Vector Bundles and K-Theory

Version 2.0, January 2003

Allen Hatcher

Table of Contents

Preface

Topological K–theory, the first generalized cohomology theory to be studied thoroughly, was introduced around 1960 by Atiyah and Hirzebruch, based on the Periodicity Theorem of Bott proved just a few years earlier. In some respects K–theory is more elementary than classical homology and cohomology, and it is also more powerful for certain purposes. Some of the best-known applications of algebraic topology in the twentieth century, such as the theorem of Bott and Milnor that there are no division algebras after the Cayley octonions, or Adams' theorem determining the maximum number of linearly independent tangent vector fields on a sphere of arbitrary dimension, have relatively elementary proofs using K–theory, much simpler than the original proofs using ordinary homology and cohomology.

The first portion of this book takes these theorems as its goals, with an exposition that should be accessible to bright undergraduates familiar with standard material in linear algebra, abstract algebra, and point-set topology. Later chapters of the book assume more, approximately the contents of a standard graduate course in algebraic topology. A concrete goal of the later chapters is to tell the full story on the stable J–homomorphism, which gives the first level of depth in the stable homotopy groups of spheres. Along the way various other topics related to vector bundles that are of interest independent of K–theory are also developed, such as the characteristic classes associated to the names Stiefel and Whitney, Chern, and Pontryagin.

Introduction

Everyone is familiar with the Möbius band, the twisted product of a circle and a line, as contrasted with an annulus which is the actual product of a circle and a line. Vector bundles are the natural generalization of the Möbius band and annulus, with the circle replaced by an arbitrary topological space, called the base space of the vector bundle, and the line replaced by a vector space of arbitrary finite dimension, called the fiber of the vector bundle. Vector bundles thus combine topology with linear algebra, and the study of vector bundles could be called Linear Algebraic Topology.

The only two vector bundles with base space a circle and one-dimensional fiber are the Möbius band and the annulus, but the classification of all the different vector bundles over a given base space with fiber of a given dimension is quite difficult in general. For example, when the base space is a high-dimensional sphere and the dimension of the fiber is at least three, then the classification is of the same order of difficulty as the fundamental but still largely unsolved problem of computing the homotopy groups of spheres.

In the absence of a full classification of all the different vector bundles over a given base space, there are two directions one can take to make some partial progress on the problem. One can either look for invariants to distinguish at least some of the different vector bundles, or one can look for a cruder classification, using a weaker equivalence relation than the natural notion of isomorphism for vector bundles. As it happens, the latter approach is more elementary in terms of prerequisites, so let us discuss this first.

There is a natural direct sum operation for vector bundles over a fixed base space X, which in each fiber reduces just to direct sum of vector spaces. Using this, one can obtain a weaker notion of isomorphism of vector bundles by defining two vector bundles over the same base space X to be stably isomorphic if they become isomorphic after direct sum with product vector bundles $X\times\mathbb{R}^n$ for some n, perhaps different n's for the two given vector bundles. Then it turns out that the set of stable isomorphism classes of vector bundles over X forms an abelian group under the direct sum operation, at least if X is compact Hausdorff. The traditional notation for this group is $\widetilde{KO}(X)$. In the case of spheres the groups $\widetilde{KO}(S^n)$ have the quite unexpected property of being periodic in n. This is called Bott Periodicity, and the values of $\widetilde{KO}(S^n)$ are given by the following table:

n mod 8	1	2	3	4	5	6	7	8
$\widetilde{KO}(S^n)$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	0	$\mathbb{Z}$	0	0	0	$\mathbb{Z}$

For example, $\widetilde{KO}(S^1)$ is $\mathbb{Z}_2$, a cyclic group of order two, and a generator for this group is the Möbius bundle. This has order two since the direct sum of two copies of the

Möbius bundle is the product $S^1 \times \mathbb{R}^2$, as one can see by embedding two Möbius bands in a solid torus so that they intersect orthogonally along the common core circle of both bands, which is also the core circle of the solid torus.

Things become simpler if one passes from the real vector spaces to complex vector spaces. The complex version of $\widetilde{KO}(X)$, called $\widetilde{K}(X)$, is constructed in the same way as $\widetilde{KO}(X)$ but using vector bundles whose fibers are vector spaces over $\mathbb{C}$ rather than $\mathbb{R}$. The complex form of Bott Periodicity asserts simply that $\widetilde{K}(S^n)$ is $\mathbb{Z}$ for n even and 0 for n odd, so the period is two rather than eight.

The groups $\widetilde{K}(X)$ and $\widetilde{KO}(X)$ for varying X share certain formal properties with the cohomology groups studied in classical algebraic topology. Using a more general form of Bott periodicity, it is in fact possible to extend the groups $\widetilde{K}(X)$ and $\widetilde{KO}(X)$ to a full cohomology theory, families of abelian groups $\widetilde{K}^n(X)$ and $\widetilde{KO}^n(X)$ for $n \in \mathbb{Z}$ that are periodic in n of period two and eight, respectively. There is more algebraic structure here than just the additive group structure, however. Tensor products of vector spaces give rise to tensor products of vector bundles, which in turn give product operations in both real and complex K–theory similar to cup product in ordinary cohomology. Furthermore, exterior powers of vector spaces give natural operations within K–theory.

With all this extra structure, K–theory becomes a powerful tool, in some ways more powerful even than ordinary cohomology. The prime example of this is the very simple proof, once the basic machinery of complex K–theory has been set up, of the theorem that there are no finite dimensional division algebras over $\mathbb{R}$ in dimensions other than 1, 2, 4, and 8, the dimensions of the classical examples of the real and complex numbers, the quaternions, and the Cayley octonions. The same proof shows also that the only spheres whose tangent bundles are product bundles are S^1, S^3, and S^7, the unit spheres in the complex numbers, quaternions, and octonions.

Another classical problem that can be solved more easily using K–theory than ordinary cohomology is to find the maximum number of linearly independent tangent vector fields on the sphere S^n. In this case complex K–theory is not enough, and the added subtlety of real K–theory is needed. There is an algebraic construction of the requisite number of vector fields using Clifford algebras, and the harder part is to show there can be no more than this construction provides. Clifford algebras also provide a nice explanation for the mysterious sequence of groups appearing the in real form of Bott periodicity.

Now let us return to the original classification problem for vector bundles over a given base space and the question of finding invariants to distinguish different vector bundles. The first such invariant is orientability, the question of whether all the fibers can be coherently oriented. For example, the Möbius bundle is not orientable since as one goes all the way around the base circle, the orientation of the fiber lines is reversed. This does not happen for the annulus, which is an orientable vector bundle.

Orientability is measured by the first of a sequence of cohomology classes associated to a vector bundle, called Stiefel-Whitney classes. The next Stiefel-Whitney class measures a more refined sort of orientability called a spin structure, and the higher Stiefel-Whitney classes measure whether the vector bundle looks more and more like a product vector bundle over succesively higher dimensional subspaces of the base space. Cohomological invariants of vector bundles such as these with nice general properties are known as characteristic classes. It turns out that Stiefel-Whitney classes generate all characteristic classes for ordinary cohomology with $\mathbb{Z}_2$ coefficients, but with $\mathbb{Z}$ coefficients there are others, called Pontryagin and Euler classes, the latter being related to the Euler characteristic. Although characteristic classes do not come close to distinguishing all the different vector bundles over a given base space, except in a few low dimensional cases, they have still proved themselves to be quite useful objects.

Chapter 1

Vector Bundles

To motivate the definition of a vector bundle let us consider tangent vectors to the unit 2-sphere S^2 in $\mathbb{R}^3$. At each point $x \in S^2$ there is a tangent plane P_x. This is a 2-dimensional vector space with the point x as its zero vector 0_x. Vectors $v_x \in P_x$ are thought of as arrows with their tail at x. If we regard a vector v_x in P_x as a vector in $\mathbb{R}^3$, then the standard convention in linear algebra would be to identify v_x with all its parallel translates, and in particular with the unique translate $\tau(v_x)$ having its tail at the origin in $\mathbb{R}^3$. The association $v_x \mapsto \tau(v_x)$ defines a function $\tau : TS^2 \to \mathbb{R}^3$ where TS^2 is the set of all tangent vectors v_x as x ranges over S^2. This function τ is surjective but certainly not injective, as every nonzero vector in $\mathbb{R}^3$ occurs as $\tau(v_x)$ for infinitely many x, in fact for all x in a great circle in S^2. Moreover $\tau(0_x) = 0$ for all $x \in S^2$, so $\tau^{-1}(0)$ is a whole sphere. On the other hand, the function $TS^2 \to S^2 \times \mathbb{R}^3$, $v_x \mapsto (x, \tau(v_x))$, is injective, and can be used to topologize TS^2 as a subspace of $S^2 \times \mathbb{R}^3$, namely the subspace consisting of pairs (x, v) with v orthogonal to x.

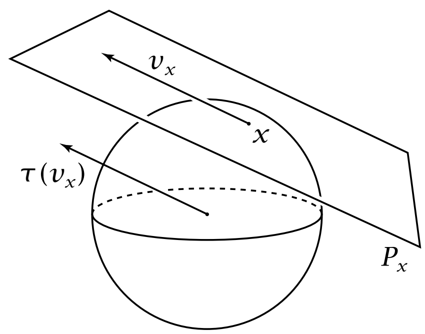

Thus TS^2 is first of all a topological space, and secondly it is the disjoint union of all the vector spaces P_x for $x \in S^2$. One can think of TS^2 as a continuous family of vector spaces parametrized by points of S^2.

The simplest continuous family of 2-dimensional vector spaces parametrized by points of S^2 is of course the product $S^2 \times \mathbb{R}^2$. Is that what TS^2 really is? More precisely we can ask whether there is a homeomorphism $h : TS^2 \to S^2 \times \mathbb{R}^2$ that takes each plane P_x to the plane $\{x\} \times \mathbb{R}^2$ by a vector space isomorphism. If we had such an h, then for each fixed nonzero vector $v \in \mathbb{R}^2$ the family of vectors $v_x = h^{-1}(x, v)$ would be a continuous field of nonzero tangent vectors to S^2. It is a classical theorem in algebraic topology that no such vector field exists. (See §2.2 for a proof using techniques from this book.) So TS^2 is genuinely twisted, and is not just a disguised form of the product $S^2 \times \mathbb{R}^2$.

Dropping down a dimension, one could consider in similar fashion the space TS^1 of tangent vectors to the unit circle S^1 in $\mathbb{R}^2$. In this case there is a continuous field v_x of nonzero tangent vectors to S^1, obtained by regarding points $x \in S^1$ as unit complex numbers and letting v_x be the translation of the vector ix that has its tail at x. This leads to a homeomorphism $S^1 \times \mathbb{R} \to TS^1$ taking (x, t) to tv_x, with $\{x\} \times \mathbb{R}$ going to the tangent line at x by a linear isomorphism. Thus TS^1 really is equivalent to the product $S^1 \times \mathbb{R}^1$.

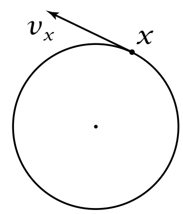

Moving up to S^3, the unit sphere in $\mathbb{R}^4$, the space TS^3 of tangent vectors is again equivalent to the product $S^3 \times \mathbb{R}^3$. Regarding $\mathbb{R}^4$ as the quaternions, an equivalence is the homeomorphism $S^3 \times \mathbb{R}^3 \to TS^3$ sending $(x, (t_1, t_2, t_3))$ to the translation of the vector $t_1 ix + t_2 jx + t_3 kx$ having its tail at x. A similar construction using Cayley octonions shows that TS^7 is equivalent to $S^7 \times \mathbb{R}^7$. It is a rather deep theorem, proved in §2.3, that S^1, S^3, and S^7 are the only spheres whose tangent bundle is equivalent to a product.

Although TS^n is not usually equivalent to the product $S^n \times \mathbb{R}^n$, there is a sense in which this is true locally. Take the case of the 2-sphere for example. For a point $x \in S^2$ let P be the translate of the tangent plane P_x that passes through the origin. For points $y \in S^2$ that are sufficiently close to x the map $\pi_y : P_y \to P$ sending a tangent vector v_y to the orthogonal projection of $\tau(v_y)$ onto P is a linear isomorphism. This is true in fact for any y on the same side of P as x. Thus for y in a suitable neighborhood U of x in S^2 the map $(y, v_y) \mapsto (y, \pi_y(v_y))$ is a homeomorphism with domain the subspace of TS^2 consisting of tangent vectors at points of U and with range the product $U \times P$. Furthermore this homeomorphism has the key property of restricting to a linear isomorphism from P_y onto P for each $y \in U$. A convenient way of rephrasing this situation, having the virtue of easily generalizing, is to let $p : TS^2 \to S^2$ be the map $(x, v_x) \mapsto x$, and then we have a homeomorphism $p^{-1}(U) \to U \times P$ that restricts to a linear isomorphism $p^{-1}(y) \to \{y\} \times P$ for each $y \in U$.

1.1 Basic Definitions and Constructions

Throughout the book we use the word map to mean a continuous function.

An *n-dimensional vector bundle* is a map $p:E\to B$ together with a real vector space structure on $p^{-1}(b)$ for each $b\in B$, such that the following local triviality condition is satisfied: There is a cover of B by open sets U_α for each of which there exists a homeomorphism $h_\alpha:p^{-1}(U_\alpha)\to U_\alpha\times\mathbb{R}^n$ taking $p^{-1}(b)$ to $\{b\}\times\mathbb{R}^n$ by a vector space isomorphism for each $b\in U_\alpha$. Such an h_α is called a *local trivialization* of the vector bundle. The space B is called the *base space*, E is the *total space*, and the vector spaces $p^{-1}(b)$ are the *fibers*. Often one abbreviates terminology by just calling the vector bundle E, letting the rest of the data be implicit.

We could equally well take $\mathbb{C}$ in place of $\mathbb{R}$ as the scalar field, obtaining the notion of a *complex vector bundle*. We will focus on real vector bundles in this chapter. Usually the complex case is entirely analogous. In the next chapter complex vector bundles will play the larger role, however.

Here are some examples of vector bundles:

(1) The *product* or *trivial* bundle $E=B\times\mathbb{R}^n$ with p the projection onto the first factor.

(2) If we let E be the quotient space of $I\times\mathbb{R}$ under the identifications $(0,t)\sim(1,-t)$, then the projection $I\times\mathbb{R}\to I$ induces a map $p:E\to S^1$ which is a 1-dimensional vector bundle, or *line bundle*. Since E is homeomorphic to a Möbius band with its boundary circle deleted, we call this bundle the *Möbius bundle*.

(3) The tangent bundle of the unit sphere S^n in $\mathbb{R}^{n+1}$, a vector bundle $p:E\to S^n$ where $E=\{\,(x,v)\in S^n\times\mathbb{R}^{n+1}\mid x\perp v\,\}$ and we think of v as a tangent vector to S^n by translating it so that its tail is at the head of x, on S^n. The map $p:E\to S^n$ sends (x,v) to x. To construct local trivializations, choose any point $x\in S^n$ and let $U_x\subset S^n$ be the open hemisphere containing x and bounded by the hyperplane through the origin orthogonal to x. Define $h_x:p^{-1}(U_x)\to U_x\times p^{-1}(x)\approx U_x\times\mathbb{R}^n$ by $h_x(y,v)=(y,\pi_x(v))$ where π_x is orthogonal projection onto the hyperplane $p^{-1}(x)$. Then h_x is a local trivialization since π_x restricts to an isomorphism of $p^{-1}(y)$ onto $p^{-1}(x)$ for each $y\in U_b$.

(4) The normal bundle to S^n in $\mathbb{R}^{n+1}$, a line bundle $p:E\to S^n$ with E consisting of pairs $(x,v)\in S^n\times\mathbb{R}^{n+1}$ such that v is perpendicular to the tangent plane to S^n at x, or in other words, $v=tx$ for some $t\in\mathbb{R}$. The map $p:E\to S^n$ is again given by $p(x,v)=x$. As in the previous example, local trivializations $h_x:p^{-1}(U_x)\to U_x\times\mathbb{R}$ can be obtained by orthogonal projection of the fibers $p^{-1}(y)$ onto $p^{-1}(x)$ for $y\in U_x$.

(5) Real projective n-space $\mathbb{R}\mathrm{P}^n$ is the space of lines in $\mathbb{R}^{n+1}$ through the origin. Since each such line intersects the unit sphere S^n in a pair of antipodal points, we

can also regard $\mathbb{R}P^n$ as the quotient space of S^n in which antipodal pairs of points are identified.. The *canonical line bundle* $p:E\to\mathbb{R}P^n$ has as its total space E the subspace of $\mathbb{R}P^n\times\mathbb{R}^{n+1}$ consisting of pairs (ℓ,v) with $v\in\ell$, and $p(\ell,v)=\ell$. Again local trivializations can be defined by orthogonal projection.

There is also an infinite-dimensional projective space $\mathbb{R}P^\infty$ which is the union of the finite-dimensional projective spaces $\mathbb{R}P^n$ under the inclusions $RP^n\subset RP^{n+1}$ coming from the natural inclusions $\mathbb{R}^{n+1}\subset\mathbb{R}^{n+2}$. The topology we use on $\mathbb{R}P^\infty$ is the weak or direct limit topology, for which a set in $\mathbb{R}P^\infty$ is open iff it intersects each $\mathbb{R}P^n$ in an open set. The inclusions $RP^n\subset\mathbb{R}P^{n+1}$ induce corresponding inclusions of canonical line bundles, and the union of all these is a canonical line bundle over $\mathbb{R}P^\infty$, again with the direct limit topology. Local trivializations work just as in the finite-dimensional case.

(6) The canonical line bundle over $\mathbb{R}P^n$ has an orthogonal complement, the space $E^\perp=\{(\ell,v)\in\mathbb{R}P^n\times\mathbb{R}^{n+1}\mid v\perp\ell\}$. The projection $p:E^\perp\to\mathbb{R}P^n$, $p(\ell,v)=\ell$, is a vector bundle with fibers the orthogonal subspaces $\ell^\perp$, of dimension n. Local trivializations can be obtained once more by orthogonal projection.

A natural generalization of $\mathbb{R}P^n$ is the so-called Grassmann manifold $G_k(\mathbb{R}^n)$, the space of all k-dimenional planes through the origin in $\mathbb{R}^n$. The topology on this space will be defined precisely in §1.2, along with a canonical k-dimensional vector bundle over it consisting of pairs (ℓ,v) where ℓ is a point in $G_k(\mathbb{R}^n)$ and v is a vector in ℓ. This too has an orthogonal complement, an $(n-k)$-dimensional vector bundle consisting of pairs (ℓ,v) with v orthogonal to ℓ.

An *isomorphism* between vector bundles $p_1:E_1\to B$ and $p_2:E_2\to B$ over the same base space B is a homeomorphism $h:E_1\to E_2$ taking each fiber $p_1^{-1}(b)$ to the corresponding fiber $p_2^{-1}(b)$ by a linear isomorphism. Thus an isomorphism preserves all the structure of a vector bundle, so isomorphic bundles are often regarded as the same. We use the notation $E_1\approx E_2$ to indicate that E_1 and E_2 are isomorphic.

For example, the normal bundle of S^n in $\mathbb{R}^{n+1}$ is isomorphic to the product bundle $S^n\times\mathbb{R}$ by the map $(x,tx)\mapsto(x,t)$. The tangent bundle to S^1 is also isomorphic to the trivial bundle $S^1\times\mathbb{R}$, via $(e^{i\theta},ite^{i\theta})\mapsto(e^{i\theta},t)$, for $e^{i\theta}\in S^1$ and $t\in\mathbb{R}$.

As a further example, the Möbius bundle in (2) above is isomorphic to the canonical line bundle over $\mathbb{R}P^1\approx S^1$. Namely, $\mathbb{R}P^1$ is swept out by a line rotating through an angle of π, so the vectors in these lines sweep out a rectangle $[0,\pi]\times\mathbb{R}$ with the two ends $\{0\}\times\mathbb{R}$ and $\{\pi\}\times\mathbb{R}$ identified. The identification is $(0,x)\sim(\pi,-x)$ since rotating a vector through an angle of π produces its negative.

One can sometimes distinguish nonisomorphic bundles by looking at the complement of the zero section since any vector bundle isomorphism $h:E_1\to E_2$ must take the zero section of E_1 onto the zero section of E_2, hence the complements of the zero sections in E_1 and E_2 must be homeomorphic. For example, the Möbius bundle is not isomorphic to the product bundle $S^1 \times \mathbb{R}$ since the complement of the zero section in the Möbius bundle is connected while for the product bundle the complement of the zero section is not connected. This method for distinguishing vector bundles can also be used with more refined topological invariants such as homology groups.

Sections

A *section* of a vector bundle $p:E \to B$ is a map $s:B \to E$ assigning to each $b \in B$ a vector $s(b)$ in the fiber $p^{-1}(b)$. The condition $s(b) \in p^{-1}(b)$ can also be written as $ps = \mathbb{1}$, the identity map of B. Every vector bundle has a canonical section, the *zero section* whose value is the zero vector in each fiber. We often identity the zero section with its image, a subspace of E which projects homeomorphically onto B by p.

One can sometimes distinguish nonisomorphic bundles by looking at the complement of the zero section since any vector bundle isomorphism $h:E_1 \to E_2$ must take the zero section of E_1 onto the zero section of E_2, so the complements of the zero sections in E_1 and E_2 must be homeomorphic. For example, we can see that the Möbius bundle is not isomorphic to the product bundle $S^1 \times \mathbb{R}$ since the complement of the zero section is connected for the Möbius bundle but not for the product bundle.

At the other extreme from the zero section would be a section whose values are all nonzero. Not all vector bundles have such a section. Consider for example the tangent bundle to S^n. Here a section is just a tangent vector field to S^n. As we shall show in §2.2, S^n has a nonvanishing vector field iff n is odd. From this it follows that the tangent bundle of S^n is not isomorphic to the trivial bundle if n is even and nonzero, since the trivial bundle obviously has a nonvanishing section, and an isomorphism between vector bundles takes nonvanishing sections to nonvanishing sections.

