser. (2)*, 17:31–38, 1966.

[AB64] Michael Atiyah and Raoul Bott. On the periodicity theorem for complex vector bundles. *Acta Math.*, 112:229–247, 1964.

[Ada60] J. F. Adams. On the non-existence of elements of hopf invariant one. *Annals of Mathematics*, 72(1):20–104, 1960.

[Ada62] J. F. Adams. Vector fields on spheres. *Ann. of Math. (2)*, 75:603–632, 1962.

[Ada74] J.F. Adams. *Stable Homotopy and Generalised Homology*. Chicago Lectures in Mathematics. University of Chicago Press, 1974.

[AP99] M. A. Aguilar and Carlos Prieto. Quasifibrations and Bott periodicity. *Topology Appl.*, 98(1-3):3–17, 1999. II Iberoamerican Conference on Topology and its Applications (Morelia, 1997).

[Are46] Richard Arens. Topologies for homeomorphism groups. *Amer. J. Math.*, 68:593–610, 1946.

[Ati66] M. F. Atiyah. K-theory and reality. *Quart. J. Math. Oxford Ser. (2)*, 17:367–386, 1966.

[Ati89] M. F. Atiyah. *K-theory*. Advanced Book Classics. Addison-Wesley Publishing Company, Advanced Book Program, Redwood City, CA, second edition, 1989. Notes by D. W. Anderson.

[Beh02] Mark J. Behrens. A new proof of the Bott periodicity theorem. *Topology Appl.*, 119(2):167–183, 2002.

[BT13] R. Bott and L.W. Tu. *Differential Forms in Algebraic Topology*. Graduate Texts in Mathematics. Springer New York, 2013.

[Coc62] W. H. Cockcroft. On the thom isomorphism theorem. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58(2):206–208, 1962.

[DL61] E. Dyer and R. Lashof. A topological proof of the Bott periodicity theorems. *Ann. Mat. Pura Appl. (4)*, 54:231–254, 1961.

[Eng89] Ryszard Engelking. *General topology*, volume 6 of *Sigma Series in Pure Mathematics*. Heldermann Verlag, Berlin, second edition, 1989. Translated from the Polish by the author.

[FG05] Eric M. Friedlander and Daniel R. Grayson, editors. *Handbook of K-theory. Vol. 1, 2*. Springer-Verlag, Berlin, 2005.

[FP90] Rudolf Fritsch and Renzo Piccinini. *Cellular Structures in Topology*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1990.

[GM80] Mark Goresky and Robert MacPherson. Intersection homology theory. *Topology*, 19(2):135–162, 1980.

[HatA] Allen Hatcher. Spectral sequences in algebraic topology. https://www.math.cornell.edu/ hatcher/AT/ATch5.pdf.

[HatB] Allen Hatcher. Vector bundles and K-theory. https://www.math.cornell.edu/ hatcher/VBKT/VBpage.html.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[Hov99] Mark Hovey. *Model categories*, volume 63 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1999.

[Hus94] Dale Husemoller. *Fibre bundles*, volume 20 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 1994.

[Jam57] I. M. James. Multiplication on spheres. II. *Trans. Amer. Math. Soc.*, 84:545–558, 1957.

[Kar08A] M. Karoubi. *K-Theory: An Introduction*. Classics in Mathematics. Springer, 2008.

[Kar08B] Max Karoubi. *K-theory.* Classics in Mathematics. Springer-Verlag, Berlin, 2008. An introduction, Reprint of the 1978 edition, With a new postface by the author and a list of errata.

[KL72] S. L. Kleiman and Dan Laksov. Schubert calculus. *Amer. Math. Monthly*, 79:1061–1082, 1972.

[KT06] Akira Kōno and Dai Tamaki. *Generalized cohomology*, volume 230. American Mathematical Soc., 2006.

[Mac15] I. G. Macdonald. *Symmetric functions and Hall polynomials.* Oxford Classic Texts in the Physical Sciences. The Clarendon Press, Oxford University Press, New York, second edition, 2015. With contribution by A. V. Zelevinsky and a foreword by Richard Stanley, Reprint of the 2008 paperback edition [MR1354144].

[May99] J.P. May. *A Concise Course in Algebraic Topology.* Chicago Lectures in Mathematics. University of Chicago Press, 1999.

[McC01] John McCleary. *A user's guide to spectral sequences*, volume 58 of *Cambridge Studies in Advanced Mathematics.* Cambridge University Press, Cambridge, second edition, 2001.

[Mil63] J.W. Milnor. *Morse Theory.* Annals of mathematics studies. Princeton University Press, 1963.

[Mil65] John Milnor. *Lectures on the h-cobordism theorem.* Notes by L. Siebenmann and J. Sondow. Princeton University Press, Princeton, N.J., 1965.

[ML98] Saunders Mac Lane. *Categories for the working mathematician*, volume 5 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, second edition, 1998.

[MS74] John W. Milnor and James D. Stasheff. *Characteristic classes.* Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1974. Annals of Mathematics Studies, No. 76.

[MT08] R.E. Mosher and M.C. Tangora. *Cohomology Operations and Applications in Homotopy Theory.* Dover Books on Mathematics Series. Dover Publications, 2008.

[NT07] S. P. Novikov and I. A. Taimanov, editors. *Topological library. Part 1: cobordisms and their applications*, volume 39 of *Series on*

Knots and Everything. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2007. Translation by V. O. Manturov.

[Ran02] A. Ranicki. *Algebraic and Geometric Surgery*. Oxford Mathematical Monographs. Clarendon Press, 2002.

[Rav92] D.C. Ravenel. *Nilpotence and Periodicity in Stable Homotopy Theory*. Annals of Mathematics Studies. Princeton University Press, 1992.

[Rav03] D.C. Ravenel. *Complex Cobordism and Stable Homotopy Groups of Spheres*. AMS Chelsea Publishing Series. AMS Chelsea Pub., 2003.

[Rie14] Emily Riehl. *Categorical homotopy theory*, volume 24 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2014.

[Rie17] E. Riehl. *Category Theory in Context*. Aurora: Dover Modern Math Originals. Dover Publications, 2017.

[Sch] Stefan Schwede. Symmetric spectra. http://www.math.uni-bonn.de/people/schwede/SymSpec-v3.pdf.

[Swi17] R.M. Switzer. *Algebraic Topology - Homotopy and Homology*. Classics in Mathematics. Springer Berlin Heidelberg, 2017.

[Tho54] René Thom. Quelques propriétés globales des variétés différentiables. *Comment. Math. Helv.*, 28:17–86, 1954.

[Wei13] Charles A. Weibel. *The K-book*, volume 145 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2013. An introduction to algebraic K-theory.

[Whi44] Hassler Whitney. The singularities of a smooth n-manifold in $(2n-1)$-space. *Ann. of Math. (2)*, 45:247–293, 1944.

[Yok56] Ichiro Yokota. On the cellular decompositions of unitary groups. *J. Inst. Polytech. Osaka City Univ. Ser. A.*, 7:39–49, 1956.

[Yok57] Ichiro Yokota. On the homology of classical Lie groups. *J. Inst. Polytech. Osaka City Univ. Ser. A*, 8:93–120, 1957.