In fact, an n-dimensional bundle $p:E \to B$ is isomorphic to the trivial bundle iff it has n sections $s_1, \cdots, s_n$ such that the vectors $s_1(b), \cdots, s_n(b)$ are linearly independent in each fiber $p^{-1}(b)$. In one direction this is evident since the trivial bundle certainly has such sections and an isomorphism of vector bundles takes linearly independent sections to linearly independent sections. Conversely, if one has n linearly independent sections s_i, the map $h:B \times \mathbb{R}^n \to E$ given by $h(b, t_1, \cdots, t_n) = \sum_i t_i s_i(b)$ is a linear isomorphism in each fiber, and is continuous since its composition with a local trivialization $p^{-1}(U) \to U \times \mathbb{R}^n$ is continuous. Hence h is an isomorphism by the following useful technical result:

Lemma 1.1. *A continuous map $h:E_1 \to E_2$ between vector bundles over the same base space B is an isomorphism if it takes each fiber $p_1^{-1}(b)$ to the corresponding fiber $p_2^{-1}(b)$ by a linear isomorphism.*

Proof: The hypothesis implies that h is one-to-one and onto. What must be checked is that h^{-1} is continuous. This is a local question, so we may restrict to an open set

$U \subset B$ over which E_1 and E_2 are trivial. Composing with local trivializations reduces to the case of an isomorphism $h: U\times\mathbb{R}^n \to U\times\mathbb{R}^n$ of the form $h(x,v) = (x, g_x(v))$. Here g_x is an element of the group $GL_n(\mathbb{R})$ of invertible linear transformations of $\mathbb{R}^n$, and g_x depends continuously on x. This means that if g_x is regarded as an $n\times n$ matrix, its n^2 entries depend continuously on x. The inverse matrix g_x^{-1} also depends continuously on x since its entries can be expressed algebraically in terms of the entries of g_x, namely, g_x^{-1} is $1/(\det g_x)$ times the classical adjoint matrix of g_x. Therefore $h^{-1}(x,v) = (x, g_x^{-1}(v))$ is continuous. □

As an example, the tangent bundle to S^1 is trivial because it has the section $(x_1,x_2) \mapsto (-x_2,x_1)$ for $(x_1,x_2) \in S^1$. In terms of complex numbers, if we set $z = x_1 + ix_2$ then this section is $z \mapsto iz$ since $iz = -x_2 + ix_1$.

There is an analogous construction using quaternions instead of complex numbers. Quaternions have the form $z = x_1 + ix_2 + jx_3 + kx_4$, and form a division algebra $\mathbb{H}$ via the multiplication rules $i^2 = j^2 = k^2 = -1$, $ij = k$, $jk = i$, $ki = j$, $ji = -k$, $kj = -i$, and $ik = -j$. If we identify $\mathbb{H}$ with $\mathbb{R}^4$ via the coordinates (x_1,x_2,x_3,x_4), then the unit sphere is S^3 and we can define three sections of its tangent bundle by the formulas

$$z \mapsto iz \quad \text{or} \quad (x_1,x_2,x_3,x_4) \mapsto (-x_2,x_1,-x_4,x_3)$$
$$z \mapsto jz \quad \text{or} \quad (x_1,x_2,x_3,x_4) \mapsto (-x_3,x_4,x_1,-x_2)$$
$$z \mapsto kz \quad \text{or} \quad (x_1,x_2,x_3,x_4) \mapsto (-x_4,-x_3,x_2,x_1)$$

It is easy to check that the three vectors in the last column are orthogonal to each other and to (x_1,x_2,x_3,x_4), so we have three linearly independent nonvanishing tangent vector fields on S^3, and hence the tangent bundle to S^3 is trivial.

The underlying reason why this works is that quaternion multiplication satisfies $|zw| = |z||w|$, where $|\cdot|$ is the usual norm of vectors in $\mathbb{R}^4$. Thus multiplication by a quaternion in the unit sphere S^3 is an isometry of $\mathbb{H}$. The quaternions $1,i,j,k$ form the standard orthonormal basis for $\mathbb{R}^4$, so when we multiply them by an arbitrary unit quaternion $z \in S^3$ we get a new orthonormal basis z, iz, jz, kz.

The same constructions work for the Cayley octonions, a division algebra structure on $\mathbb{R}^8$. Thinking of $\mathbb{R}^8$ as $\mathbb{H}\times\mathbb{H}$, multiplication of octonions is defined by $(z_1,z_2)(w_1,w_2) = (z_1w_1 - \overline{w}_2z_2, z_2\overline{w}_1 + w_2z_1)$ and satisfies the key property $|zw| = |z||w|$. This leads to the construction of seven orthogonal tangent vector fields on the unit sphere S^7, so the tangent bundle to S^7 is also trivial. As we shall show in §2.3, the only spheres with trivial tangent bundle are S^1, S^3, and S^7.

Another way of characterizing the trivial bundle $E \approx B\times\mathbb{R}^n$ is to say that there is a continuous projection map $E \to \mathbb{R}^n$ which is a linear isomorphism on each fiber, since such a projection together with the bundle projection $E \to B$ gives an isomorphism $E \approx B\times\mathbb{R}^n$, by Lemma 1.1.

Direct Sums

Given two vector bundles $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ over the same base space B, we would like to create a third vector bundle over B whose fiber over each point of B is the direct sum of the fibers of E_1 and E_2 over this point. This leads us to define the *direct sum* of E_1 and E_2 as the space

$$E_1 \oplus E_2 = \{ (v_1, v_2) \in E_1 \times E_2 \mid p_1(v_1) = p_2(v_2) \}$$

There is then a projection $E_1 \oplus E_2 \to B$ sending (v_1, v_2) to the point $p_1(v_1) = p_2(v_2)$. The fibers of this projection are the direct sums of the fibers of E_1 and E_2, as we wanted. For a relatively painless verification of the local triviality condition we make two preliminary observations:

(a) Given a vector bundle $p : E \to B$ and a subspace $A \subset B$, then $p : p^{-1}(A) \to A$ is clearly a vector bundle. We call this the *restriction of E over A*.

(b) Given vector bundles $p_1 : E_1 \to B_1$ and $p_2 : E_2 \to B_2$, then $p_1 \times p_2 : E_1 \times E_2 \to B_1 \times B_2$ is also a vector bundle, with fibers the products $p_1^{-1}(b_1) \times p_2^{-1}(b_2)$. For if we have local trivializations $h_\alpha : p_1^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n$ and $h_\beta : p_2^{-1}(U_\beta) \to U_\beta \times \mathbb{R}^m$ for E_1 and E_2, then $h_\alpha \times h_\beta$ is a local trivialization for $E_1 \times E_2$.

Then if E_1 and E_2 both have the same base space B, the restriction of the product $E_1 \times E_2$ over the diagonal $B = \{(b,b) \in B \times B\}$ is exactly $E_1 \oplus E_2$.

The direct sum of two trivial bundles is again a trivial bundle, clearly, but the direct sum of nontrivial bundles can also be trivial. For example, the direct sum of the tangent and normal bundles to S^n in $\mathbb{R}^{n+1}$ is the trivial bundle $S^n \times \mathbb{R}^{n+1}$ since elements of the direct sum are triples $(x, v, tx) \in S^n \times \mathbb{R}^{n+1} \times \mathbb{R}^{n+1}$ with $x \perp v$, and the map $(x, v, tx) \mapsto (x, v + tx)$ gives an isomorphism of the direct sum bundle with $S^n \times \mathbb{R}^{n+1}$. So the tangent bundle to S^n is *stably trivial*: it becomes trivial after taking the direct sum with a trivial bundle.

As another example, the direct sum $E \oplus E^\perp$ of the canonical line bundle $E \to \mathbb{R}P^n$ with its orthogonal complement, defined in example (6) above, is isomorphic to the trivial bundle $\mathbb{R}P^n \times \mathbb{R}^{n+1}$ via the map $(\ell, v, w) \mapsto (\ell, v + w)$ for $v \in \ell$ and $w \perp \ell$. Specializing to the case $n = 1$, the bundle $E^\perp$ is isomorphic to E itself by the map that rotates each vector in the plane by 90 degrees. We noted earlier that E is isomorphic to the Möbius bundle over $S^1 = \mathbb{R}P^1$, so it follows that the direct sum of the Möbius bundle with itself is the trivial bundle. To see this geometrically, embed the Möbius bundle in the product bundle $S^1 \times \mathbb{R}^2$ by taking the line in the fiber $\{\theta\} \times \mathbb{R}^2$ that makes an angle of $\theta/2$ with the x-axis, and then the orthogonal lines in the fibers form a second copy of the Möbius bundle, giving a decomposition of the product $S^1 \times \mathbb{R}^2$ as the direct sum of two Möbius bundles.

Example: The tangent bundle of real projective space. Starting with the isomorphism $S^n \times \mathbb{R}^{n+1} \approx TS^n \oplus NS^n$, where NS^n is the normal bundle of S^n in $\mathbb{R}^{n+1}$, suppose we factor out by the identifications $(x,v) \sim (-x,-v)$ on both sides of this isomorphism. Applied to TS^n this identification yields $T\mathbb{R}\mathrm{P}^n$, the tangent bundle to $\mathbb{R}\mathrm{P}^n$. This is saying that a tangent vector to $\mathbb{R}\mathrm{P}^n$ is equivalent to a pair of antipodal tangent vectors to S^n. A moment's reflection shows this to be entirely reasonable, although a formal proof would require a significant digression on what precisely tangent vectors to a smooth manifold are, a digression we shall skip here. What we will show is that even though the direct sum of $T\mathbb{R}\mathrm{P}^n$ with a trivial line bundle may not be trivial as it is for a sphere, it does split in an interesting way as a direct sum of nontrivial line bundles.

In the normal bundle NS^n the identification $(x,v) \sim (-x,-v)$ can be written as $(x,tx) \sim (-x,t(-x))$. This identification yields the product bundle $\mathbb{R}\mathrm{P}^n \times \mathbb{R}$ since the section $x \mapsto (-x,-x)$ is well-defined in the quotient. Now let us consider the identification $(x,v) \sim (-x,-v)$ in $S^n \times \mathbb{R}^{n+1}$. This identification respects the coordinate factors of $\mathbb{R}^{n+1}$, so the quotient is the direct sum of $n+1$ copies of the line bundle E over $\mathbb{R}\mathrm{P}^n$ obtained by making the identifications $(x,t) \sim (-x,-t)$ in $S^n \times \mathbb{R}$. The claim is that E is just the canonical line bundle over $\mathbb{R}\mathrm{P}^n$. To see this, let us identify $S^n \times \mathbb{R}$ with NS^n by the isomorphism $(x,t) \mapsto (x,tx)$. hence $(-x,-t) \mapsto ((-x,(-t)(-x)) = (-x,tx)$. Thus we have the identification $(x,tx) \sim (-x,tx)$ in NS^n. The quotient is the canonical line bundle over $\mathbb{R}\mathrm{P}^n$ since the identifications $x \sim -x$ in the first coordinate give lines through the origin in $\mathbb{R}^{n+1}$, and in the second coordinate there are no identifications so we have well-defined vectors tx in these lines.

Thus we have shown that the tangent bundle $T\mathbb{R}\mathrm{P}^n$ is stably isomorphic to the direct sum of $n+1$ copies of the canonical line bundle over $\mathbb{R}\mathrm{P}^n$. When $n=3$, for example, $T\mathbb{R}\mathrm{P}^3$ is trivial since the three linearly independent tangent vector fields on S^3 defined earlier in terms of quaternions pass down to linearly independent tangent vector fields on the quotient $\mathbb{R}\mathrm{P}^3$. Hence the direct sum of four copies of the canonical line bundle over $\mathbb{R}\mathrm{P}^3$ is trivial. Similarly using octonions we can see that the direct sum of eight copies of the canonical line bundle over $\mathbb{R}\mathrm{P}^7$ is trivial. In §2.5 we will determine when the sum of k copies of the canonical line bundle over $\mathbb{R}\mathrm{P}^n$ is at least stably trivial. The answer turns out to be rather subtle: This happens exactly when k is a multiple of $2^{\varphi(n)}$ where $\varphi(n)$ is the number of integers i in the range $0 < i \le n$ with i congruent to 0, 1, 2, or 4 modulo 8. For $n = 3,7$ this gives $2^{\varphi(n)} = 4,8$, the same numbers we obtained from the triviality of $T\mathbb{R}\mathrm{P}^3$ and $T\mathbb{R}\mathrm{P}^7$. If there were a 16-dimensional division algebra after the octonions then one might expect the sum of 16 copies of the canonical line bundle over $\mathbb{R}\mathrm{P}^{15}$ to be trivial. However this is not the case since $2^{\varphi(15)} = 2^7 = 128$.

Inner Products

An *inner product* on a vector bundle $p:E\to B$ is a map $\langle\,,\rangle:E\oplus E\to\mathbb{R}$ which restricts in each fiber to an inner product, a positive definite symmetric bilinear form.

Proposition 1.2. *An inner product exists for a vector bundle $p:E\to B$ if B is compact Hausdorff or more generally paracompact.*

The definition of paracompactness we are using is that a space X is paracompact if it is Hausdorff and every open cover has a subordinate partition of unity, a collection of maps $\varphi_\beta:X\to[0,1]$ each supported in some set of the open cover, and with $\sum_\beta \varphi_\beta = 1$, only finitely many of the φ_β being nonzero near each point of X. Constructing such functions is easy when X is compact Hausdorff, using Urysohn's Lemma. This is done in the appendix to this chapter, where we also show that certain classes of noncompact spaces are paracompact. Most spaces that arise naturally in algebraic topology are paracompact.

Proof: An inner product for $p:E\to B$ can be constructed by first using local trivializations $h_\alpha:p^{-1}(U_\alpha)\to U_\alpha\times\mathbb{R}^n$, to pull back the standard inner product in $\mathbb{R}^n$ to an inner product $\langle\cdot,\cdot\rangle_\alpha$ on $p^{-1}(U_\alpha)$, then setting $\langle v,w\rangle = \sum_\beta \varphi_\beta p(v)\langle v,w\rangle_{\alpha(\beta)}$ where $\{\varphi_\beta\}$ is a partition of unity with the support of φ_β contained in $U_{\alpha(\beta)}$. $\square$

In the case of complex vector bundles one can construct Hermitian inner products in the same way.

In linear algebra one can show that a vector subspace is always a direct summand by taking its orthogonal complement. We will show now that the corresponding result holds for vector bundles over a paracompact base. A *vector subbundle* of a vector bundle $p:E\to B$ has the natural definition: a subspace $E_0\subset E$ intersecting each fiber of E in a vector subspace, such that the restriction $p:E_0\to B$ is a vector bundle.

Proposition 1.3. *If $E\to B$ is a vector bundle over a paracompact base B and $E_0\subset E$ is a vector subbundle, then there is a vector subbundle $E_0^\perp\subset E$ such that $E_0\oplus E_0^\perp\approx E$.*

Proof: With respect to a chosen inner product on E, let $E_0^\perp$ be the subspace of E which in each fiber consists of all vectors orthogonal to vectors in E_0. We claim that the natural projection $E_0^\perp\to B$ is a vector bundle. If this is so, then $E_0\oplus E_0^\perp$ is isomorphic to E via the map $(v,w)\mapsto v+w$, using Lemma 1.1.

To see that $E_0^\perp$ satisfies the local triviality condition for a vector bundle, note first that we may assume E is the product $B\times\mathbb{R}^n$ since the question is local in B. Since E_0 is a vector bundle, of dimension m say, it has m independent local sections $b\mapsto(b,s_i(b))$ near each point $b_0\in B$. We may enlarge this set of m independent local sections of E_0 to a set of n independent local sections $b\mapsto(b,s_i(b))$ of E by choosing $s_{m+1},\cdots,s_n$ first in the fiber $p^{-1}(b_0)$, then taking the same vectors for all nearby fibers, since if $s_1,\cdots,s_m,s_{m+1},\cdots,s_n$ are independent at b_0, they will remain independent for nearby b by continuity of the determinant function. Apply the Gram-Schmidt orthogonalization process to $s_1, \cdots, s_m, s_{m+1}, \cdots, s_n$ in each fiber, using the given inner product, to obtain new sections s_i'. The explicit formulas for the Gram-Schmidt process show the s_i''s are continuous, and the first m of them are a basis for E_0 in each fiber. The sections s_i' allow us to define a local trivialization $h:p^{-1}(U)\to U\times\mathbb{R}^n$ with $h(b,s_i'(b))$ equal to the i^{th} standard basis vector of $\mathbb{R}^n$. This h carries E_0 to $U\times\mathbb{R}^m$ and $E_0^{\perp}$ to $U\times\mathbb{R}^{n-m}$, so $h|E_0^{\perp}$ is a local trivialization of $E_0^{\perp}$. □

Note that when the subbundle E_0 is equal to E itself, the last part of the proof shows that for any vector bundle with an inner product it is always possible to choose local trivializations that carry the inner product to the standard inner product, so the local trivializations are by isometries.

We have seen several cases where the sum of two bundles, one or both of which may be nontrivial, is the trivial bundle. Here is a general result result along these lines:

Proposition 1.4. *For each vector bundle $E\to B$ with B compact Hausdorff there exists a vector bundle $E'\to B$ such that $E\oplus E'$ is the trivial bundle.*

This can fail when B is noncompact. An example is the canonical line bundle over $\mathbb{R}\mathrm{P}^\infty$, as we shall see in Example 3.6.

Proof: To motivate the construction, suppose first that the result holds and hence that E is a subbundle of a trivial bundle $B\times\mathbb{R}^N$. Composing the inclusion of E into this product with the projection of the product onto $\mathbb{R}^N$ yields a map $E\to\mathbb{R}^N$ that is a linear injection on each fiber. Our strategy will be to reverse the logic here, first constructing a map $E\to\mathbb{R}^N$ that is a linear injection on each fiber, then showing that this gives an embedding of E in $B\times\mathbb{R}^N$ as a direct summand.

Each point $x\in B$ has a neighborhood U_x over which E is trivial. By Urysohn's Lemma there is a map $\varphi_x:B\to[0,1]$ that is 0 outside U_x and nonzero at x. Letting x vary, the sets $\varphi_x^{-1}(0,1]$ form an open cover of B. By compactness this has a finite subcover. Let the corresponding U_x's and φ_x's be relabeled U_i and φ_i. Define $g_i:E\to\mathbb{R}^n$ by $g_i(v)=\varphi_i(p(v))[\pi_i h_i(v)]$ where p is the projection $E\to B$ and $\pi_i h_i$ is the composition of a local trivialization $h_i:p^{-1}(U_i)\to U_i\times\mathbb{R}^n$ with the projection π_i to $\mathbb{R}^n$. Then g_i is a linear injection on each fiber over $\varphi_i^{-1}(0,1]$, so if we make the various g_i's the coordinates of a map $g:E\to\mathbb{R}^N$ with $\mathbb{R}^N$ a product of copies of $\mathbb{R}^n$, then g is a linear injection on each fiber.

The map g is the second coordinate of a map $f:E\to B\times\mathbb{R}^N$ with first coordinate p. The image of f is a subbundle of the product $B\times\mathbb{R}^N$ since projection of $\mathbb{R}^N$ onto the i^{th} $\mathbb{R}^n$ factor gives the second coordinate of a local trivialization over $\varphi_i^{-1}(0,1]$. Thus we have E isomorphic to a subbundle of $B\times\mathbb{R}^N$ so by preceding proposition there is a complementary subbundle E' with $E\oplus E'$ isomorphic to $B\times\mathbb{R}^N$. □

Tensor Products

In addition to direct sum, a number of other algebraic constructions with vector spaces can be extended to vector bundles. One which is particularly important for K–theory is tensor product. For vector bundles $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$, let $E_1 \otimes E_2$, as a set, be the disjoint union of the vector spaces $p_1^{-1}(x) \otimes p_2^{-1}(x)$ for $x \in B$. The topology on this set is defined in the following way. Choose isomorphisms $h_i : p_i^{-1}(U) \to U \times \mathbb{R}^{n_i}$ for each open set $U \subset B$ over which E_1 and E_2 are trivial. Then a topology $\mathcal{T}_U$ on the set $p_1^{-1}(U) \otimes p_2^{-1}(U)$ is defined by letting the fiberwise tensor product map $h_1 \otimes h_2 : p_1^{-1}(U) \otimes p_2^{-1}(U) \to U \times (\mathbb{R}^{n_1} \otimes \mathbb{R}^{n_2})$ be a homeomorphism. The topology $\mathcal{T}_U$ is independent of the choice of the h_i's since any other choices are obtained by composing with isomorphisms of $U \times \mathbb{R}^{n_i}$ of the form $(x, v) \mapsto (x, g_i(x)(v))$ for continuous maps $g_i : U \to GL_{n_i}(\mathbb{R})$, hence $h_1 \otimes h_2$ changes by composing with analogous isomorphisms of $U \times (\mathbb{R}^{n_1} \otimes \mathbb{R}^{n_2})$ whose second coordinates $g_1 \otimes g_2$ are continuous maps $U \to GL_{n_1 n_2}(\mathbb{R})$, since the entries of the matrices $g_1(x) \otimes g_2(x)$ are the products of the entries of $g_1(x)$ and $g_2(x)$. When we replace U by an open subset V, the topology on $p_1^{-1}(V) \otimes p_2^{-1}(V)$ induced by $\mathcal{T}_U$ is the same as the topology $\mathcal{T}_V$ since local trivializations over U restrict to local trivializations over V. Hence we get a well-defined topology on $E_1 \otimes E_2$ making it a vector bundle over B.

There is another way to look at this construction that takes as its point of departure a general method for constructing vector bundles we have not mentioned previously. If we are given a vector bundle $p : E \to B$ and an open cover $\{U_\alpha\}$ of B with local trivializations $h_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n$, then we can reconstruct E as the quotient space of the disjoint union $\coprod_\alpha (U_\alpha \times \mathbb{R}^n)$ obtained by identifying $(x, v) \in U_\alpha \times \mathbb{R}^n$ with $h_\beta h_\alpha^{-1}(x, v) \in U_\beta \times \mathbb{R}^n$ whenever $x \in U_\alpha \cap U_\beta$. The functions $h_\beta h_\alpha^{-1}$ can be viewed as maps $g_{\beta\alpha} : U_\alpha \cap U_\beta \to GL_n(\mathbb{R})$. These satisfy the 'cocycle condition' $g_{\gamma\beta} g_{\beta\alpha} = g_{\gamma\alpha}$ on $U_\alpha \cap U_\beta \cap U_\gamma$. Any collection of 'gluing functions' $g_{\beta\alpha}$ satisfying this condition can be used to construct a vector bundle $E \to B$.

In the case of tensor products, suppose we have two vector bundles $E_1 \to B$ and $E_2 \to B$. We can choose an open cover $\{U_\alpha\}$ with both E_1 and E_2 trivial over each U_α, and so obtain gluing functions $g^i_{\beta\alpha} : U_\alpha \cap U_\beta \to GL_{n_i}(\mathbb{R})$ for each E_i. Then the gluing functions for the bundle $E_1 \otimes E_2$ are the tensor product functions $g^1_{\beta\alpha} \otimes g^2_{\beta\alpha}$ assigning to each $x \in U_\alpha \cap U_\beta$ the tensor product of the two matrices $g^1_{\beta\alpha}(x)$ and $g^2_{\beta\alpha}(x)$.

It is routine to verify that the tensor product operation for vector bundles over a fixed base space is commutative, associative, and has an identity element, the trivial line bundle. It is also distributive with respect to direct sum.

If we restrict attention to line bundles, then the set $\mathrm{Vect}^1(B)$ of isomorphism classes of one-dimensional vector bundles over B is an abelian group with respect to the tensor product operation. The inverse of a line bundle $E \to B$ is obtained by replacing its gluing matrices $g_{\beta\alpha}(x) \in GL_1(\mathbb{R})$ with their inverses. The cocycle condition is preserved since 1×1 matrices commute. If we give E an inner product, we may rescale local trivializations h_α to be isometries, taking vectors in fibers of E to vectors in $\mathbb{R}^1$ of the same length. Then all the values of the gluing functions $g_{\beta\alpha}$ are ± 1, being isometries of $\mathbb{R}$. The gluing functions for $E \otimes E$ are the squares of these $g_{\beta\alpha}$'s, hence are identically 1, so $E \otimes E$ is the trivial line bundle. Thus each element of $\mathrm{Vect}^1(B)$ is its own inverse. As we shall see in §3.1, the group $\mathrm{Vect}^1(B)$ is isomorphic to a standard object in algebraic topology, the cohomology group $H^1(B;\mathbb{Z}_2)$ when B is homotopy equivalent to a CW complex.

These tensor product constructions work equally well for complex vector bundles. Tensor product again makes the complex analog $\mathrm{Vect}^1_{\mathbb{C}}(B)$ of $\mathrm{Vect}^1(B)$ into an abelian group, but after rescaling the gluing functions $g_{\beta\alpha}$ for a complex line bundle E, the values are complex numbers of norm 1, not necessarily ± 1, so we cannot expect $E \otimes E$ to be trivial. In §3.1 we will show that the group $\mathrm{Vect}^1_{\mathbb{C}}(B)$ is isomorphic to $H^2(B;\mathbb{Z})$ when B is homotopy equivalent to a CW complex.

We may as well mention here another general construction for complex vector bundles $E \to B$, the notion of the *conjugate bundle* $\overline{E} \to B$. As a topological space, $\overline{E}$ is the same as E, but the vector space structure in the fibers is modified by redefining scalar multiplication by the rule $\lambda(v) = \overline{\lambda} v$ where the right side of this equation means scalar multiplication in E and the left side means scalar multiplication in $\overline{E}$. This implies that local trivializations for $\overline{E}$ are obtained from local trivializations for E by composing with the coordinatewise conjugation map $\mathbb{C}^n \to \mathbb{C}^n$ in each fiber. The effect on the gluing maps $g_{\beta\alpha}$ is to replace them by their complex conjugates as well. Specializing to line bundles, we then have $E \otimes \overline{E}$ isomorphic to the trivial line bundle since its gluing maps have values $z\overline{z} = 1$ for z a unit complex number. Thus conjugate bundles provide inverses in $\mathrm{Vect}^1_{\mathbb{C}}(B)$.

Besides tensor product of vector bundles, another construction useful in K–theory is the exterior power $\lambda^k(E)$ of a vector bundle E. Recall from linear algebra that the exterior power $\lambda^k(V)$ of a vector space V is the quotient of the k-fold tensor product $V \otimes \cdots \otimes V$ by the subspace generated by vectors of the form $v_1 \otimes \cdots \otimes v_k - \mathrm{sgn}(\sigma) v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(k)}$ where σ is a permutation of the subscripts and $\mathrm{sgn}(\sigma) = \pm 1$ is its sign, $+1$ for an even permutation and -1 for an odd permutation. If V has dimension n then $\lambda^k(V)$ has dimension $\binom{n}{k}$. Now to define $\lambda^k(E)$ for a vector bundle $p : E \to B$ the procedure follows closely what we did for tensor product. We first form the disjoint union of the exterior powers $\lambda^k(p^{-1}(x))$ of all the fibers $p^{-1}(x)$, then we define a topology on this set via local trivializations. The key fact about tensor product which we needed before was that the tensor product $\varphi \otimes \psi$ of linear transformations φ and ψ depends continuously on φ and ψ. For exterior powers the analogous fact is that a linear map $\varphi : \mathbb{R}^n \to \mathbb{R}^n$ induces a linear map $\lambda^k(\varphi) : \lambda^k(\mathbb{R}^n) \to \lambda^k(\mathbb{R}^n)$ which depends continuously on φ. This holds since $\lambda^k(\varphi)$ is a quotient map of the k-fold tensor product of φ with itself.

Associated Fiber Bundles

If we modify the definition of a vector bundle by dropping all references to vector spaces and replace the model fiber $\mathbb{R}^n$ by an arbitrary space F, then we have the more general notion of a *fiber bundle* with fiber F. This is a map $p:E\to B$ such that there is a cover of B by open sets U_α for each of which there exists a homeomorphism $h_\alpha:p^{-1}(U_\alpha)\to U_\alpha\times F$ taking $p^{-1}(b)$ to $\{b\}\times F$ for each $b\in U_\alpha$. We will describe now several different ways of constructing fiber bundles from vector bundles.

Having an inner product on a vector bundle E, lengths of vectors are defined, and we can consider the subspace $S(E)$ consisting of the unit spheres in all the fibers. The natural projection $S(E)\to B$ is a fiber bundle with sphere fibers since we have observed that local trivializations for E can be chosen to be isometries in each fiber, so these local trivializations restrict to local trivializations for $S(E)$. Similarly there is a disk bundle $D(E)\to B$ with fibers the disks of vectors of length less than or equal to 1. It is possible to describe $S(E)$ without reference to an inner product, as the quotient of the complement of the zero section in E obtained by identifying each nonzero vector with all positive scalar multiples of itself. It follows that $D(E)$ can also be defined without invoking a metric, namely as the mapping cylinder of the projection $S(E)\to B$. This is the quotient space of $S(E)\times[0,1]$ obtained by identifying two points in $S(E)\times\{0\}$ if they have the same image in B.

The canonical line bundle $E\to\mathbb{R}\mathrm{P}^n$ has as its unit sphere bundle $S(E)$ the space of unit vectors in lines through the origin in $\mathbb{R}^{n+1}$. Since each unit vector uniquely determines the line containing it, $S(E)$ is the same as the space of unit vectors in $\mathbb{R}^{n+1}$, i.e., S^n.

Here are some more examples.

(1) Associated to a vector bundle $E\to B$ is the *projective bundle* $P(E)\to B$, where $P(E)$ is the space of all lines through the origin in all the fibers of E. We topologize $P(E)$ as the quotient of the sphere bundle $S(E)$ obtained by factoring out scalar multiplication in each fiber. Over a neighborhood U in B where E is a product $U\times\mathbb{R}^n$, this quotient is $U\times\mathbb{R}\mathrm{P}^{n-1}$, so $P(E)$ is a fiber bundle over B with fiber $\mathbb{R}\mathrm{P}^{n-1}$, with respect to the projection $P(E)\to B$ which sends each line in the fiber of E over a point $b\in B$ to b.

(2) For an n-dimensional vector bundle $E\to B$, the associated *flag bundle* $F(E)\to B$ has total space $F(E)$ the subspace of the n-fold product of $P(E)$ with itself consisting of n-tuples of orthogonal lines in fibers of E. The fiber of $F(E)$ is thus the flag manifold $F(\mathbb{R}^n)$ consisting of n-tuples of orthogonal lines through the origin in $\mathbb{R}^n$. Local triviality follows as in the preceding example. More generally, for any $k\le n$ one could take k-tuples of orthogonal lines in fibers of E and get a bundle $F_k(E)\to B$.

(3) As a refinement of the last example, one could form the *Stiefel bundle* $V_k(E)\to B$, where points of $V_k(E)$ are k-tuples of orthogonal unit vectors in fibers of E, so $V_k(E)$

is a subspace of the product of k copies of $S(E)$. The fiber of $V_k(E)$ is the Stiefel manifold $V_k(\mathbb{R}^n)$ of orthonormal k-frames in $\mathbb{R}^n$.

(4) Generalizing $P(E)$, there is the *Grassmann bundle* $G_k(E) \to B$ of k-dimensional linear subspaces of fibers of E. This is the quotient space of $V_k(E)$ obtained by identifying two k-frames in a fiber if they span the same subspace of the fiber. The fiber of $G_k(E)$ is the Grassmann manifold $G_k(\mathbb{R}^n)$ of k-planes through the origin in $\mathbb{R}^n$.

Exercises

1. Show that a vector bundle $E \to X$ has k independent sections iff it has a trivial k-dimensional subbundle.

2. For a vector bundle $E \to X$ with a subbundle $E' \subset E$, construct a quotient vector bundle $E/E' \to X$.

3. Show that the orthogonal complement of a subbundle is independent of the choice of inner product, up to isomorphism.

1.2 Classifying Vector Bundles

As was stated in the Introduction, it is usually a difficult problem to classify all the different vector bundles over a given base space. The goal of this section will be to rephrase the problem in terms of a standard concept of algebraic topology, the idea of homotopy classes of maps. This will allow the problem to be solved directly in a few very simple cases. Using machinery of algebraic topology, other more difficult cases can be handled as well, as is explained in §??. The reformulation in terms of homotopy also offers some conceptual enlightment about the structure of vector bundles.

For the reader who is unfamiliar with the notion of homotopy we give the basic definitions in the Glossary [not yet written], and more details can be found in the introductory chapter of the author's book *Algebraic Topology*.

Pullback Bundles

We will denote the set of isomorphism classes of n-dimensional real vector bundles over B by $\mathrm{Vect}^n(B)$. The complex analogue, when we need it, will be denoted by $\mathrm{Vect}^n_{\mathbb{C}}(B)$. Our first task is to show how a map $f:A\to B$ gives rise to a function $f^*:\mathrm{Vect}(B)\to\mathrm{Vect}(A)$, in the reverse direction.

Proposition 1.5. *Given a map $f:A\to B$ and a vector bundle $p:E\to B$, then there exists a vector bundle $p':E'\to A$ with a map $f':E'\to E$ taking the fiber of E' over each point $a\in A$ isomorphically onto the fiber of E over $f(a)$, and such a vector bundle E' is unique up to isomorphism.*

From the uniqueness statement it follows that the isomorphism type of E' depends only on the isomorphism type of E since we can compose the map f' with an isomorphism of E with another vector bundle over B. Thus we have a function $f^*:\mathrm{Vect}(B)\to\mathrm{Vect}(A)$ taking the isomorphism class of E to the isomorphism class of E'. Often the vector bundle E' is written as $f^*(E)$ and called the bundle *induced* by f, or the *pullback* of E by f.

Proof: First we construct an explicit pullback by setting $E'=\{(a,v)\in A\times E \mid f(a)=p(v)\}$. This subspace of $A\times E$ fits into the diagram at the right where $p'(a,v)=a$ and $f'(a,v)=v$. Notice that the two compositions fp' and pf' from E' to B are equal since they both since a pair (a,v) to $f(a)$. The formula $fp'=pf'$ looks a bit like a commutativity relation, which may explain why the word *commutative* is used to describe a diagram like this one in which any two compositions of maps from one point in the diagram to another are equal.

$$\begin{array}{ccc} E' & \xrightarrow{\tilde{f}'} & E \\ \downarrow{\scriptstyle p'} & & \downarrow{\scriptstyle p} \\ A & \xrightarrow{f} & B \end{array}$$

If we let Γ_f denote the graph of f, all points $(a,f(a))$ in $A\times B$, then p' factors as the composition $E'\to\Gamma_f\to A$, $(a,v)\mapsto(a,p(v))=(a,f(a))\mapsto a$. The first of these two maps is the restriction of the vector bundle $\mathbb{1}\times p:A\times E\to A\times B$ over the graph Γ_f, so it is a vector bundle, and the second map is a homeomorphism, so their composition $p':E'\to A$ is a vector bundle. The map f' obviously takes the fiber E' over a isomorphically onto the fiber of E over $f(a)$.

For the uniqueness statement, we can construct an isomorphism from an arbitrary E' satisfying the conditions in the proposition to the particular one just constructed by sending $v'\in E'$ to the pair $(p'(v'),f'(v'))$. This map obviously takes each fiber of E' to the corresponding fiber of $f^*(E)$ by a vector space isomorphism, so by Lemma 1.1 it is an isomorphism of vector bundles. □

One can be more explicit about local trivializations in the pullback bundle $f^*(E)$ constructed above. If E is trivial over a subspace $U\subset B$ then $f^*(E)$ is trivial over $f^{-1}(U)$ since linearly independent sections s_i of E over U give rise to independent sections $a\mapsto (a,s_i(f(a)))$ of $f^*(E)$ over $f^{-1}(U)$. In particular, the pullback of a trivial bundle is a trivial bundle. This can also be seen directly from the definition, which in the case $E=B\times\mathbb{R}^n$ just says that $f^*(E)$ consists of the triples (a,b,v) in $A\times B\times\mathbb{R}^n$ with $b=f(a)$, so b is redundant and we have just the product $A\times\mathbb{R}^n$.

Now let us give some examples of pullbacks.

(1) The restriction of a vector bundle $p:E\to B$ over a subspace $A\subset B$ can be viewed as a pullback with respect to the inclusion map $A\hookrightarrow B$ since the inclusion $p^{-1}(A)\hookrightarrow E$ is certainly an isomorphism on each fiber.

(2) Another very special case is when the map f is a constant map, having image a single point $b\in B$. Then $f^*(E)$ is just the product $A\times p^{-1}(b)$, a trivial bundle.

(3) The tangent bundle TS^n is the pullback of the tangent bundle $T\mathbb{R}\mathrm{P}^n$ via the quotient map $S^n\to\mathbb{R}\mathrm{P}^n$. Indeed, $T\mathbb{R}\mathrm{P}^n$ was defined as a quotient space of TS^n and the quotient map takes fibers isomorphically to fibers.

(4) An interesting example which is small enough to be visualized completely is the pullback of the Möbius bundle $E\to S^1$ by the two-to-one map $f:S^1\to S^1$ given by $f(z)=z^2$ in complex notation. One can regard the Möbius bundle as the quotient of $S^1\times\mathbb{R}$ under the identifications $(z,t)\sim(-z,-t)$, and the quotient map for this identification is the map f' exhibiting the annulus $S^1\times\mathbb{R}$ as the pullback of the Möbius bundle. More concretely, the pullback bundle can be thought of as a coat of paint applied to 'both sides' of the Möbius bundle, and the map f' sends each molecule of paint to the point of the Möbius band to which it adheres. Since E has one half-twist, the pullback has two half-twists, hence is the trivial bundle. More generally, if E_n is the pullback of the Möbius bundle by the map $z\mapsto z^n$, then E_n is the trivial bundle for n even and the Möbius bundle for n odd. This can be seen by viewing the Möbius bundle as the quotient of a strip $[0,1]\times\mathbb{R}$ obtained by identifying the two edges $\{0\}\times\mathbb{R}$ and $\{1\}\times\mathbb{R}$ by a reflection of $\mathbb{R}$, and then the bundle E_n can be constructed from n such strips by identifying the right edge of the i^{th} strip to the

left edge of the $i+1^{st}$ strip by a reflection, the number i being taken modulo n so that the last strip is glued back to the first.

(5) At the end of the previous section we defined the flag bundle $F(E)$ associated to an n-dimensional vector bundle $E \to B$ to be the space of orthogonal direct sum decompositions of fibers of E into lines. The vectors in the i^{th} line form a line bundle $L_i \to F(E)$, and the direct sum $L_1 \oplus \cdots \oplus L_n$ is nothing but the pullback of E with respect to the projection $F(E) \to E$ since a point in the pullback consists of an n-tuple of lines $\ell_1 \perp \cdots \perp \ell_n$ in a fiber of E together with a vector v in this fiber, and v can be expressed uniquely as a sum $v = v_1 + \cdots + v_n$ with $v_i \in \ell_i$. This construction for pulling an arbitrary bundle back to a sum of line bundles is a key ingredient the so-called 'splitting principle' which is important in §2.3 and §3.1.

$$\begin{array}{ccc} L_1 \oplus \cdots \oplus L_n & \longrightarrow & E \\ \downarrow & & \downarrow \\ F(E) & \longrightarrow & B \end{array}$$

Here are some elementary properties of pullbacks:

(i) $(fg)^*(E) \approx g^*(f^*(E))$.
(ii) $\mathbb{1}^*(E) \approx E$.
(iii) $f^*(E_1 \oplus E_2) \approx f^*(E_1) \oplus f^*(E_2)$.
(iv) $f^*(E_1 \otimes E_2) \approx f^*(E_1) \otimes f^*(E_2)$.

The proofs are easy applications of the preceding proposition. In each case one just checks that the bundle on the right satisfies the characteristic property of a pullback. For example in (iv) there is a natural map from $f^*(E_1) \otimes f^*(E_2)$ to $E_1 \otimes E_2$ that is an isomorphism on each fiber, so $f^*(E_1) \otimes f^*(E_2)$ satisfies the condition for being the pullback $f^*(E_1 \otimes E_2)$.

Now we come to the make technical result about pullbacks:

Theorem 1.6. *Given a vector bundle $p : E \to B$ and homotopic maps $f_0, f_1 : A \to B$, then the induced bundles $f_0^*(E)$ and $f_1^*(E)$ are isomorphic if A is compact Hausdorff or more generally paracompact.*

Proof: Let $F : A \times I \to B$ be a homotopy from f_0 to f_1. The restrictions of $F^*(E)$ over $A \times \{0\}$ and $A \times \{1\}$ are then $f_0^*(E)$ and $f_1^*(E)$. So the theorem will be an immediate consequence of the following result:

Proposition 1.7. *The restrictions of a vector bundle $E \to X \times I$ over $X \times \{0\}$ and $X \times \{1\}$ are isomorphic if X is paracompact.*

Proof: We need two preliminary facts:

(1) A vector bundle $p : E \to X \times [a,b]$ is trivial if its restrictions over $X \times [a,c]$ and $X \times [c,b]$ are both trivial for some $c \in (a,b)$. To see this, let these restrictions be $E_1 = p^{-1}(X \times [a,c])$ and $E_2 = p^{-1}(X \times [c,b])$, and let $h_1 : E_1 \to X \times [a,c] \times \mathbb{R}^n$ and $h_2 : E_2 \to X \times [c,b] \times \mathbb{R}^n$ be isomorphisms. These isomorphisms may not agree on $p^{-1}(X \times \{c\})$, but they can be made to agree by replacing h_2 by its composition with

the isomorphism $X \times [c,b] \times \mathbb{R}^n \to X \times [c,b] \times \mathbb{R}^n$ which on each slice $X \times \{x\} \times \mathbb{R}^n$ is given by $h_1 h_2^{-1} : X \times \{c\} \times \mathbb{R}^n \to X \times \{c\} \times \mathbb{R}^n$. Once h_1 and h_2 agree on $E_1 \cap E_2$, they define a trivialization of E.

(2) For a vector bundle $p : E \to X \times I$, there exists an open cover $\{U_\alpha\}$ of X so that each restriction $p^{-1}(U_\alpha \times I) \to U_\alpha \times I$ is trivial. This is because for each $x \in X$ we can find open neighborhoods $U_{x,1}, \cdots, U_{x,k}$ in X and a partition $0 = t_0 < t_1 < \cdots < t_k = 1$ of $[0,1]$ such that the bundle is trivial over $U_{x,i} \times [t_{i-1}, t_i]$, using compactness of $[0,1]$. Then by (1) the bundle is trivial over $U_\alpha \times I$ where $U_\alpha = U_{x,1} \cap \cdots \cap U_{x,k}$.

Now we prove the proposition. By (2), we can choose an open cover $\{U_\alpha\}$ of X so that E is trivial over each $U_\alpha \times I$. Let us first deal with the simpler case that X is compact Hausdorff. Then a finite number of U_α's cover X. Relabel these as U_i for $i = 1, 2, \cdots, m$. As shown in Proposition 1.4 there is a corresponding partition of unity by functions φ_i with the support of φ_i contained in U_i. For $i \ge 0$, let $\psi_i = \varphi_1 + \cdots + \varphi_i$, so in particular $\psi_0 = 0$ and $\psi_m = 1$. Let X_i be the graph of ψ_i, the subspace of $X \times I$ consisting of points of the form $(x, \psi_i(x))$, and let $p_i : E_i \to X_i$ be the restriction of the bundle E over X_i. Since E is trivial over $U_i \times I$, the natural projection homeomorphism $X_i \to X_{i-1}$ lifts to a homeomorphism $h_i : E_i \to E_{i-1}$ which is the identity outside $p_i^{-1}(U_i)$ and which takes each fiber of E_i isomorphically onto the corresponding fiber of E_{i-1}. The composition $h = h_1 h_2 \cdots h_m$ is then an isomorphism from the restriction of E over $X \times \{1\}$ to the restriction over $X \times \{0\}$.

In the general case that X is only paracompact, Lemma 1.21 asserts that there is a countable cover $\{V_i\}_{i \ge 1}$ of X and a partition of unity $\{\varphi_i\}$ with φ_i supported in V_i, such that each V_i is a disjoint union of open sets each contained in some U_α. This means that E is trivial over each $V_i \times I$. As before we let $\psi_i = \varphi_1 + \cdots + \varphi_i$ and let $p_i : E_i \to X_i$ be the restriction of E over the graph of ψ_i. Also as before we construct $h_i : E_i \to E_{i-1}$ using the fact that E is trivial over $V_i \times I$. The infinite composition $h = h_1 h_2 \cdots$ is then a well-defined isomorphism from the restriction of E over $X \times \{1\}$ to the restriction over $X \times \{0\}$ since near each point $x \in X$ only finitely many φ_i's are nonzero, which implies that for large enough i, $h_i = \mathbb{1}$ over a neighborhood of x. $\square$

Corollary 1.8. *A homotopy equivalence $f : A \to B$ of paracompact spaces induces a bijection $f^* : \mathrm{Vect}^n(B) \to \mathrm{Vect}^n(A)$. In particular, every vector bundle over a contractible paracompact base is trivial.*

Proof: If g is a homotopy inverse of f then we have $f^* g^* = \mathbb{1}^* = \mathbb{1}$ and $g^* f^* = \mathbb{1}^* = \mathbb{1}$. $\square$

We might remark that Theorem 1.6 holds for fiber bundles as well as vector bundles, with the same proof.

Clutching Functions

Let us describe a way to construct vector bundles $E \to S^k$ with base space a sphere. Write S^k as the union of its upper and lower hemispheres D^k_+ and D^k_-, with $D^k_+ \cap D^k_- = S^{k-1}$. Given a map $f: S^{k-1} \to GL_n(\mathbb{R})$, let E_f be the quotient of the disjoint union $D^k_+ \times \mathbb{R}^n \amalg D^k_- \times \mathbb{R}^n$ obtained by identifying $(x, v) \in \partial D^k_+ \times \mathbb{R}^n$ with $(x, f(x)(v)) \in \partial D^k_- \times \mathbb{R}^n$. There is then a natural projection $E_f \to S^k$ and this is an n-dimensional vector bundle, as one can most easily see by taking an equivalent definition in which the two hemispheres of S^k are enlarged slightly to open balls and the identification occurs over their intersection, a product $S^{k-1} \times (-\varepsilon, \varepsilon)$, with the map f used in each slice $S^{k-1} \times \{t\}$. From this viewpoint the construction of E_f is a special case of the general construction of vector bundles by gluing together products described earlier in the discussion of tensor products.

The map f is called the *clutching function.* for E_f, presumably in reference to the mechanical clutch which engages and disengages gears in machinery. The same construction works equally well with $\mathbb{C}$ in place of $\mathbb{R}$, so from a map $f: S^{k-1} \to GL_n(\mathbb{C})$ one obtains a complex vector bundle $E_f \to S^k$.

Example 1.9. Let us see how the tangent bundle TS^2 to S^2 can be described in these terms. Define two orthogonal vector fields v_+ and w_+ on the northern hemisphere D^2_+ of S^2 in the following way. Start with a standard pair of orthogonal vectors at each point of a flat disk D^2 as in the left-hand figure below, then stretch the disk over the northern hemisphere of S^2, carrying the vectors along as tangent vectors to the resulting curved disk. As we travel around the equator of S^2 the vectors v_+ and w_+ then rotate through an angle of 2π relative to the equatorial direction, as in the right half of the figure.

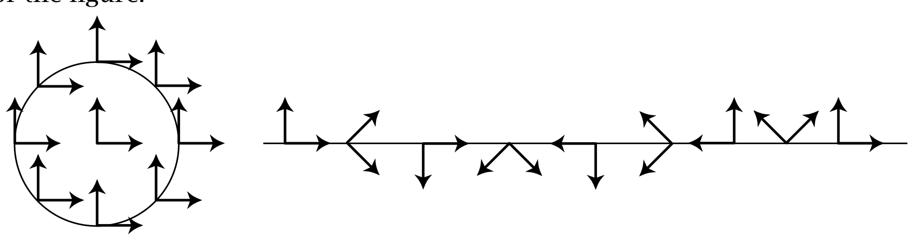

Reflecting everything across the equatorial plane, we obtain orthogonal vector fields v_- and w_- on the southern hemisphere D^2_-. The restrictions of v_- and w_- to the equator also rotate through an angle of 2π, but in the opposite direction from v_+ and w_+ since we have reflected across the equator. The pair $(v_\pm, w_\pm)$ defines a trivialization of TS^2 over $D^2_\pm$ taking $(v_\pm, w_\pm)$ to the standard basis for $\mathbb{R}^2$. Over the equator S^1 we then have two trivializations, and the function $f: S^1 \to GL_2(\mathbb{R})$ which rotates (v_+, w_+) to (v_-, w_-) sends $\theta \in S^1$, regarded as an angle, to rotation through the angle 2θ. For this map f we then have $E_f = TS^2$.

Example 1.10. Let us find a clutching function for the canonical complex line bundle over $\mathbb{C}P^1 = S^2$. The space $\mathbb{C}P^1$ is the quotient of $\mathbb{C}^2 - \{0\}$ under the equivalence relation $(z_0, z_1) \sim \lambda(z_0, z_1)$. Denote the equivalence class of (z_0, z_1) by $[z_0, z_1]$. We can also write points of $\mathbb{C}P^1$ as ratios $z = z_1/z_0 \in \mathbb{C} \cup \{\infty\} = S^2$. Points in the disk D^2_- inside the unit circle $S^1 \subset \mathbb{C}$ can be expressed uniquely in the form $[1, z_1/z_0] = [1, z]$ with $|z| \le 1$, and points in the disk D^2_+ outside S^1 can be written uniquely in the form $[z_0/z_1, 1] = [z^{-1}, 1]$ with $|z^{-1}| \le 1$. Over D^2_- a section of the canonical line bundle is then given by $[1, z_1/z_0] \mapsto (1, z_1/z_0)$ and over D^2_+ a section is $[z_0/z_1, 1] \mapsto (z_0/z_1, 1)$. These sections determine trivializations of the canonical line bundle over these two disks, and over their common boundary S^1 we pass from the D^2_+ trivialization to the D^2_- trivialization by multiplying by $z = z_1/z_0$. Thus the canonical line bundle is E_f for the clutching function $f : S^1 \to GL_1(\mathbb{C})$ defined by $f(z) = (z)$.

A basic property of the construction of bundles $E_f \to S^k$ via clutching functions $f : S^{k-1} \to GL_n(\mathbb{R})$ is that $E_f \approx E_g$ if f and g are homotopic. For if we have a homotopy $F : S^{k-1} \times I \to GL_n(\mathbb{R})$ from f to g, then we can construct by the same sort of clutching construction a vector bundle $E_F \to S^k \times I$ restricting to E_f over $S^k \times \{0\}$ and E_g over $S^k \times \{1\}$. Hence E_f and E_g are isomorphic by Proposition 1.7. Thus if we denote by $[X, Y]$ the set of homotopy classes of maps $X \to Y$, then the association $f \mapsto E_f$ gives a well-defined map $\Phi : [S^{k-1}, GL_n(\mathbb{R})] \longrightarrow \mathrm{Vect}^n(S^k)$.

There is a similar map for complex vector bundles, and this happens to have simpler behavior than in the real case:

Proposition 1.11. *The map* $\Phi : [S^{k-1}, GL_n(\mathbb{C})] \to \mathrm{Vect}^n_{\mathbb{C}}(S^k)$ *is a bijection.*

Proof: An inverse mapping Ψ can be constructed as follows. Given an n-dimensional vector bundle $p : E \to S^k$, its restrictions E_+ and E_- over D^k_+ and D^k_- are trivial since D^k_+ and D^k_- are contractible. Choose trivializations $h_\pm : E_\pm \to D^k_\pm \times \mathbb{C}^n$. Then $h_- h_+^{-1}$ defines a map $S^{k-1} \to GL_n(\mathbb{C})$, whose homotopy class is by definition $\Psi(E) \in \pi_{k-1} GL_n(\mathbb{C})$. To see that $\Psi(E)$ is well-defined, note first that any two choices of $h_\pm$ differ by a map $D^k_\pm \to GL_n(\mathbb{C})$. Since $D^k_\pm$ is contractible, such a map is homotopic to a constant map by composing with a contraction of $D^k_\pm$. Now we need the fact that $GL_n(\mathbb{C})$ is path-connected. The elementary row operation of modifying a matrix in $GL_n(\mathbb{C})$ by adding a scalar multiple of one row to another row can be realized by a path in $GL_n(\mathbb{C})$, just by inserting a factor of t in front of the scalar multiple and letting t go from 0 to 1. By such operations any matrix in $GL_n(\mathbb{C})$ can be diagonalized. The set of diagonal matrices in $GL_n(\mathbb{C})$ is path-connected since it is homeomorphic to the product of n copies of $\mathbb{C} - \{0\}$.

From this we conclude that h_+ and h_- are unique up to homotopy, hence the composition $h_- h_+^{-1} \to S^{k-1} \to GL_n(\mathbb{C})$ is also unique up to homotopy, which means that Ψ is a well-defined map $\mathrm{Vect}^n_{\mathbb{C}}(S^k) \to [S^{k-1}, GL_n(\mathbb{C})]$. It is clear that Ψ and Φ are inverses of each other. □

Example 1.12. Every complex vector bundle over S^1 is trivial, since by the proposition this is equivalent to saying that $[S^0, GL_n(\mathbb{C})]$ has a single element, or in other words that $GL_n(\mathbb{C})$ is path-connected, which is the case as we saw in the proof of the proposition.

Example 1.13. Let us show that the canonical line bundle $H \to \mathbb{C}\mathrm{P}^1$ satisfies the relation $(H \otimes H) \oplus 1 \approx H \oplus H$ where 1 is the trivial one-dimensional bundle. This can be seen by looking at the clutching functions for these two bundles, which are the maps $S^1 \to GL_2(\mathbb{C})$ given by

$$z \mapsto \begin{pmatrix} z^2 & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad z \mapsto \begin{pmatrix} z & 0 \\ 0 & z \end{pmatrix}$$

More generally, let us show the formula $E_{fg} \oplus n \approx E_f \oplus E_g$ for n-dimensional vector bundles E_f and E_g over S^k with clutching functions $f, g : S^{k-1} \to GL_n(\mathbb{C})$, where fg is the clutching function obtained by pointwise matrix multiplication, $fg(x) = f(x)g(x)$.

The bundle $E_f \oplus E_g$ has clutching function the map $f \oplus g : S^{k-1} \to GL_{2n}(\mathbb{C})$ having the matrices $f(x)$ in the upper left $n \times n$ block and the matrices $g(x)$ in the lower right $n \times n$ block, the other two blocks being zero. Since $GL_{2n}(\mathbb{C})$ is path-connected, there is a path $\alpha_t \in GL_{2n}(\mathbb{C})$ from the identity matrix to the matrix of the transformation which interchanges the two factors of $\mathbb{C}^n \times \mathbb{C}^n$. Then the matrix product $(f \oplus \mathbb{1})\alpha_t(\mathbb{1} \oplus g)\alpha_t$ gives a homotopy from $f \oplus g$ to $fg \oplus \mathbb{1}$, which is the clutching function for $E_{fg} \oplus n$.

The preceding analysis does not quite work for real vector bundles since $GL_n(\mathbb{R})$ is not path-connected. We can see that there are at least two path-components since the determinant function is a continuous surjection $GL_n(\mathbb{R}) \to \mathbb{R} - \{0\}$ to a space with two path-components. In fact $GL_n(\mathbb{R})$ has exactly two path-components, as we will now show. Just as in the complex case we can construct a path from an arbitrary matrix in $GL_n(\mathbb{R})$ to a diagonal matrix. Then by a path of diagonal matrices we can make all the diagonal entries $+1$ or -1. Two -1's represent a 180 degree rotation of a plane so they can be replaced by $+1$'s via a path in $GL_n(\mathbb{R})$. This shows that the subgroup $GL_n^+(\mathbb{R})$ consisting of matrices of positive determinant is path-connected. This subgroup has index 2, and $GL_n(\mathbb{R})$ is the disjoint union of the two cosets $GL_n^+(\mathbb{R})$ and $\alpha GL_n^+(\mathbb{R})$ for α a fixed matrix of determinant -1. The two cosets are homeomorphic since the map $\beta \mapsto \alpha\beta$ is a homeomorphism of $GL_n(\mathbb{R})$ with inverse the map $\beta \mapsto \alpha^{-1}\beta$. Thus both cosets are path-connected and so $GL_n(\mathbb{R})$ has two path-components.

The closest analogy with the complex case is obtained by considering oriented real vector bundles. Recall from linear algebra that an orientation of a real vector space is an equivalence class of ordered bases, two ordered bases being equiivalent if the invertible matrix taking the first basis to the second has positive determinant.

An *orientation* of a real vector bundle $p: E \to B$ is a function assigning an orientation to each fiber in such a way that near each point of B there is a local trivialization $h: p^{-1}(U) \to U \times \mathbb{R}^n$ carrying the orientations of fibers in $p^{-1}(U)$ to the standard orientation of $\mathbb{R}^n$ in the fibers of $U \times \mathbb{R}^n$. Another way of stating this condition would be to say that the orientations of fibers of E can be defined by ordered n-tuples of independent local sections. Not all vector bundles can be given an orientation. For example the Möbius bundle is not orientable. This is because an oriented line bundle over a paracompact base is always trivial since it has a canonical section formed by the unit vectors having positive orientation.

Let $\mathrm{Vect}^n_+(B)$ denote the set of isomorphism classes of oriented n-dimensional real vector bundles over B, where isomorphisms are required to preserve orientations. The clutching construction defines a map $\Phi: [S^{k-1}, GL^+_n(\mathbb{R})] \longrightarrow \mathrm{Vect}^n_+(S^k)$, and since $GL^+_n(\mathbb{R})$ is path-connected, the argument from the complex case shows:

Proposition 1.14. *The map* $\Phi: [S^{k-1}, GL^+_n(\mathbb{R})] \longrightarrow \mathrm{Vect}^n_+(S^k)$ *is a bijection.* □

To analyze $\mathrm{Vect}^n(S^k)$ let us introduct an *ad hoc* hybrid object $\mathrm{Vect}^n_0(S^k)$, the n-dimensional vector bundles over S^k with an orientation specified in the fiber over one point $x_0 \in S^{k-1}$, with the equivalence relation of isomorphism preserving the orientation of the fiber over x_0. We can choose the trivializations $h_\pm$ over $D^k_\pm$ to carry this orientation to a standard orientation of $\mathbb{R}^n$, and then $h_\pm$ is unique up to homotopy as before, so we obtain a bijection of $\mathrm{Vect}^n_0(S^k)$ with the homotopy classes of maps $S^{k-1} \to GL_n(\mathbb{R})$ taking x_0 to $GL^+_n(\mathbb{R})$.

The map $\mathrm{Vect}^n_0(S^k) \to \mathrm{Vect}^n(S^k)$ that forgets the orientation over x_0 is a surjection that is two-to-one except on vector bundles that have an automorphism (an isomorphism from the bundle to itself) reversing the orientation of the fiber over x_0, where it is one-to-one.

When $k = 1$ there are just two homotopy classes of maps $S^0 \to GL_n(\mathbb{R})$ taking x_0 to $GL^+_n(\mathbb{R})$, represented by maps taking x_0 to the identity and the other point of S^0 to either the identity or a reflection. The corresponding bundles are the trivial bundle and, when $n = 1$, the Möbius bundle, or the direct sum of the Möbius bundle with a trivial bundle when $n > 1$. These bundles all have automorphisms reversing orientations of fibers. We conclude that $\mathrm{Vect}^n(S^1)$ has exactly two elements for each $n \geq 1$. The nontrivial bundle is nonorientable since an orientable bundle over S^1 is trivial, any two clutching functions $S^0 \to GL^+_n(\mathbb{R})$ being homotopic.

When $k > 1$ the sphere S^{k-1} is path-connected, so maps $S^{k-1} \to GL_n(\mathbb{R})$ taking x_0 to $GL^+_n(\mathbb{R})$ take all of S^{k-1} to $GL^+_n(\mathbb{R})$. This implies that the natural map $\mathrm{Vect}^n_+(S^k) \to \mathrm{Vect}^n_0(S^k)$ is a bijection. Thus every vector bundle over S^k is orientable and has exactly two different orientations, determined by an orientation in one fiber. It follows that the map $\mathrm{Vect}^n_+(S^k) \to \mathrm{Vect}^n(S^k)$ is surjective and at most two-to-one.

It is one-to-one on bundles having an orientation-reversing automorphism, and two-to-one otherwise.

Inside $GL_n(\mathbb{R})$ is the orthogonal group $O(n)$, the subgroup consisting of isometries, or in matrix terms, the $n \times n$ matrices A with $AA^T = I$, that is, matrices whose columns form an orthonormal basis. The Gram-Schmidt orthogonalization process provides a deformation retraction of $GL_n(\mathbb{R})$ onto the subspace $O(n)$. Each step of the process is either normalizing the i^{th} vector of a basis to have length one or subtracting a linear combination of the preceding vectors from the i^{th} vector in order to make it orthogonal to them. Both operations are realizable by a homotopy. This is obvious for rescaling, and for the second operation one can just insert a scalar factor of t in front of the vectors being subtracted. and let t go from 0 to 1. The explicit algebraic formulas show that both operations depend continuously on the initial basis and have no effect on an orthonormal basis. Hence the sequence of operations provides a deformation retraction of $GL_n(\mathbb{R})$ onto $O(n)$. Restricting to the component $GL_n^+(\mathbb{R})$, we have a deformation retraction of this component onto the special orthogonal group $SO(n)$, the orthogonal matrices of determinant $+1$. This is the path-component of $O(n)$ containing the identity matrix since deformation retractions preserve path-components. The other component consists of orthogonal matrices of determinant -1.

In the complex case the same argument applied to $GL_n(\mathbb{C})$ gives a deformation retraction onto the unitary group $U(n)$, the $n \times n$ matrices over $\mathbb{C}$ whose columns form an orthonormal basis using the standard Hermitian inner product.

Composing maps $S^{k-1} \to SO(n)$ with the inclusion $SO(n) \hookrightarrow GL_n^+(\mathbb{R})$ gives a function $[S^{k-1}, SO(n)] \to [S^{k-1}, GL_n^+(\mathbb{R})]$. This is a bijection since the deformation retraction of $GL_n^+(\mathbb{R})$ onto $SO(n)$ gives a homotopy of any map $S^{k-1} \to GL_n^+(\mathbb{R})$ to a map into $SO(n)$, and if two such maps to $SO(n)$ are homotopic in $GL_n^+(\mathbb{R})$ then they are homotopic in $SO(n)$ by composing each stage of the homotopy with the retraction $GL_n^+(\mathbb{R}) \to SO(n)$ produced by the deformation retraction.

The advantage of $SO(n)$ over $GL_n^+(\mathbb{R})$ is that it is considerably smaller. For example, $SO(2)$ is just a circle, since orientation-preserving isometries of $\mathbb{R}^2$ are rotations, determined by the angle of rotation, a point in S^1. Taking $k = 2$ as well, we have for each integer $m \in \mathbb{Z}$ the map $f_m : S^1 \to S^1$, $z \mapsto z^m$, which is the clutching function for an oriented 2-dimensional vector bundle $E_m \to S^2$. We encountered E_2 in Example 1.9 as TS^2, and E_1 is the real vector bundle underlying the canonical complex line bundle over $\mathbb{C}\mathrm{P}^1 = S^2$, as we saw in Example 1.10. It is a fact proved in every introductory algebraic topology course that an arbitrary map $f : S^1 \to S^1$ is homotopic to one of the maps $z \mapsto z^m$ for some m uniquely determined by f, called the degree of f. Thus via the bundles E_m we have a bijection $\mathrm{Vect}_+^2(S^2) \approx \mathbb{Z}$. Reversing the orientation of E_m changes it to E_{-m}, so elements of $\mathrm{Vect}^2(S^2)$ correspond bijectively with integers $m \geq 0$, via the bundles E_m. Since $U(1)$ is also homeomorphic to S^1,we obtain a similar bijection $\text{Vect}^2_{\mathbb{C}}(S^2) \approx \mathbb{Z}$.

Another fact proved near the beginning of algebraic topology is that any two maps $S^{k-1} \to S^1$ are homotopic if $k > 2$. (Such a map lifts to the covering space $\mathbb{R}$ of S^1, and $\mathbb{R}$ is contractible.) Hence every 2-dimensional vector bundle over S^k is trivial when $k > 2$, as is any complex line bundle over S^k for $k > 2$.

It is a fact that any two maps $S^2 \to SO(n)$ or $S^2 \to U(n)$ are homotopic, so all vector bundles over S^3 are trivial. This is not too hard to show using a little basic machinery of algebraic topology, as we show in §??, where we also compute how many homotopy classes of maps $S^3 \to SO(n)$ and $S^3 \to U(n)$ there are.

For any space X the set $[X, SO(n)]$ has a natural group structure coming from the group structure in $SO(n)$. Namely, the product of two maps $f, g : X \to SO(n)$ is the map $x \mapsto f(x)g(x)$. Similarly $[X, U(n)]$ is a group. For example the bijection $[S^1, SO(2)] \approx \mathbb{Z}$ is a group isomorphism since for maps $S^1 \to S^1$, the product of $z \mapsto z^m$ and $z \mapsto z^n$ is $z \mapsto z^m z^n$, and $z^m z^n = z^{m+n}$.

The groups $[S^i, SO(n)]$ and $[S^i, U(n)]$ are isomorphic to the homotopy groups $\pi_i SO(n)$ and $\pi_i U(n)$, special cases of the homotopy groups $\pi_i X$ defined for all spaces X and central to algebraic topology. To define $\pi_i X$ requires choosing a basepoint $x_0 \in X$, and then $\pi_i X$ is the set of homotopy classes of maps $S^i \to X$ taking a chosen basepoint $s_0 \in S^i$ to x_0, where homotopies are also required to take s_0 to x_0 at all times. This basepoint data is needed in order to define a group structure in $\pi_i X$ when $i > 0$. (When $i = 0$ there is no natural group structure in general.) However, the group structure on $\pi_i X$ turns out to be independent of the choice of x_0 when X is path-connected. Taking $X = SO(n)$, there is an evident map $\pi_i SO(n) \to [S^i, SO(n)]$ obtained by ignoring basepoint data. To see that this is surjective, suppose we are given a map $f : S^i \to SO(n)$. Since $SO(n)$ is path-connected, there is a path α_t in $SO(n)$ from the identity matrix to $f(s_0)$, and then the matrix product $\alpha_t^{-1} f$ defines a homotopy from f to a map $S^i \to SO(n)$ taking s_0 to the identity matrix, which we choose as the basepoint of $SO(n)$. For injectivity of the map $\pi_i SO(n) \to [S^i, SO(n)]$ suppose that f_0 and f_1 are two maps $S^i \to SO(n)$ taking s_0 to the identity matrix, and suppose f_t is a homotopy from f_0 to f_1 that may not take basepoint to basepoint at all times. Then $[f_t(s_0)]^{-1} f_t$ is a new homotopy from f_0 to f_1 that does take basepoint to basepoint for all t. The same argument also applies for $U(n)$.

Note to the reader: The rest of this chapter will not be needed until Chapter 3.

The Universal Bundle

We will show that there is a special n-dimensional vector bundle $E_n \to G_n$ with the property that all n-dimensional bundles over paracompact base spaces are obtainable as pullbacks of this single bundle. When $n = 1$ this bundle will be just the canonical line bundle over $\mathbb{R}P^\infty$, defined earlier. The generalization to $n > 1$ will consist in replacing $\mathbb{R}P^\infty$, the space of 1-dimensional vector subspaces of $\mathbb{R}^\infty$, by the space of n-dimensional vector subspaces of $\mathbb{R}^\infty$.

First we define the *Grassmann manifold* $G_n(\mathbb{R}^k)$ for nonnegative integers $n \le k$. As a set this is the collection of all n-dimensional vector subspaces of $\mathbb{R}^k$, that is, n-dimensional planes in $\mathbb{R}^k$ passing through the origin. To define a topology on $G_n(\mathbb{R}^k)$ we first define the *Stiefel manifold* $V_n(\mathbb{R}^k)$ to be the space of orthonormal n-frames in $\mathbb{R}^k$, in other words, n-tuples of orthonormal vectors in $\mathbb{R}^k$. This is a subspace of the product of n copies of the unit sphere S^{k-1}, namely, the subspace of orthogonal n-tuples. It is a closed subspace since orthogonality of two vectors can be expressed by an algebraic equation. Hence $V_n(\mathbb{R}^k)$ is compact since the product of spheres is compact. There is a natural surjection $V_n(\mathbb{R}^k)\to G_n(\mathbb{R}^k)$ sending an n-frame to the subspace it spans, and $G_n(\mathbb{R}^k)$ is topologized by giving it the quotient topology with respect to this surjection. So $G_n(\mathbb{R}^k)$ is compact as well. Later in this section we will construct a finite CW complex structure on $G_n(\mathbb{R}^k)$ and in the process show that it is Hausdorff and a manifold of dimension $n(k-n)$.

The inclusions $\mathbb{R}^k \subset \mathbb{R}^{k+1} \subset \cdots$ give inclusions $G_n(\mathbb{R}^k) \subset G_n(\mathbb{R}^{k+1}) \subset \cdots$, and we let $G_n(\mathbb{R}^\infty) = \bigcup_k G_n(\mathbb{R}^k)$. This is the set of all n-dimensional vector subspaces of the vector space $\mathbb{R}^\infty$ since if we choose a basis for a finite-dimensional subspace of $\mathbb{R}^\infty$, each basis vector will lie in some $\mathbb{R}^k$, hence there will be an $\mathbb{R}^k$ containing all the basis vectors and therefore the whole subspace. We give $G_n(\mathbb{R}^\infty)$ the weak or direct limit topology, so a set in $G_n(\mathbb{R}^\infty)$ is open iff it intersects each $G_n(\mathbb{R}^k)$ in an open set.

There are canonical n-dimensional vector bundles over $G_n(\mathbb{R}^k)$ and $G_n(\mathbb{R}^\infty)$. Define $E_n(\mathbb{R}^k) = \{\,(\ell,v) \in G_n(\mathbb{R}^k)\times\mathbb{R}^k \mid v \in \ell\,\}$. The inclusions $\mathbb{R}^k \subset \mathbb{R}^{k+1} \subset \cdots$ give inclusions $E_n(\mathbb{R}^k) \subset E_n(\mathbb{R}^{k+1}) \subset \cdots$ and we set $E_n(\mathbb{R}^\infty) = \bigcup_k E_n(\mathbb{R}^k)$, again with the weak or direct limit topology.

Lemma 1.15. *The projection $p:E_n(\mathbb{R}^k)\to G_n(\mathbb{R}^k)$, $p(\ell,v)=\ell$, is a vector bundle., both for finite and infinite k.*

Proof: First suppose k is finite. For $\ell \in G_n(\mathbb{R}^k)$, let $\pi_\ell:\mathbb{R}^k\to\ell$ be orthogonal projection and let $U_\ell = \{\,\ell' \in G_n(\mathbb{R}^k) \mid \pi_\ell(\ell') \text{ has dimension } n\,\}$. In particular, $\ell \in U_\ell$. We will show that U_ℓ is open in $G_n(\mathbb{R}^k)$ and that the map $h:p^{-1}(U_\ell)\to U_\ell\times\ell \approx U_\ell\times\mathbb{R}^n$ defined by $h(\ell',v) = (\ell',\pi_\ell(v))$ is a local trivialization of $E_n(\mathbb{R}^k)$.

For U_ℓ to be open is equivalent to its preimage in $V_n(\mathbb{R}^k)$ being open. This preimage consists of orthonormal frames $v_1,\cdots,v_n$ such that $\pi_\ell(v_1),\cdots,\pi_\ell(v_n)$ are independent. Let A be the matrix of π_ℓ with respect to the standard basis in the domain $\mathbb{R}^k$ and any fixed basis in the range ℓ. The condition on $v_1,\cdots,v_n$ is then that the $n\times n$ matrix with columns $Av_1,\cdots,Av_n$ have nonzero determinant. Since the value of this determinant is obviously a continuous function of $v_1,\cdots,v_n$, it follows that the frames $v_1,\cdots,v_n$ yielding a nonzero determinant form an open set in $V_n(\mathbb{R}^k)$.

It is clear that h is a bijection which is a linear isomorphism on each fiber. We need to check that h and h^{-1} are continuous. For $\ell' \in U_\ell$ there is a unique invertible linear map $L_{\ell'} : \mathbb{R}^k \to \mathbb{R}^k$ restricting to π_ℓ on ℓ' and the identity on $\ell^\perp = \operatorname{Ker} \pi_\ell$. We claim that $L_{\ell'}$, regarded as a $k \times k$ matrix, depends continuously on ℓ'. Namely, we can write $L_{\ell'}$ as a product AB^{-1} where:

— B sends the standard basis to $v_1, \cdots, v_n, v_{n+1}, \cdots, v_k$ with $v_1, \cdots, v_n$ an orthonormal basis for ℓ' and $v_{n+1}, \cdots, v_k$ a fixed basis for $\ell^\perp$.

— A sends the standard basis to $\pi_\ell(v_1), \cdots, \pi_\ell(v_n), v_{n+1}, \cdots, v_k$.

Both A and B depend continuously on $v_1, \cdots, v_n$. Since matrix multiplication and matrix inversion are continuous operations (think of the 'classical adjoint' formula for the inverse of a matrix), it follows that the product $L_{\ell'} = AB^{-1}$ depends continuously on $v_1, \cdots, v_n$. But since $L_{\ell'}$ depends only on ℓ', not on the basis $v_1, \cdots, v_n$ for ℓ', it follows that $L_{\ell'}$ depends continuously on ℓ' since $G_n(\mathbb{R}^k)$ has the quotient topology from $V_n(\mathbb{R}^k)$. Since we have $h(\ell', v) = (\ell', \pi_\ell(v)) = (\ell', L_{\ell'}(v))$, we see that h is continuous. Similarly, $h^{-1}(\ell', w) = (\ell', L_{\ell'}^{-1}(w))$ and $L_{\ell'}^{-1}$ depends continuously on ℓ', matrix inversion being continuous, so h^{-1} is continuous.

This finishes the proof for finite k. When $k = \infty$ one takes U_ℓ to be the union of the U_ℓ's for increasing k. The local trivializations h constructed above for finite k then fit together to give a local trivialization over this U_ℓ, continuity being automatic since we use the weak topology. □

We will mainly be interested in the case $k = \infty$ now, and to simplify notation we will write G_n for $G_n(\mathbb{R}^\infty)$ and E_n for $E_n(\mathbb{R}^\infty)$. As earlier in this section, we use the notation $[X, Y]$ for the set of homotopy classes of maps $f : X \to Y$.

Theorem 1.16. *For paracompact X, the map $[X, G_n] \to \operatorname{Vect}^n(X)$, $[f] \mapsto f^*(E_n)$, is a bijection.*

Thus, vector bundles over a fixed base space are classified by homotopy classes of maps into G_n. Because of this, G_n is called the *classifying space* for n-dimensional vector bundles and $E_n \to G_n$ is called the *universal bundle*.

As an example of how a vector bundle could be isomorphic to a pullback $f^*(E_n)$, consider the tangent bundle to S^n. This is the vector bundle $p : E \to S^n$ where $E = \{ (x, v) \in S^n \times \mathbb{R}^{n+1} \mid x \perp v \}$. Each fiber $p^{-1}(x)$ is a point in $G_n(\mathbb{R}^{n+1})$, so we have a map $S^n \to G_n(\mathbb{R}^{n+1})$, $x \mapsto p^{-1}(x)$. Via the inclusion $\mathbb{R}^{n+1} \hookrightarrow \mathbb{R}^\infty$ we can view this as a map $f : S^n \to G_n(\mathbb{R}^\infty) = G_n$, and E is exactly the pullback $f^*(E_n)$.

Proof of 1.16: The key observation is the following: For an n-dimensional vector bundle $p : E \to X$, an isomorphism $E \approx f^*(E_n)$ is equivalent to a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber. To see this, suppose first that we have a map $f : X \to G_n$ and an isomorphism $E \approx f^*(E_n)$. Then we have a commutative diagram

$$\begin{array}{ccccc} E \approx f^*(E_n) & \xrightarrow{\widetilde{f}} & E_n & \xrightarrow{\pi} & \mathbb{R}^\infty \\ {\scriptstyle p}\searrow \quad \downarrow & & \downarrow & & \\ X & \xrightarrow{f} & G_n & & \end{array}$$

where $\pi(\ell, v) = v$. The composition across the top row is a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber, since both $\widetilde{f}$ and π have this property. Conversely, given a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber, define $f : X \to G_n$ by letting $f(x)$ be the n-plane $g(p^{-1}(x))$. This clearly yields a commutative diagram as above.

To show surjectivity of the map $[X, G_n] \longrightarrow \mathrm{Vect}^n(X)$, suppose $p : E \to X$ is an n-dimensional vector bundle. Let $\{U_\alpha\}$ be an open cover of X such that E is trivial over each U_α. By Lemma 1.21 in the Appendix to this chapter there is a countable open cover $\{U_i\}$ of X such that E is trivial over each U_i, and there is a partition of unity $\{\varphi_i\}$ with φ_i supported in U_i. Let $g_i : p^{-1}(U_i) \to \mathbb{R}^n$ be the composition of a trivialization $p^{-1}(U_i) \to U_i \times \mathbb{R}^n$ with projection onto $\mathbb{R}^n$. The map $(\varphi_i p) g_i$, $v \mapsto \varphi_i(p(v)) g_i(v)$, extends to a map $E \to \mathbb{R}^n$ that is zero outside $p^{-1}(U_i)$. Near each point of X only finitely many φ_i's are nonzero, and at least one φ_i is nonzero, so these extended $(\varphi_i p) g_i$'s are the coordinates of a map $g : E \to (\mathbb{R}^n)^\infty = \mathbb{R}^\infty$ that is a linear injection on each fiber.

For injectivity, if we have isomorphisms $E \approx f_0^*(E_n)$ and $E \approx f_1^*(E_n)$ for two maps $f_0, f_1 : X \to G_n$, then these give maps $g_0, g_1 : E \to \mathbb{R}^\infty$ that are linear injections on fibers, as in the first paragraph of the proof. We claim g_0 and g_1 are homotopic through maps g_t that are linear injections on fibers. If this is so, then f_0 and f_1 will be homotopic via $f_t(x) = g_t(p^{-1}(x))$.

The first step in constructing a homotopy g_t is to compose g_0 with the homotopy $L_t : \mathbb{R}^\infty \to \mathbb{R}^\infty$ defined by $L_t(x_1, x_2, \cdots) = (1-t)(x_1, x_2, \cdots) + t(x_1, 0, x_2, 0, \cdots)$. For each t this is a linear map whose kernel is easily computed to be 0, so L_t is injective. Composing the homotopy L_t with g_0 moves the image of g_0 into the odd-numbered coordinates. Similarly we can homotope g_1 into the even-numbered coordinates. Still calling the new g's g_0 and g_1, let $g_t = (1-t)g_0 + t g_1$. This is linear and injective on fibers for each t since g_0 and g_1 are linear and injective on fibers. $\square$

An explicit calculation of $[X, G_n]$ is usually beyond the reach of what is possible technically, so this theorem is of limited usefulness in enumerating all the different vector bundles over a given base space. Its importance is due more to its theoretical implications. Among other things, it can reduce the proof of a general statement to the special case of the universal bundle. For example, it is easy to deduce that vector bundles over a paracompact base have inner products, since the bundle $E_n \to G_n$ has an obvious inner product obtained by restricting the standard inner product in $\mathbb{R}^\infty$ to each n-plane, and this inner product on E_n induces an inner product on every pullback $f^*(E_n)$.

The preceding constructions and results hold equally well for vector bundles over $\mathbb{C}$, with $G_n(\mathbb{C}^k)$ the space of n-dimensional $\mathbb{C}$-linear subspaces of $\mathbb{C}^k$, and so on. In particular, the proof of Theorem 1.16 translates directly to complex vector bundles, showing that $\mathrm{Vect}^n_{\mathbb{C}}(X) \approx [X, G_n(\mathbb{C}^\infty)]$.

There is also a version of Theorem 1.16 for oriented vector bundles. Let $\widetilde{G}_n(\mathbb{R}^k)$ be the space of oriented n-planes in $\mathbb{R}^k$, the quotient space of $V_n(\mathbb{R}^k)$ obtained by identifying two n-frames when they determine the same oriented subspace of $\mathbb{R}^k$. Forming the union over increasing k we then have the space $\widetilde{G}_n(\mathbb{R}^\infty)$. The universal oriented bundle $\widetilde{E}_n(\mathbb{R}^\infty)$ over $\widetilde{G}_n(\mathbb{R}^\infty)$ consists of pairs $(\ell, v) \in \widetilde{G}_n(\mathbb{R}^\infty) \times \mathbb{R}^\infty$ with $v \in \ell$. In other words, $\widetilde{E}_n(\mathbb{R}^\infty)$ is the pullback of $E_n(\mathbb{R}^\infty)$ via the natural projection $\widetilde{G}_n(\mathbb{R}^\infty) \to G_n(\mathbb{R}^\infty)$. Small modifications in the proof that $\mathrm{Vect}^n(X) \approx [X, G_n(\mathbb{R}^\infty)]$ show that $\mathrm{Vect}^n_+(X) \approx [X, \widetilde{G}_n(\mathbb{R}^\infty)]$.

Both $G_n(\mathbb{R}^\infty)$ and $\widetilde{G}_n(\mathbb{R}^\infty)$ are path-connected since $\mathrm{Vect}^n(X)$ and $\mathrm{Vect}^n_+(X)$ have a single element when X is a point. The natural projection $\widetilde{G}_n(\mathbb{R}^\infty) \to G_n(\mathbb{R}^\infty)$ obtained by ignoring orientations is two-to-one, and readers familiar with the notion of a covering space will have no trouble in recognizing that this two-to-one projection map is a covering space, using for example the local trivializations constructed in Lemma 1.15. In fact $\widetilde{G}_n(\mathbb{R}^\infty)$ is the universal cover of $G_n(\mathbb{R}^\infty)$ since it is simply-connected, because of the triviality of $\mathrm{Vect}^n_+(S^1) \approx [S^1, \widetilde{G}_n(\mathbb{R}^\infty)]$. One can also observe that a vector bundle $E \approx f^*(E_n(\mathbb{R}^\infty)$ is orientable iff its classifying map $f : X \to G_n(\mathbb{R}^\infty)$ lifts to a map $\widetilde{f} : X \to \widetilde{G}_n(\mathbb{R}^\infty)$, and in fact orientations of E correspond bijectively with lifts $\widetilde{f}$.

Cell Structures on Grassmannians

Since Grassmann manifolds play a fundamental role in vector bundle theory, it would be good to have a better grasp on their topology. Here we show that $G_n(\mathbb{R}^\infty)$ has the structure of a CW complex with each $G_n(\mathbb{R}^k)$ a finite subcomplex. We will also see that $G_n(\mathbb{R}^k)$ is a closed manifold of dimension $n(k-n)$. Similar statements hold in the complex case as well, with $G_n(\mathbb{C}^k)$ a closed manifold of dimension $2n(k-n)$.

For a start let us show that $G_n(\mathbb{R}^k)$ is Hausdorff, since we will need this fact later when we construct the CW structure. Given two n-planes ℓ and ℓ' in $G_n(\mathbb{R}^k)$, it suffices to find a continuous map $f : G_n(\mathbb{R}^k) \to \mathbb{R}$ taking different values on ℓ and ℓ'. For a vector $v \in \mathbb{R}^k$ let $f_v(\ell)$ be the length of the orthogonal projection of v onto ℓ. This is a continuous function of ℓ since if we choose an orthonormal basis $v_1, \cdots, v_n$ for ℓ then $f_v(\ell) = ((v \cdot v_1)^2 + \cdots + (v \cdot v_n)^2)^{1/2}$, which is certainly continuous in $v_1, \cdots, v_n$ hence in ℓ since $G_n(\mathbb{R}^k)$ has the quotient topology from $V_n(\mathbb{R}^k)$. Now for an n-plane $\ell' \neq \ell$ choose $v \in \ell - \ell'$, and then $f_v(\ell) = |v| > f_v(\ell')$.

There is a nice description of the cells in the CW structure on $G_n(\mathbb{R}^k)$ in terms of the familiar concept of echelon form for matrices. Recall that any $n \times k$ matrix A can be put into an echelon form by a finite sequence of elementary row operations consisting either adding a multiple of one row to another, multiplying a row by a nonzero scalar, or permuting two rows. In the standard version of echelon form one strives to make zeros in the lower left corner of the matrix, but for our purposes it will be more convenient to make zeros in the upper right corner instead. Thus a matrix in echelon form will look like the following, with the asterisks denoting entries that are arbitrary numbers:

$$\begin{bmatrix} * & * & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 1 & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & 1 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & 0 & * & * & 1 & 0 \end{bmatrix}$$

We are assuming that our given $n\times k$ matrix A has rank n and that $n \le k$. The shape of the echelon form is specified by which columns contain the special entries 1, say in the columns numbered $\sigma_1 < \cdots < \sigma_n$. The n-tuple $(\sigma_1,\cdots,\sigma_n)$ is called the *Schubert symbol* $\sigma(A)$. To see that this depends only on the givne matrix A and not any particular reduction of A to echelon form, consider the n-plane ℓ_A in $\mathbb{R}^k$ spanned by the rows of A. This is the same as the n-plane spanned by the rows of an echelon form of A since elementary row operations do not change the subspace spanned by the rows. Let $p_i:\mathbb{R}^k\to\mathbb{R}^{n-i}$ be projection onto the last $n-i$ coordinates. As i goes from 0 to k the dimension of the subspace $p_i(\ell_A)$ of $\mathbb{R}^k$ decreases from n to 0 in n steps and the decreases occur precisely when i hits one of the values σ_j, as is apparent from the echelon form since $p_i(\ell_A)$ has basis the rows of the matrix obtained by deleting the first i columns of the echelon form. Thus the Schubert symbol only depends on the n-plane spanned by the rows of A.

In fact the echelon form itself depends only on the n-plane spanned by the rows. To see this, consider projection onto the n coordinates of $\mathbb{R}^k$ given by the numbers σ_i. This projection is surjective when restricted to the n-plane, hence is also injective, and the n rows of the echelon form are the vectors in the n-plane that project to the standard basis vectors.

As noted above, the numbers σ_j in the Schubert symbol of an n-plane ℓ are the numbers i for which the dimension of the image $p_i(\ell)$ decreases as i goes from 0 to n. Equivalently, these are the i's for which the dimension of the kernel of the restriction $p_i|\ell$ increases. This kernel is just $\ell\cap\mathbb{R}^i$ where $\mathbb{R}^i$ is the kernel of p_i, the subspace of $\mathbb{R}^k$ spanned by the first i standard basis vectors. This gives an alternative definition of the Schubert symbol $\sigma(\ell)$.

Given a Schubert symbol σ one can consider the set $e(\sigma)$ of all n-planes in $\mathbb{R}^k$ having σ as their Schubert symbol. In terms of echelon forms, the various n-planes in $e(\sigma)$ are parametrized by the arbitrary entries in the echelon form. There are $\sigma_i - i$ of these entries in the i^{th} row, for a total of $(ß_1-1)+\cdots(\sigma_n-n)$ entries. Thus $e(\sigma)$ is homeomorphic to a euclidean space of this dimension, or equivalently an open cell.

Proposition 1.17. *The cells $e(\sigma)$ are the cells of a CW structure on $G_n(\mathbb{R}^k)$.*

For example $G_2(\mathbb{R}^4)$ has six cells corresponding to the Schubert symbols $(1,2)$, $(1,3)$, $(1,4)$, $(2,3)$, $(2,4)$, $(3,4)$, and these cells have dimensions 0, 1, 2, 2, 3, 4 respectively. In general the number of cells of $G_n(\mathbb{R}^k)$ is $\binom{k}{n}$, the number of ways of choosing the n distinct numbers $\sigma_i \le k$.

Proof: Our main task will be to find a characteristic map for $e(\sigma)$. This is a map from a closed ball of the same dimension as $e(\sigma)$ into $G_n(\mathbb{R}^k)$ whose restriction to the interior of the ball is a homeomorphism onto $e(\sigma)$. From the echelon forms described above it is not clear how to do this, so we will use a slightly different sort of echelon form. We allow the special 1's to be arbitrary nonzero numbers and we allow the entries below these 1's to be nonzero. Then we impose the condition that the rows be orthonormal and that the last nonzero entry in each row be positive. Let us call this an *orthonormal echelon form*. Once again there is a unique orthonormal echelon form for each n-plane ℓ since if we let ℓ_i denote the subspace of ℓ spanned by the first i rows of the standard echelon form, or in other words $\ell_i = \ell \cap \mathbb{R}^{\sigma_i}$, then there is a unique unit vector in ℓ_i orthogonal to ℓ_{i-1} and having positive $\sigma_i{}^{th}$ coordinate.

The i^{th} row of the orthonormal echelon form then belongs to the hemisphere H_i in the unit sphere $S^{\sigma_i - 1} \subset \mathbb{R}^{\sigma_i} \subset \mathbb{R}^k$ consisting of unit vectors with non-negative $\sigma_i{}^{th}$ coordinate. In the Stiefel manifold $V_n(\mathbb{R}^k)$ let $E(\sigma)$ be the subspace of orthonormal frames $(v_1, \cdots, v_n) \in (S^{k-1})^n$ such that $v_i \in H_i$ for each i. We claim that $E(\sigma)$ is homeomorphic to a closed ball. To prove this the main step is to show that the projection $\pi : E(\sigma) \to H_1$, $\pi(v_1, \cdots, v_n) = v_1$, is a trivial fiber bundle. This is equivalent to finding a projection $p : E(\sigma) \to \pi^{-1}(v_0)$ which is a homeomorphism on fibers of π, where $v_0 = (0, \cdots, 0, 1) \in \mathbb{R}^{\sigma_1} \subset \mathbb{R}^k$, since the map $\pi \times p : E(\sigma) \to H_1 \times \pi^{-1}(v_0)$ is then a continuous bijection of compact Hausdorff spaces, hence a homeomorphism. The map $p : \pi^{-1}(v) \to \pi^{-1}(v_0)$ is obtained by applying the rotation ρ_v of $\mathbb{R}^k$ that takes v to v_0 and fixes the $(k-2)$-dimensional subspace orthogonal to v and v_0. This rotation takes H_i to itself for $i > 1$ since it affects only the first σ_1 coordinates of vectors in $\mathbb{R}^k$. Hence p takes $\pi^{-1}(v)$ onto $\pi^{-1}(v_0)$.

The fiber $\pi^{-1}(v_0)$ can be identified with $E(\sigma')$ for $\sigma' = (\sigma_2 - 1, \cdots, \sigma_n - 1)$. By induction on n this is homeomorphic to a closed ball of dimension $(\sigma_2 - 2) + \cdots + (\sigma_n - n)$, so $E(\sigma)$ is a closed ball of dimension $(\sigma_1 - 1) + \cdots + (\sigma_n - n)$. The boundary of this ball consists of points in $E(\sigma)$ having v_i in ∂H_i for at least one i. This too follows by induction since the rotation ρ_v takes ∂H_i to itself for $i > 1$.

The natural map $E(\sigma) \to G_n$ sending an orthonormal n-tuple to the n-plane it spans takes the interior of the ball $E(\sigma)$ to $e(\sigma)$ bijectively. Since G_n has the quotient topology from V_n, the map $\operatorname{int} E(\sigma) \to e(\sigma)$ is a homeomorphism. The boundary of

$E(\sigma)$ maps to cells $e(\sigma')$ of G_n where σ' is obtained from σ by decreasing some σ_i's, so these cells $e(\sigma')$ have lower dimension than $e(\sigma)$.

To see that the maps $E(\sigma)\to G_n(\mathbb{R}^k)$ for the cells $e(\sigma)$ are the characteristic maps for a CW structure on $G_n(\mathbb{R}^k)$ we can argue as follows. Let X^i be the union of the cells $e(\sigma)$ in $G_n(\mathbb{R}^k)$ having dimension at most i. Suppose by induction on i that X^i is a CW complex with these cells. Attaching the $(i+1)$-cells $e(\sigma)$ of X^{i+1} to X^i via the maps $\partial E(\sigma)\to X^i$ produces a CW complex Y and a natural continuous bijection $Y\to X^{i+1}$. Since Y is a finite CW complex it is compact, and X^{i+1} is Hausdorff as a subspace of $G_n(\mathbb{R}^k)$, so the map $Y\to X^{i+1}$ is a homeomorphism and X^{i+1} is a CW complex, finishing the induction. Thus we have a CW structure on $G_n(\mathbb{R}^k)$. □

Since the inclusions $G_n(\mathbb{R}^k)\subset G_n(\mathbb{R}^{k+1})$ for varying k are inclusions of subcomplexes and $G_n(\mathbb{R}^\infty)$ has the weak topology with respect to these subspaces, it follows that we have also a CW structure on $G_n(\mathbb{R}^\infty)$.

Similar constructions work to give CW structures on complex Grassmann manifolds, but here $e(\sigma)$ will be a cell of dimension $(2\sigma_1-2)+(2\sigma_2-4)+\cdots+(2\sigma_n-2n)$. The hemisphere H_i is defined to be the subspace of the unit sphere $S^{2\sigma_i-1}$ in $\mathbb{C}^{\sigma_i}$ consisting of vectors whose ${\sigma_i}^{th}$ coordinate is real and nonnegative, so H_i is a ball of dimension $2\sigma_i-2$. The transformation $\rho_v\in SU(k)$ is uniquely determined by specifying that it takes v to v_0 and fixes the orthogonal $(k-2)$-dimensional complex subspace, since an element of $U(2)$ of determinant 1 is determined by where it sends one unit vector.

The highest-dimensional cell of $G_n(\mathbb{R}^k)$ is $e(\sigma)$ for $\sigma=(k-n+1,k-n+2,\cdots,k)$, of dimension $n(k-n)$, so this is the dimension of $G_n(\mathbb{R}^k)$. Near points in these top-dimensional cells $G_n(\mathbb{R}^k)$ is a manifold. But $G_n(\mathbb{R}^k)$ is homogeneous in the sense that given any two points in $G_n(\mathbb{R}^k)$ there is a homeomorphism $G_n(\mathbb{R}^k)\to G_n(\mathbb{R}^k)$ taking one point to the other, namely, the homeomorphism induced by an invertible linear map $\mathbb{R}^k\to\mathbb{R}^k$ taking one n-plane to the other. From this homogeneity it follows that $G_n(\mathbb{R}^k)$ is a manifold near all points. Since it is compact, it is a closed manifold.

There is a natural inclusion $i:G_n\hookrightarrow G_{n+1}$, $i(\ell)=\mathbb{R}\times j(\ell)$ where $j:\mathbb{R}^\infty\to\mathbb{R}^\infty$ is the embedding $j(x_1,x_2,\cdots)=(0,x_1,x_2,\cdots)$. If $\sigma(\ell)=(\sigma_1,\cdots,\sigma_n)$ then $\sigma(i(\ell))=(1,\sigma_1+1,\cdots,\sigma_n+1)$, so i takes cells of G_n to cells of G_{n+1} of the same dimension, making $i(G_n)$ a subcomplex of G_{n+1}. Identifying G_n with the subcomplex $i(G_n)$, we obtain an increasing sequence of CW complexes $G_1\subset G_2\subset\cdots$ whose union $G_\infty=\bigcup_n G_n$ is therefore also a CW complex. Similar remarks apply as well in the complex case.

Appendix: Paracompactness

A Hausdorff space X is *paracompact* if for each open cover $\{U_\alpha\}$ of X there is a partition of unity $\{\varphi_\beta\}$ subordinate to the cover. This means that the φ_β's are maps $X \to I$ such that each φ_β has support (the closure of the set where $\varphi_\beta \neq 0$) contained in some U_α, each $x \in X$ has a neighborhood in which only finitely many φ_β's are nonzero, and $\sum_\beta \varphi_\beta = 1$. An equivalent definition which is often given is that X is Hausdorff and every open cover of X has a locally finite open refinement. The first definition clearly implies the second by taking the cover $\{\varphi_\beta^{-1}(0,1]\}$. For the converse, see [Dugundji] or [Lundell-Weingram]. It is the former definition which is most useful in algebraic topology, and the fact that the two definitions are equivalent is rarely if ever needed. So we shall use the first definition.

A paracompact space X is normal, for let A_1 and A_2 be disjoint closed sets in X, and let $\{\varphi_\beta\}$ be a partition of unity subordinate to the cover $\{X - A_1, X - A_2\}$. Let φ_i be the sum of the φ_β's which are nonzero at some point of A_i. Then $\varphi_i(A_i) = 1$, and $\varphi_1 + \varphi_2 \leq 1$ since no φ_β can be a summand of both φ_1 and φ_2. Hence $\varphi_1^{-1}(1/2, 1]$ and $\varphi_2^{-1}(1/2, 1]$ are disjoint open sets containing A_1 and A_2, respectively.

Most of the spaces one meets in algebraic topology are paracompact, including:

(1) compact Hausdorff spaces

(2) unions of increasing sequences $X_1 \subset X_2 \subset \cdots$ of compact Hausdorff spaces X_i, with the weak or direct limit topology (a set is open iff it intersects each X_i in an open set)

(3) CW complexes

(4) metric spaces

Note that (2) includes (3) for CW complexes with countably many cells, since such a CW complex can be expressed as an increasing union of finite subcomplexes. Using (1) and (2), it can be shown that many manifolds are paracompact, for example $\mathbb{R}^n$.

The next three propositions verify that the spaces in (1), (2), and (3) are paracompact.

Proposition 1.18. *A compact Hausdorff space X is paracompact.*

Proof: Let $\{U_\alpha\}$ be an open cover of X. Since X is normal, each $x \in X$ has an open neighborhood V_x with closure contained in some U_α. By Urysohn's lemma there is a map $\varphi_x : X \to I$ with $\varphi_x(x) = 1$ and $\varphi_x(X - V_x) = 0$. The open cover $\{\varphi_x^{-1}(0,1]\}$ of X contains a finite subcover, and we relabel the corresponding φ_x's as φ_β's. Then $\sum_\beta \varphi_\beta(x) > 0$ for all x, and we obtain the desired partition of unity subordinate to $\{U_\alpha\}$ by normalizing each φ_β by dividing it by $\sum_\beta \varphi_\beta$. □

Proposition 1.19. *If X is the direct limit of an increasing sequence $X_1 \subset X_2 \subset \cdots$ of compact Hausdorff spaces X_i, then X is paracompact.*

Proof: A preliminary observation is that X is normal. To show this, it suffices to find a map $f:X\to I$ with $f(A)=0$ and $f(B)=1$ for any two disjoint closed sets A and B. Such an f can be constructed inductively over the X_i's, using normality of the X_i's. For the induction step one has f defined on the closed set $X_i\cup(A\cap X_{i+1})\cup(B\cap X_{i+1})$ and one extends over X_{i+1} by Tietze's theorem.

To prove that X is paracompact, let an open cover $\{U_\alpha\}$ be given. Since X_i is compact Hausdorff, there is a finite partition of unity $\{\varphi_{ij}\}$ on X_i subordinate to $\{U_\alpha\cap X_i\}$. Using normality of X, extend each φ_{ij} to a map $\varphi_{ij}:X\to I$ with support in the same U_α. Let $\sigma_i=\sum_j\varphi_{ij}$. This sum is 1 on X_i, so if we normalize each φ_{ij} by dividing it by $\max\{1/2,\sigma_i\}$, we get new maps φ_{ij} with $\sigma_i=1$ in a neighborhood V_i of X_i. Let $\psi_{ij}=\max\{0,\varphi_{ij}-\sum_{k<i}\sigma_k\}$. Since $0\le\psi_{ij}\le\varphi_{ij}$, the collection $\{\psi_{ij}\}$ is subordinate to $\{U_\alpha\}$. In V_i all ψ_{kj}'s with $k>i$ are zero, so each point of X has a neighborhood in which only finitely many ψ_{ij}'s are nonzero. For each $x\in X$ there is a ψ_{ij} with $\psi_{ij}(x)>0$, since if $\varphi_{ij}(x)>0$ and i is minimal with respect to this condition, then $\psi_{ij}(x)=\varphi_{ij}(x)$. Thus when we normalize the collection $\{\psi_{ij}\}$ by dividing by $\sum_{i,j}\psi_{ij}$ we obtain a partition of unity on X subordinate to $\{U_\alpha\}$. ◻

Proposition 1.20. *Every CW complex is paracompact.*

Proof: Given an open cover $\{U_\alpha\}$ of a CW complex X, suppose inductively that we have a partition of unity $\{\varphi_\beta\}$ on X^n subordinate to the cover $\{U_\alpha\cap X^n\}$. For a cell e^{n+1}_γ with characteristic map $\Phi_\gamma:D^{n+1}\to X$, $\{\varphi_\beta\Phi_\gamma\}$ is a partition of unity on $S^n=\partial D^{n+1}$. Since S^n is compact, only finitely many of these compositions $\varphi_\beta\Phi_\gamma$ can be nonzero, for fixed γ. We extend these functions $\varphi_\beta\Phi_\gamma$ over D^{n+1} by the formula $\rho_\varepsilon(r)\varphi_\beta\Phi_\gamma(x)$ using 'spherical coordinates' $(r,x)\in I\times S^n$ on D^{n+1}, where $\rho_\varepsilon:I\to I$ is 0 on $[0,1-\varepsilon]$ and 1 on $[1-\varepsilon/2,1]$. If $\varepsilon=\varepsilon_\gamma$ is chosen small enough, these extended functions $\rho_\varepsilon\varphi_\beta\Phi_\gamma$ will be subordinate to the cover $\{\Phi_\gamma^{-1}(U_\alpha)\}$. Let $\{\psi_{\gamma j}\}$ be a finite partition of unity on D^{n+1} subordinate to $\{\Phi_\gamma^{-1}(U_\alpha)\}$. Then $\{\rho_\varepsilon\varphi_\beta\Phi_\gamma,(1-\rho_\varepsilon)\psi_{\gamma j}\}$ is a partition of unity on D^{n+1} subordinate to $\{\Phi_\gamma^{-1}(U_\alpha)\}$. This partition of unity extends the partition of unity $\{\varphi_\beta\Phi_\gamma\}$ on S^n and induces an extension of $\{\varphi_\beta\}$ to a partition of unity defined on $X^n\cup e^{n+1}_\gamma$ and subordinate to $\{U_\alpha\}$. Doing this for all $(n+1)$-cells e^{n+1}_γ gives a partition of unity on X^{n+1}. The local finiteness condition continues to hold since near a point in X^n only the extensions of the φ_β's in the original partition of unity on X^n are nonzero, while in a cell e^{n+1}_γ the only other functions that can be nonzero are the ones coming from $\psi_{\gamma j}$'s. After we make such extensions for all n, we obtain a partition of unity defined on all of X and subordinate to $\{U_\alpha\}$. ◻

Here is a technical fact about paracompact spaces that is occasionally useful:

Lemma 1.21. *Given an open cover $\{U_\alpha\}$ of the paracompact space X, there is a countable open cover $\{V_k\}$ such that each V_k is a disjoint union of open sets each contained in some U_α, and there is a partition of unity $\{\varphi_k\}$ with φ_k supported in V_k.*

Proof: Let $\{\varphi_\beta\}$ be a partition of unity subordinate to $\{U_\alpha\}$. For each finite set S of functions φ_β let V_S be the subset of X where all the φ_β's in S are strictly greater than all the φ_β's not in S. Since only finitely many φ_β's are nonzero near any $x \in X$, V_S is defined by finitely many inequalities among φ_β's near x, so V_S is open. Also, V_S is contained in some U_α, namely, any U_α containing the support of any $\varphi_\beta \in S$, since $\varphi_\beta \in S$ implies $\varphi_\beta > 0$ on V_S. Let V_k be the union of all the open sets V_S such that S has k elements. This is clearly a disjoint union. The collection $\{V_k\}$ is a cover of X since if $x \in X$ then $x \in V_S$ where $S = \{\,\varphi_\beta \mid \varphi_\beta(x) > 0\,\}$.

For the second statement, let $\{\varphi_\gamma\}$ be a partition of unity subordinate to the cover $\{V_k\}$, and let φ_k be the sum of those φ_γ's supported in V_k but not in V_j for $j < k$. □

Exercises

1. Show that the projection $V_n(R^k) \to G_n(\mathbb{R}^k)$ is a fiber bundle with fiber $O(n)$ by showing that it is the orthonormal n-frame bundle associated to the vector bundle $E_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$.

Chapter 2

K-Theory

The idea of K–theory is to make the direct sum operation on real or complex vector bundles over a fixed base space X into the addition operation in a group. There are two slightly different ways of doing this, producing, in the case of complex vector bundles, groups $K(X)$ and $\widetilde{K}(X)$ with $K(X) \approx \widetilde{K}(X) \oplus \mathbb{Z}$, and for real vector bundles, groups $KO(X)$ and $\widetilde{KO}(X)$ with $KO(X) \approx \widetilde{KO}(X) \oplus \mathbb{Z}$. Complex K–theory turns out to be somewhat simpler than real K–theory, so we concentrate on this case in the present chapter.

Computing $\widetilde{K}(X)$ even for simple spaces X requires some work. The case $X = S^n$ is the Bott Periodicity Theorem, which gives isomorphisms $\widetilde{K}(S^n) \approx \widetilde{K}(S^{n+2})$ for all n, and more generally $\widetilde{K}(X) \approx \widetilde{K}(S^2 X)$ where $S^2 X$ is the double suspension of X. This is a deep theorem, so it is not surprising that it has applications of real substance. We give some of these in §2.3, notably:

(1) The nonexistence of division algebras over $\mathbb{R}$ in dimensions other than 1, 2, 4, and 8, the dimensions of the real and complex numbers, quaternions, and Cayley octonions.

(2) The nonparallelizability of spheres other than S^1, S^3, and S^7.

The proof of the Bott Periodicity Theorem divides into two parts. The first is the hard technical work, proving an isomorphism $K(X \times S^2) \approx K(X) \otimes K(S^2)$. This takes about ten pages, forming the bulk of §2.1. The other half is easier, being more formal in nature, and this is contained in §2.2 where the cohomological aspects of $K(X)$ are the main focus.

2.1. The Functor K(X)

Since we shall be dealing exclusively with complex vector bundles in this chapter, let us take 'vector bundle' to mean 'complex vector bundle' unless otherwise specified. Base spaces will always be assumed to be compact Hausdorff, so that the results of the preceding chapter which have this hypothesis will be available to us.

For the purposes of K-theory it is convenient to take a slightly broader definition of 'vector bundle' which allows the fibers of a vector bundle $p:E\to X$ to be vector spaces of different dimensions. We still assume local trivializations of the form $h:p^{-1}(U)\to U\times\mathbb{C}^n$, so the dimensions of fibers must be locally constant over X, but if X is disconnected the dimensions of fibers need not be globally constant.

Consider vector bundles over a fixed base space X. The trivial n-dimensional vector bundle we write as $\varepsilon^n\to X$. Define two vector bundles E_1 and E_2 over X to be *stably isomorphic*, written $E_1\approx_s E_2$, if $E_1\oplus\varepsilon^n\approx E_2\oplus\varepsilon^n$ for some n. In a similar vein we set $E_1\sim E_2$ if $E_1\oplus\varepsilon^m\approx E_2\oplus\varepsilon^n$ for some m and n. It is easy to see that both $\approx_s$ and $\sim$ are equivalence relations. On equivalence classes of either sort the operation of direct sum is well-defined, commutative, and associative. A zero element is the class of ε^0.

Proposition 2.1. *If X is compact Hausdorff, then the set of $\sim$-equivalence classes of vector bundles over X forms an abelian group with respect to $\oplus$.*

This group is called $\widetilde{K}(X)$.

Proof: Only the existence of inverses needs to be shown, which we do by showing that for each vector bundle $\pi:E\to X$ there is a bundle $E'\to X$ such that $E\oplus E'\approx\varepsilon^m$ for some m. If all the fibers of E have the same dimension, this is Proposition 1.4. In the general case let $X_i=\{x\in X\mid \dim\pi^{-1}(x)=i\}$. These X_i's are disjoint open sets in X, hence are finite in number by compactness. By adding to E a bundle which over each X_i is a trivial bundle of suitable dimension we can produce a bundle whose fibers all have the same dimension. □

For the direct sum operation on $\approx_s$-equivalence classes, only the zero element, the class of ε^0, can have an inverse since $E\oplus E'\approx_s\varepsilon^0$ implies $E\oplus E'\oplus\varepsilon^n\approx\varepsilon^n$ for some n, which can only happen if E and E' are 0-dimensional. However, even though inverses do not exist, we do have the cancellation property that $E_1\oplus E_2\approx_s E_1\oplus E_3$ implies $E_2\approx_s E_3$ over a compact base space X, since we can add to both sides of $E_1\oplus E_2\approx_s E_1\oplus E_3$ a bundle E_1' such that $E_1\oplus E_1'\approx\varepsilon^n$ for some n.

Just as the positive rational numbers are constructed from the positive integers by forming quotients a/b with the equivalence relation $a/b=c/d$ iff $ad=bc$, so we can form for compact X an abelian group $K(X)$ consisting of formal differences $E-E'$ of vector bundles E and E' over X, with the equivalence relation $E_1-E_1'=E_2-E_2'$ iff $E_1\oplus E_2'\approx_s E_2\oplus E_1'$. Verifying transitivity of this relation involves the cancellation property, which is why compactness of X is needed. With the obvious addition rule $(E_1 - E_1') + (E_2 - E_2') = (E_1 \oplus E_2) - (E_1' \oplus E_2')$, $K(X)$ is then a group. The zero element is the equivalence class of $E - E$ for any E, and the inverse of $E - E'$ is $E' - E$. Note that every element of $K(X)$ can be represented as a difference $E - \varepsilon^n$ since if we start with $E - E'$ we can add to both E and E' a bundle E'' such that $E' \oplus E'' \approx \varepsilon^n$ for some n.

There is a natural homomorphism $K(X) \to \widetilde{K}(X)$ sending $E - \varepsilon^n$ to the $\sim$-class of E. This is well-defined since if $E - \varepsilon^n = E' - \varepsilon^m$ in $K(X)$, then $E \oplus \varepsilon^m \approx_s E' \oplus \varepsilon^n$, hence $E \sim E'$. The map $K(X) \to \widetilde{K}(X)$ is obviously surjective, and its kernel consists of elements $E - \varepsilon^n$ with $E \sim \varepsilon^0$, hence $E \approx_s \varepsilon^m$ for some m, so the kernel consists of the elements of the form $\varepsilon^m - \varepsilon^n$. This subgroup $\{\varepsilon^m - \varepsilon^n\}$ of $K(X)$ is isomorphic to $\mathbb{Z}$. In fact, restriction of vector bundles to a basepoint $x_0 \in X$ defines a homomorphism $K(X) \to K(x_0) \approx \mathbb{Z}$ which restricts to an isomorphism on the subgroup $\{\varepsilon^m - \varepsilon^n\}$. Thus we have a splitting $K(X) \approx \widetilde{K}(X) \oplus \mathbb{Z}$, depending on the choice of x_0. The group $\widetilde{K}(X)$ is sometimes called *reduced*, to distinguish it from $K(X)$.

Ring Structure

Besides the additive structure in $K(X)$ there is also a natural multiplication coming from tensor product of vector bundles. For elements of $K(X)$ represented by vector bundles E_1 and E_2 their product in $K(X)$ will be represented by the bundle $E_1 \otimes E_2$, so for arbitrary elements of $K(X)$ represented by differences of vector bundles, their product in $K(X)$ is defined by the formula

$$(E_1 - E_1')(E_2 - E_2') = E_1 \otimes E_2 - E_1 \otimes E_2' - E_1' \otimes E_2 + E_1' \otimes E_2'$$

It is routine to verify that this is well-defined and makes $K(X)$ into a commutative ring with identity ε^1, the trivial line bundle, using the basic properties of tensor product of vector bundles described in §1.1. We can simplify notation by writing the element $\varepsilon^n \in K(X)$ just as n. This is consistent with familiar arithmetic rules. For example, the product nE is the sum of n copies of E.

If we choose a basepoint $x_0 \in X$, then the map $K(X) \to K(x_0)$ obtained by restricting vector bundles to their fibers over x_0 is a ring homomorphism. Its kernel, which can be identified with $\widetilde{K}(X)$ as we have seen, is an ideal, hence also a ring in its own right, though not necessarily a ring with identity.

The rings $K(X)$ and $\widetilde{K}(X)$ can be regarded as functors of X. A map $f : X \to Y$ induces a map $f^* : K(Y) \to K(X)$, sending $E - E'$ to $f^*(E) - f^*(E')$. This is a ring homomorphism since $f^*(E_1 \oplus E_2) \approx f^*(E_1) \oplus f^*(E_2)$ and $f^*(E_1 \otimes E_2) \approx f^*(E_1) \otimes f^*(E_2)$. The functor properties $(fg)^* = g^* f^*$ and $\mathbb{1}^* = \mathbb{1}$ as well as the fact that $f \simeq g$ implies $f^* = g^*$ all follow from the corresponding properties for pullbacks of vector bundles. Similarly, we have induced maps $f^* : \widetilde{K}(Y) \to \widetilde{K}(X)$ with the same properties, except that for f^* to be a ring homomorphism we must be in the category of basepointed spaces and basepoint-preserving maps since our definition of multiplication for $\widetilde{K}$ required basepoints.

An *external product* $\mu : K(X)\otimes K(Y)\to K(X\times Y)$ can be defined by $\mu(a\otimes b) = p_1^*(a)p_2^*(b)$ where p_1 and p_2 are the projections of $X\times Y$ onto X and Y. The tensor product of rings is a ring, with multiplication defined by $(a\otimes b)(c\otimes d) = ac\otimes bd$, and μ is a ring homomorphism since $\mu((a\otimes b)(c\otimes d)) = \mu(ac\otimes bd) = p_1^*(ac)p_2^*(bd) = p_1^*(a)p_1^*(c)p_2^*(b)p_2^*(d) = p_1^*(a)p_2^*(b)p_1^*(c)p_2^*(d) = \mu(a\otimes b)\mu(c\otimes d)$.

Taking Y to be S^2 we have an external product

$$\mu : K(X)\otimes K(S^2)\to K(X\times S^2)$$

Most of the remainder of §2.1 will be devoted to showing that this map is an isomorphism. This is the essential core of the proof of Bott Periodicity.

The external product in ordinary cohomology is called 'cross product' and written $a\times b$, but to use this symbol for the K–theory external product might lead to confusion with Cartesian product of vector bundles, which is quite different from tensor product. Instead we will sometimes use the notation $a * b$ as shorthand for $\mu(a\otimes b)$.

The Fundamental Product Theorem

The key result allowing the calculation of $K(X)$ in nontrivial cases is a certain formula that computes $K(X\times S^2)$ in terms of $K(X)$. For example, when X is a point this will yield a calculation of $K(S^2)$. In the next section we will deduce other calculations, in particular $K(S^n)$ for all n.

Let H be the canonical line bundle over $S^2 = \mathbb{C}\mathrm{P}^1$. We showed in Example ?? that $(H\otimes H)\oplus 1 \approx H\oplus H$. In $K(S^2)$ this is the formula $H^2 + 1 = 2H$, so $H^2 = 2H - 1$. We can also write this as $(H-1)^2 = 0$, so we have a natural ring homomorphism $\mathbb{Z}[H]/(H-1)^2\to K(S^2)$ whose domain is the quotient ring of the polynomial ring $\mathbb{Z}[H]$ by the ideal generated by $(H-1)^2$. In particular, note that an additive basis for $\mathbb{Z}[H]/(H-1)^2$ is $\{1,H\}$.

We define a homomorphism μ as the composition

$$\mu : K(X)\otimes\mathbb{Z}[H]/(H-1)^2 \longrightarrow K(X)\otimes K(S^2)\longrightarrow K(X\times S^2)$$

where the second map is the external product.

Theorem 2.2. *The homomorphism $\mu : K(X)\otimes\mathbb{Z}[H]/(H-1)^2\to K(X\times S^2)$ is an isomorphism of rings for all compact Hausdorff spaces X.*

Taking X to be a point we obtain:

Corollary 2.3. *The map $\mathbb{Z}[H]/(H-1)^2\to K(S^2)$ is an isomorphism of rings.*

Thus if we regard $\widetilde{K}(S^2)$ as the kernel of $K(S^2)\to K(x_0)$, then it is generated as an abelian group by $H-1$. Since we have the relation $(H-1)^2 = 0$, this means that the multiplication in $\widetilde{K}(S^2)$ is completely trivial: The product of any two elements is zero. Readers familiar with cup product in ordinary cohomology will recognize that the situation is exactly the same as in $H^*(S^2;\mathbb{Z})$ and $\widetilde{H}^*(S^2;\mathbb{Z})$, with $H-1$ behaving exactly like the generator of $H^2(S^2;\mathbb{Z})$. In the case of ordinary cohomology the cup product of a generator of $H^2(S^2;\mathbb{Z})$ with itself is automatically zero since $H^4(S^2;\mathbb{Z})=0$, whereas with K–theory a calculation is required.

The present section will be devoted entirely to the proof of the preceding theorem. Nothing in the proof will be used elsewhere in the book except in the proof of Bott periodicity for real K–theory in §2.4, so the reader who wishes to defer a careful reading of the proof may skip ahead to §2.2 without any loss of continuity.

The main work in proving the theorem will be to prove the surjectivity of μ. Injectivity will then be proved by a closer examination of the surjectivity argument.

Clutching Functions

From the classification of vector bundles over spheres in §1.2 we know that vector bundles over S^2 correspond exactly to homotopy classes of maps $S^1 \to GL_n(\mathbb{C})$, which we called clutching functions. To prove the theorem we will generalize this construction, creating vector bundles over $X\times S^2$ by gluing together two vector bundles over $X\times D^2$ by means of a generalized clutching function.

We begin by describing this more general clutching construction. Given a vector bundle $p:E\to X$, let $f:E\times S^1\to E\times S^1$ be an automorphism of the product vector bundle $p\times \mathbb{1}:E\times S^1\to X\times S^1$. Thus for each $x\in X$ and $z\in S^1$, f specifies an isomorphism $f(x,z):p^{-1}(x)\to p^{-1}(x)$. From E and f we construct a vector bundle over $X\times S^2$ by taking two copies of $E\times D^2$ and identifying the subspaces $E\times S^1$ via f. We write this bundle as $[E,f]$, and call f a *clutching function* for $[E,f]$. If $f_t:E\times S^1\to E\times S^1$ is a homotopy of clutching functions, then $[E,f_0]\approx[E,f_1]$ since from the homotopy f_t we can construct a vector bundle over $X\times S^2\times I$ restricting to $[E,f_0]$ and $[E,f_1]$ over $X\times S^2\times\{0\}$ and $X\times S^2\times\{1\}$. From the definitions it is evident that $[E_1,f_1]\oplus[E_2,f_2]\approx[E_1\oplus E_2,f_1\oplus f_2]$.

Here are some examples of bundles built using clutching functions:

1. $[E,\mathbb{1}]$ is the external product $E*1=\mu(E,1)$, or equivalently the pullback of E via the projection $X\times S^2\to X$.

2. Taking X to be a point, then we showed in Example 1.10 that $[1,z]\approx H$ where '1' is the trivial line bundle over X, 'z' means scalar multiplication by $z\in S^1\subset\mathbb{C}$, and H is the canonical line bundle over $S^2=\mathbb{C}\mathrm{P}^1$. More generally we have $[1,z^n]\approx H^n$, the n-fold tensor product of H with itself. The formula $[1,z^n]\approx H^n$ holds also for negative n if we define $H^{-1}=[1,z^{-1}]$, which is justified by the fact that $H\otimes H^{-1}\approx 1$.

3. $[E,z^n]\approx E*H^n=\mu(E,H^n)$ for $n\in\mathbb{Z}$.

4. Generalizing this, $[E,z^nf]\approx[E,f]\otimes\widehat{H}^n$ where $\widehat{H}^n$ denotes the pullback of H^n via the projection $X\times S^2\to S^2$.

Every vector bundle $E' \to X \times S^2$ is isomorphic to $[E, f]$ for some E and f. To see this, let the unit circle $S^1 \subset \mathbb{C} \cup \{\infty\} = S^2$ decompose S^2 into the two disks D_0 and D_∞, and let E_α for $\alpha = 0, \infty$ be the restriction of E' over $X \times D_\alpha$, with E the restriction of E' over $X \times \{1\}$. The projection $X \times D_\alpha \to X \times \{1\}$ is homotopic to the identity map of $X \times D_\alpha$, so the bundle E_α is isomorphic to the pullback of E by the projection, and this pullback is $E \times D_\alpha$, so we have an isomorphism $h_\alpha : E_\alpha \to E \times D_\alpha$. Then $f = h_0 h_\infty^{-1}$ is a clutching function for E'.

We may assume a clutching function f is normalized to be the identity over $X \times \{1\}$ since we may normalize any isomorphism $h_\alpha : E_\alpha \to E \times D_\alpha$ by composing it over each $X \times \{z\}$ with the inverse of its restriction over $X \times \{1\}$. Any two choices of normalized h_α are homotopic through normalized h_α's since they differ by a map g_α from D_α to the automorphisms of E, with $g_\alpha(1) = \mathbb{1}$, and such a g_α is homotopic to the constant map $\mathbb{1}$ by composing it with a deformation retraction of D_α to 1. Thus any two choices f_0 and f_1 of normalized clutching functions are joined by a homotopy of normalized clutching functions f_t.

The strategy of the proof will be to reduce from arbitrary clutching functions to successively simpler clutching functions. The first step is to reduce to *Laurent polynomial* clutching functions, which have the form $\ell(x, z) = \sum_{|i| \le n} a_i(x) z^i$ where $a_i : E \to E$ restricts to a linear transformation $a_i(x)$ in each fiber $p^{-1}(x)$. We call such an a_i an *endomorphism* of E since the linear transformations $a_i(x)$ need not be invertible, though their linear combination $\sum_i a_i(x) z^i$ is since clutching functions are automorphisms.

Proposition 2.4. *Every vector bundle $[E, f]$ is isomorphic to $[E, \ell]$ for some Laurent polynomial clutching function ℓ. Laurent polynomial clutching functions ℓ_0 and ℓ_1 which are homotopic through clutching functions are homotopic by a Laurent polynomial clutching function homotopy $\ell_t(x, z) = \sum_i a_i(x, t) z^i$.*

Before beginning the proof we need a lemma. For a compact space X we wish to approximate a continuous function $f : X \times S^1 \to \mathbb{C}$ by Laurent polynomial functions $\sum_{|n| \le N} a_n(x) z^n = \sum_{|n| \le N} a_n(x) e^{in\theta}$, where each a_n is a continuous function $X \to \mathbb{C}$. Motivated by Fourier series, we set

$$a_n(x) = \frac{1}{2\pi} \int_{S^1} f(x, \theta) e^{-in\theta} \, d\theta$$

For positive real r let $u(x, r, \theta) = \sum_{n \in \mathbb{Z}} a_n(x) r^{|n|} e^{in\theta}$. For fixed $r < 1$, this series converges absolutely and uniformly as (x, θ) ranges over $X \times S^1$, by comparison with the geometric series $\sum_n r^n$, since compactness of $X \times S^1$ implies that $|f(x, \theta)|$ is bounded and hence also $|a_n(x)|$. If we can show that $u(x, r, \theta)$ approaches $f(x, \theta)$ uniformly in x and θ as r goes to 1, then sums of finitely many terms in the series for $u(x, r, \theta)$ with r near 1 will give the desired approximations to f by Laurent polynomial functions.

Lemma 2.5. *As* $r \to 1$, $u(x,r,\theta) \to f(x,\theta)$ *uniformly in* x *and* θ.

Proof: For $r < 1$ we have

$$u(x,r,\theta) = \sum_{n=-\infty}^{\infty} \frac{1}{2\pi} \int_{S^1} r^{|n|} e^{in(\theta-t)} f(x,t)\,dt = \int_{S^1} \frac{1}{2\pi} \sum_{n=-\infty}^{\infty} r^{|n|} e^{in(\theta-t)} f(x,t)\,dt$$

where the order of summation and integration can be interchanged since the series in the latter formula converges uniformly, by comparison with the geometric series $\sum_n r^n$. Define the Poisson kernel

$$P(r,\varphi) = \frac{1}{2\pi} \sum_{n=-\infty}^{\infty} r^{|n|} e^{in\varphi} \qquad \text{for } 0 \le r < 1 \text{ and } \varphi \in \mathbb{R}$$

Then $u(x,r,\theta) = \int_{S^1} P(r,\theta-t) f(x,t)\,dt$. By summing the two geometric series for positive and negative n in the formula for $P(r,\varphi)$, one computes that

$$P(r,\varphi) = \frac{1}{2\pi} \cdot \frac{1-r^2}{1-2r\cos\varphi + r^2}$$

Three basic facts about $P(r,\varphi)$ which we shall need are:

(a) As a function of φ, $P(r,\varphi)$ is even, of period 2π, and monotone decreasing on $[0,\pi]$, since the same is true of $\cos\varphi$ which appears in the denominator of $P(r,\varphi)$ with a minus sign. In particular we have $P(r,\varphi) \ge P(r,\pi) > 0$ for all $r < 1$.

(b) $\int_{S^1} P(r,\varphi)\,d\varphi = 1$ for each $r < 1$, as one sees by integrating the series for $P(r,\varphi)$ term by term.

(c) For fixed $\varphi \in (0,\pi)$, $P(r,\varphi) \to 0$ as $r \to 1$ since the numerator of $P(r,\varphi)$ approaches 0 and the denominator approaches $2 - 2\cos\varphi \ne 0$.

Now to show uniform convergence of $u(x,r,\theta)$ to $f(x,\theta)$ we first observe that, using (b), we have

$$|u(x,r,\theta) - f(x,\theta)| = \Big| \int_{S^1} P(r,\theta-t) f(x,t)\,dt - \int_{S^1} P(r,\theta-t) f(x,\theta)\,dt \Big| \le \int_{S^1} P(r,\theta-t)\,|f(x,t) - f(x,\theta)|\,dt$$

Given $\varepsilon > 0$, there exists a $\delta > 0$ such that $|f(x,t) - f(x,\theta)| < \varepsilon$ for $|t-\theta| < \delta$ and all x, since f is uniformly continuous on the compact space $X \times S^1$. Let I_δ denote the integral $\int P(r,\theta-t)\,|f(x,t)-f(x,\theta)|\,dt$ over the interval $|t-\theta| \le \delta$ and let I'_δ denote this integral over the rest of S^1. Then we have

$$I_\delta \le \int_{|t-\theta|\le\delta} P(r,\theta-t)\,\varepsilon\,dt \le \varepsilon \int_{S^1} P(r,\theta-t)\,dt = \varepsilon$$

By (a) the maximum value of $P(r,\theta-t)$ on $|t-\theta| \ge \delta$ is $P(r,\delta)$. So

$$I'_\delta \le P(r,\delta) \int_{S^1} |f(x,t) - f(x,\theta)|\,dt$$

The integral here has a uniform bound for all x and θ since f is bounded. Thus by (c) we can make $I'_\delta \le \varepsilon$ by taking r close enough to 1. Therefore $|u(x,r,\theta) - f(x,\theta)| \le I_\delta + I'_\delta \le 2\varepsilon$. □

Proof of Proposition 2.4: Choosing a Hermitian inner product on E, the endomorphisms of $E\times S^1$ form a vector space $\mathrm{End}(E\times S^1)$ with a norm $\|\alpha\| = \sup_{|v|=1}|\alpha(v)|$. The triangle inequality holds for this norm, so balls in $\mathrm{End}(E\times S^1)$ are convex. The subspace $\mathrm{Aut}(E\times S^1)$ of automorphisms is open in the topology defined by this norm since it is the preimage of $(0,\infty)$ under the continuous map $\mathrm{End}(E\times S^1)\to[0,\infty)$, $\alpha \mapsto \inf_{(x,z)\in X\times S^1}|\det(\alpha(x,z))|$. Thus to prove the first statement of the lemma it will suffice to show that Laurent polynomials are dense in $\mathrm{End}(E\times S^1)$, since a sufficiently close Laurent polynomial approximation ℓ to f will then be homotopic to f via the linear homotopy $t\ell + (1-t)f$ through clutching functions. The second statement follows similarly by approximating a homotopy from ℓ_0 to ℓ_1, viewed as an automorphism of $E\times S^1\times I$, by a Laurent polynomial homotopy ℓ'_t, then combining this with linear homotopies from ℓ_0 to ℓ'_0 and ℓ_1 to ℓ'_1 to obtain a homotopy ℓ_t from ℓ_0 to ℓ_1.

To show that every $f \in \mathrm{End}(E\times S^1)$ can be approximated by Laurent polynomial endomorphisms, first choose open sets U_i covering X together with isomorphisms $h_i : p^{-1}(U_i)\to U_i\times\mathbb{C}^{n_i}$. We may assume h_i takes the chosen inner product in $p^{-1}(U_i)$ to the standard inner product in $\mathbb{C}^{n_i}$, by applying the Gram-Schmidt process to h_i^{-1} of the standard basis vectors. Let $\{\varphi_i\}$ be a partition of unity subordinate to $\{U_i\}$ and let X_i be the support of φ_i, a compact set in U_i. Via h_i, the linear maps $f(x,z)$ for $x \in X_i$ can be viewed as matrices. The entries of these matrices define functions $X_i\times S^1\to\mathbb{C}$. By the lemma we can find Laurent polynomial matrices $\ell_i(x,z)$ whose entries uniformly approximate those of $f(x,z)$ for $x \in X_i$. It follows easily that ℓ_i approximates f in the $\|\cdot\|$ norm. From the Laurent polynomial approximations ℓ_i over X_i we form the convex linear combination $\ell = \sum_i \varphi_i\ell_i$, a Laurent polynomial approximating f over all of $X\times S^1$. □

A Laurent polynomial clutching function can be written $\ell = z^{-m}q$ for a polynomial clutching function q, and then we have $[E,\ell] \approx [E,q]\otimes\widehat{H}^{-m}$. The next step is to reduce polynomial clutching functions to linear clutching functions.

Proposition 2.6. *If q is a polynomial clutching function of degree at most n, then $[E,q]\oplus[nE,\mathbb{1}] \approx [(n+1)E, L^nq]$ for a linear clutching function L^nq.*

Proof: Let $q(x,z) = a_n(x)z^n + \cdots + a_0(x)$. Each of the matrices

$$A = \begin{pmatrix} 1 & -z & 0 & \cdots & 0 & 0 \\ 0 & 1 & -z & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -z \\ a_n & a_{n-1} & a_{n-2} & \cdots & a_1 & a_0 \end{pmatrix} \qquad B = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & q \end{pmatrix}$$

defines an endomorphism of $(n+1)E$ by interpreting the (i,j) entry of the matrix as a linear map from the j^{th} summand of $(n+1)E$ to the i^{th} summand, with the entries 1 denoting the identity $E \to E$ and z denoting z times the identity, for $z \in S^1$. We can pass from the matrix A to the matrix B by a sequence of elementary row and column operations in the following way. In A, add z times the first column to the second column, then z times the second column to the third, and so on. This produces 0's above the diagonal and the polynomial q in the lower right corner. Then for each $i \le n$, subtract the appropriate multiple of the i^{th} row from the last row to make all the entries in the last row 0 except for the final q. These row and column operations are not quite elementary row and column operations in the traditional sense since the entries of the matrices are not numbers but linear maps. However, by restricting to a fiber of E and choosing a basis in this fiber, each entry in A becomes a matrix itself, and then each of the preceding row and column operations can be realized by a sequence of traditional row and column operations on the expanded matrices.

The matrix B is a clutching function for $[nE, \mathbb{1}] \oplus [E, q]$, hence in each fiber the expanded version of B has nonzero determinant. Since elementary row and column operations preserve determinant, the expanded version of A is also invertible in each fiber. This means that A is an automorphism of $(n+1)E$ for each $z \in S^1$, and therefore determines a clutching function which we denote by $L^n q$. Since $L^n q$ has the form $A(x)z + B(x)$, it is a linear clutching function. The matrices A and B define homotopic clutching functions since the elementary row and column operations can be achieved by continuous one-parameter families of such operations. For example the first operation can be achieved by adding tz times the first column to the second, with t ranging from 0 to 1. Since homotopic clutching functions produce isomorphic bundles, we obtain an isomorphism $[E, q] \oplus [nE, \mathbb{1}] \approx [(n+1)E, L^n q]$. □

Linear Clutching Functions

For linear clutching functions $a(x)z + b(x)$ we have the following key fact:

Proposition 2.7. *Given a bundle $[E, a(x)z + b(x)]$, there is a splitting $E \approx E_+ \oplus E_-$ with $[E, a(x)z + b(x)] \approx [E_+, \mathbb{1}] \oplus [E_-, z]$.*

Proof: The first step is to reduce to the case that $a(x)$ is the identity for all x. Consider the expression:

$$(*) \qquad (1+tz)\left[a(x)\frac{z+t}{1+tz} + b(x)\right] = \left[a(x) + tb(x)\right]z + ta(x) + b(x)$$

When $t = 0$ this equals $a(x)z + b(x)$. For $0 \le t < 1$, $(*)$ defines an invertible linear transformation since the left-hand side is obtained from $a(x)z + b(x)$ by first applying the substitution $z \mapsto (z+t)/(1+tz)$ which takes S^1 to itself (because if $|z| = 1$ then $|(z+t)/(1+tz)| = |\overline{z}(z+t)/(1+tz)| = |(1+t\overline{z})/(1+tz)| = |\overline{w}/w| = 1$), and then multiplying by the nonzero scalar $1+tz$. Therefore $(*)$ defines a homotopy of clutching functions as t goes from 0 to $t_0 < 1$. In the right-hand side of $(*)$ the term $a(x) + tb(x)$ is invertible for $t = 1$ since it is the restriction of $a(x)z + b(x)$ to $z = 1$. Therefore $a(x) + tb(x)$ is invertible for $t = t_0$ near 1, as the continuous function $t \mapsto \inf_{x\in X}|\det[a(x) + tb(x)]|$ is nonzero for $t = 1$, hence also for t near 1. Now we use the simple fact that $[E, fg] \approx [E, f]$ for any isomorphism $g : E \to E$. This allows us to replace the clutching function on the right-hand side of $(*)$ by the clutching function $z + [t_0 a(x) + b(x)][a(x) + t_0 b(x)]^{-1}$, reducing to the case of clutching functions of the form $z + b(x)$.

Since $z + b(x)$ is invertible for all x, $b(x)$ has no eigenvalues on the unit circle S^1.

Lemma 2.8. *Let $b : E \to E$ be an endomorphism having no eigenvalues on the unit circle S^1. Then there are unique subbundles E_+ and E_- of E such that*:

(a) $E = E_+ \oplus E_-$.

(b) $b(E_\pm) \subset E_\pm$.

(c) *The eigenvalues of $b|E_+$ all lie outside S^1 and the eigenvalues of $b|E_-$ all lie inside S^1.*

Proof: Consider first the algebraic situation of a linear transformation $T : V \to V$ with characteristic polynomial $q(t)$. Assuming $q(t)$ has no roots on S^1, we may factor $q(t)$ in $\mathbb{C}[t]$ as $q_+(t)q_-(t)$ where $q_+(t)$ has all its roots outside S^1 and $q_-(t)$ has all its roots inside S^1. Let $V_\pm$ be the kernel of $q_\pm(T) : V \to V$. Since q_+ and q_- are relatively prime in $\mathbb{C}[t]$, there are polynomials r and s with $rq_+ + sq_- = 1$. From $q_+(T)q_-(T) = q(T) = 0$, we have $\operatorname{Im} q_-(T) \subset \operatorname{Ker} q_+(T)$, and the opposite inclusion follows from $r(T)q_+(T) + q_-(T)s(T) = \mathbb{1}$. Thus $\operatorname{Ker} q_+(T) = \operatorname{Im} q_-(T)$, and similarly $\operatorname{Ker} q_-(T) = \operatorname{Im} q_+(T)$. From $q_+(T)r(T) + q_-(T)s(T) = \mathbb{1}$ we see that $\operatorname{Im} q_+(T) + \operatorname{Im} q_-(T) = V$, and from $r(T)q_+(T) + s(T)q_-(T) = \mathbb{1}$ we deduce that $\operatorname{Ker} q_+(T) \cap \operatorname{Ker} q_-(T) = 0$. Hence $V = V_+ \oplus V_-$. We have $T(V_\pm) \subset V_\pm$ since $q_\pm(T)(v) = 0$ implies $q_\pm(T)(T(v)) = T(q_\pm(T)(v)) = 0$. All eigenvalues of $T|V_\pm$ are roots of $q_\pm$ since $q_\pm(T) = 0$ on $V_\pm$. Thus conditions (a)–(c) hold for V_+ and V_-.

To see the uniqueness of V_+ and V_- satisfying (a)–(c), let $q'_\pm$ be the characteristic polynomial of $T|V_\pm$, so $q = q'_+ q'_-$. All the linear factors of $q'_\pm$ must be factors of $q_\pm$ by condition (c), so the factorizations $q = q'_+ q'_-$ and $q = q_+ q_-$ must coincide up to scalar factors. Since $q'_\pm(T)$ is identically zero on $V_\pm$, so must be $q_\pm(T)$, hence $V_\pm \subset \operatorname{Ker} q_\pm(T)$. Since $V = V_+ \oplus V_-$ and $V = \operatorname{Ker} q_+(T) \oplus \operatorname{Ker} q_-(T)$, we must have $V_\pm = \operatorname{Ker} q_\pm(T)$. This establishes the uniqueness of $V_\pm$.

As T varies continuously through linear transformations without eigenvalues on S^1, its characteristic polynomial $q(t)$ varies continuously through polynomials without roots in S^1. In this situation we assert that the factors $q_\pm$ of q vary continuously with q, assuming that q, q_+, and q_- are normalized to be monic polynomials. To see this we shall use the fact that for any circle C in $\mathbb{C}$ disjoint from the roots of q, the number of roots of q inside C, counted with multiplicity, equals the degree of the map $\gamma : C \to S^1$, $\gamma(z) = q(z)/|q(z)|$. To prove this fact it suffices to consider the case of a small circle C about a root $z = a$ of multiplicity m, so $q(t) = p(t)(t-a)^m$ with $p(a) \neq 0$. The homotopy

$$\gamma_s(z) = \frac{p(sa + (1-s)z)(z-a)^m}{|p(sa + (1-s)z)(z-a)^m|}$$

gives a reduction to the case $(t-a)^m$, where it is clear that the degree is m.

Thus for a small circle C about a root $z = a$ of q of multiplicity m, small perturbations of q produce polynomials q' which also have m roots $a_1, \cdots, a_m$ inside C, so the factor $(z-a)^m$ of q becomes a factor $(z-a_1)\cdots(z-a_m)$ of the nearby q'. Since the a_i's are near a, these factors of q and q' are close, and so $q'_\pm$ is close to $q_\pm$.

Next we observe that as T varies continuously through transformations without eigenvalues in S^1, the splitting $V = V_+ \oplus V_-$ also varies continuously. To see this, recall that $V_+ = \operatorname{Im} q_-(T)$ and $V_- = \operatorname{Im} q_+(T)$. Choose a basis $v_1, \cdots, v_n$ for V such that $q_-(T)(v_1), \cdots, q_-(T)(v_k)$ is a basis for V_+ and $q_+(T)(v_{k+1}), \cdots, q_+(T)(v_n)$ is a basis for V_-. For nearby T these vectors vary continuously, hence remain independent. Thus the splitting $V = \operatorname{Im} q_-(T) \oplus \operatorname{Im} q_+(T)$ continues to hold for nearby T, and so the splitting $V = V_+ \oplus V_-$ varies continuously with T.

It follows that the union $E_\pm$ of the subspaces $V_\pm$ in all the fibers V of E is a subbundle, and so the proof of the lemma is complete. □

To finish the proof of Proposition 2.7, note that the lemma gives a splitting $[E, z + b(x)] \approx [E_+, z + b_+(x)] \oplus [E_-, z + b_-(x)]$ where b_+ and b_- are the restrictions of b. Since $b_+(x)$ has all its eigenvalues outside S^1, the formula $tz + b_+(x)$ for $0 \leq t \leq 1$ defines a homotopy of clutching functions from $z + b_+(x)$ to $b_+(x)$. Hence $[E_+, z + b_+(x)] \approx [E_+, b_+(x)] \approx [E_+, \mathbb{1}]$. Similarly, $z + tb_-(x)$ defines a homotopy of clutching functions from $z + b_-(x)$ to z, so $[E_-, z + b_-(x)] \approx [E_-, z]$. □

For future reference we note that the splitting $[E, az + b] \approx [E_+, \mathbb{1}] \oplus [E_-, z]$ constructed in the proof of Proposition 2.7 preserves direct sums, in the sense that the splitting for a sum $[E_1 \oplus E_2, (a_1 z + b_1) \oplus (a_2 z + b_2)]$ has $(E_1 \oplus E_2)_\pm = (E_1)_\pm \oplus (E_2)_\pm$. This is because the first step of reducing to the case $a = \mathbb{1}$ clearly respects sums, and the uniqueness of the $\pm$-splitting in Lemma 2.8 guarantees that it preserves sums.

Conclusion of the Proof

Now we address the question of showing that the homomorphism

$$\mu: K(X)\otimes\mathbb{Z}[H]/(H-1)^2 \to K(X\times S^2$$

is an isomorphism. The preceding propositions imply that in $K(X\times S^2)$ we have

$$\begin{aligned}[E,f] &= [E,z^{-m}q]\\ &= [E,q]\otimes\widehat{H}^{-m}\\ &= [(n+1)E,L^nq]\otimes\widehat{H}^{-m} - [nE,\mathbb{1}]\otimes\widehat{H}^{-m}\\ &= [((n+1)E)_+,\mathbb{1}]\otimes\widehat{H}^{-m} + [((n+1)E)_-,z]\otimes\widehat{H}^{-m} - [nE,\mathbb{1}]\otimes\widehat{H}^{-m}\\ &= ((n+1)E)_+\otimes H^{-m} + ((n+1)E)_-\otimes H^{1-m} - nE\otimes H^{-m}\end{aligned}$$

This last expression is in the image of $\mu: K(X)\otimes K(S^2)\to K(X\times S^2)$. Since every vector bundle over $X\times S^2$ has the form $[E,f]$, it follows that μ is surjective.

To show μ is injective we shall construct $\nu: K(X\times S^2)\to K(X)\otimes\mathbb{Z}[H]/(H-1)^2$ such that $\nu\mu = \mathbb{1}$. The idea will be to define $\nu([E,f])$ as some linear combination of terms $E\otimes H^k$ and $((n+1)E)_\pm\otimes H^k$ which is independent of all choices.

To investigate the dependence of the terms in the formula for $[E,f]$ displayed above on m and n we first derive the following two formulas, where $\deg q\le n$:

(1) $[(n+2)E,L^{n+1}q]\approx[(n+1)E,L^nq]\oplus[E,\mathbb{1}]$

(2) $[(n+2)E,L^{n+1}(zq)]\approx[(n+1)E,L^nq]\oplus[E,z]$

The matrix representations of $L^{n+1}q$ and $L^{n+1}(zq)$ are:

$$\begin{pmatrix}1 & -z & 0 & \cdots & 0\\ 0 & 1 & -z & \cdots & 0\\ \vdots & \vdots & \vdots & & \vdots\\ 0 & 0 & 0 & 1 & -z\\ 0 & a_n & a_{n-1} & \cdots & a_0\end{pmatrix} \quad\text{and}\quad \begin{pmatrix}1 & -z & 0 & \cdots & 0 & 0\\ 0 & 1 & -z & \cdots & 0 & 0\\ \vdots & \vdots & \vdots & & \vdots & \vdots\\ 0 & 0 & 0 & \cdots & 1 & -z\\ a_n & a_{n-1} & a_{n-2} & \cdots & a_0 & 0\end{pmatrix}$$

In the first matrix we can add z times the first column to the second column to eliminate the $-z$ in the first row, and then the first row and column give the summand $[E,\mathbb{1}]$ while the rest of the matrix gives $[(n+1)E,L^nq]$. This proves (1). Similarly, in the second matrix we add z^{-1} times the last column to the next-to-last column to make the $-z$ in the last column have all zeros in its row and column, which gives the splitting in (2) since $[E,-z]\approx[E,z]$, the clutching function $-z$ being the composition of the clutching function z with the automorphism -1 of E.

In view of the appearance of the correction terms $[E,\mathbb{1}]$ and $[E,z]$ in (1) and (2), it will be useful to know the '$\pm$' splittings for these two bundles:

(3) For $[E,\mathbb{1}]$ the summand E_- is 0 and $E_+ = E$.

(4) For $[E,z]$ the summand E_+ is 0 and $E_- = E$.

Statement (4) is obvious from the definitions since the clutching function z is already in the form $z+b(x)$ with $b(x)=0$, so 0 is the only eigenvalue of $b(x)$ and hence

$E_+ = 0$. To obtain (3) we first apply the procedure at the beginning of the proof of Proposition 2.7 which replaces a clutching function $a(x)z + b(x)$ by the clutching function $z + [t_0a(x) + b(x)][a(x) + t_0b(x)]^{-1}$ with $0 < t_0 < 1$. Specializing to the case $a(x)z + b(x) = \mathbb{1}$ this yields $z + t_0^{-1}\mathbb{1}$. Since $t_0^{-1}\mathbb{1}$ has only the one eigenvalue $t_0^{-1} > 1$, we have $E_- = 0$.

Formulas (1) and (3) give $((n + 2)E)_- \approx ((n + 1)E)_-$, using the fact that the $\pm$-splitting preserves direct sums. So the 'minus' summand is independent of n.

Suppose we define

$$\nu([E, z^{-m}q]) = ((n+1)E)_- \otimes (H-1) + E \otimes H^{-m} \in K(X) \otimes \mathbb{Z}[H]/(H-1)^2$$

for $n \ge \deg q$. We claim that this is well-defined. We have just noted that 'minus' summands are independent of n, so $\nu([E, z^{-m}q])$ does not depend on n. To see that it is independent of m we must see that it is unchanged when $z^{-m}q$ is replaced by $z^{-m-1}(zq)$. By (2) and (4) we have the first of the following equalities:

$$\begin{aligned} \nu([E, z^{-m-1}(zq)]) &= ((n+1)E)_- \otimes (H-1) + E \otimes (H-1) + E \otimes H^{-m-1} \\ &= ((n+1)E)_- \otimes (H-1) + E \otimes (H^{-m} - H^{-m-1}) + E \otimes H^{-m-1} \\ &= ((n+1)E)_- \otimes (H-1) + E \otimes H^{-m} \\ &= \nu([E, z^{-m}q]) \end{aligned}$$

To obtain the second equality we use the relation $(H-1)^2 = 0$ which implies $H(H - 1) = H - 1$ and hence $H - 1 = H^{-m} - H^{-m-1}$ for all m. The third and fourth equalities are evident.

Another choice which might perhaps affect the value of $\nu([E, z^{-m}q])$ is the constant $t_0 < 1$ in the proof of Proposition 2.7. This could be any number sufficiently close to 1, so varying t_0 gives a homotopy of the endomorphism b in Lemma 2.8. This has no effect on the $\pm$-splitting since we can apply Lemma 2.8 to the endomorphism of $E \times I$ given by the homotopy. Hence the choice of t_0 does not affect $\nu([E, z^{-m}q])$.

It remains to see that $\nu([E, z^{-m}q])$ depends only on the bundle $[E, z^{-m}q]$, not on the clutching function $z^{-m}q$ for this bundle. We showed that every bundle over $X \times S^2$ has the form $[E, f]$ for a normalized clutching function f which was unique up to homotopy, and in Proposition 2.5 we showed that Laurent polynomial approximations to homotopic f's are Laurent-polynomial-homotopic. If we apply Propositions 2.6 and 2.7 over $X \times I$ with a Laurent polynomial homotopy as clutching function, we conclude that the two bundles $((n+1)E)_-$ over $X \times \{0\}$ and $X \times \{1\}$ are isomorphic. This finishes the verification that $\nu([E, z^{-m}q])$ is well-defined.

It is easy to check through the definitions to see that ν takes sums to sums since $L^n(q_1 \oplus q_2) = L^n q_1 \oplus L^n q_2$ and, as previously noted, the $\pm$-splitting in Proposition 2.7 preserves sums. So ν extends to a homomorphism $K(X \times S^2) \to K(X) \otimes \mathbb{Z}[H]/(H - 1)^2$.

The last thing to verify is that $\nu\mu = \mathbb{1}$. The group $\mathbb{Z}[H]/(H-1)^2$ is generated by 1 and H, so in view of the relation $H + H^{-1} = 2$, which follows from $(H-1)^2 = 0$, we see that $K(S^2)$ is also generated by 1 and H^{-1}. Thus it suffices to show $\nu\mu = 1\!\!1$ on elements $E\otimes H^{-m}$ for $m \ge 0$. We have $\nu\mu(E\otimes H^{-m}) = \nu([E, z^{-m}]) = E_-\otimes(H - 1) + E\otimes H^{-m} = E\otimes H^{-m}$ since $E_- = 0$, the polynomial q being $1\!\!1$ so that (3) applies.

This completes the proof of Theorem 2.2. □

2.2 Bott Periodicity

In this section we develop a few basic tools that make it possible to compute $K(X)$ for more complicated spaces X, and in particular for all spheres, yielding the Bott Periodicity Theorem. These tools are formally very similar to the basic machinery of algebraic topology involving cohomology, and in fact they make K-theory into what was classically called a generalized cohomology theory, but now is called simply a cohomology theory.

Exact Sequences

We will be going back and forth between the two versions of K-theory, $K(X)$ and $\widetilde{K}(X)$. Sometimes one is more convenient for stating a result, sometimes the other. We begin by examining a key property of the reduced groups $\widetilde{K}(X)$.

Proposition 2.9. *If X is compact Hausdorff and $A\subset X$ is a closed subspace, then the inclusion and quotient maps $A \xrightarrow{i} X \xrightarrow{q} X/A$ induce homomorphisms $\widetilde{K}(X/A) \xrightarrow{q^*} \widetilde{K}(X) \xrightarrow{i^*} \widetilde{K}(A)$ for which the kernel of i^* equals the image of q^*.*

Since we assume A is a closed subspace of a compact Hausdorff space, it is also compact Hausdorff. The quotient space X/A is compact Hausdorff as well, with the Hausdorff property following from the fact that compact Hausdorff spaces are normal, hence a point $x\in X - A$ and the subspace A have disjoint neighborhoods in X, projecting to disjoint neighborhoods of x and the point A/A in X/A.

There is some standard terminology to described the conclusion of the proposition. A sequence of homomorphisms $G_1 \to G_2 \to \cdots \to G_n$ is said to be *exact* if at each intermediate group G_i the kernel of the outgoing map equals the image of incoming map. Thus the proposition states that $\widetilde{K}(X/A) \xrightarrow{q^*} \widetilde{K}(X) \xrightarrow{i^*} \widetilde{K}(A)$ is exact. We will be deriving some longer exact sequences below.

Proof: The inclusion $\operatorname{Im} q^* \subset \operatorname{Ker} i^*$ is equivalent to $i^*q^* = 0$. Since qi is equal to the composition $A \to A/A \hookrightarrow X/A$ and $\widetilde{K}(A/A) = 0$, it follows that $i^*q^* = 0$.

For the opposite inclusion $\operatorname{Ker} i^* \subset \operatorname{Im} q^*$, suppose the restriction over A of a vector bundle $p: E\to X$ is stably trivial. Adding a trivial bundle to E, we may assume that E itself is trivial over A. Choosing a trivialization $h: p^{-1}(A) \to A\times\mathbb{C}^n$, let E/h be the quotient space of E under the identifications $h^{-1}(x,v) \sim h^{-1}(y,v)$ for $x, y \in A$.

There is then an induced projection $E/h \to X/A$. To see that this is a vector bundle we need to find a local trivialization over a neighborhood of the point A/A.

We claim that since E is trivial over A, it is trivial over some neighborhood of A. In many cases this holds because there is a neighborhood which deformation retracts onto A, so the restriction of E over this neighborhood is trivial since it is isomorphic to the pullback of $p^{-1}(A)$ via the retraction. In the absence of such a deformation retraction one can make the following more complicated argument. A trivialization of E over A determines sections $s_i : A \to E$ which form a basis in each fiber over A. Choose a cover of A by open sets U_j in X over each of which E is trivial. Via a local trivialization, each section s_i can be regarded as a map from $A \cap U_j$ to a single fiber, so by the Tietze extension theorem we obtain a section $s_{ij} : U_j \to E$ extending s_i. If $\{\varphi_j, \varphi\}$ is a partition of unity subordinate to the cover $\{U_j, X - A\}$ of X, the sum $\sum_j \varphi_j s_{ij}$ gives an extension of s_i to a section defined on all of X. Since these sections form a basis in each fiber over A, they must form a basis in all nearby fibers. Namely, over U_j the extended s_i's can be viewed as a square-matrix-valued function having nonzero determinant at each point of A, hence at nearby points as well.

Thus we have a trivialization h of E over a neighborhood U of A. This induces a trivialization of E/h over U/A, so E/h is a vector bundle. It remains only to verify that $E \approx q^*(E/h)$. In the commutative diagram at the right the quotient map $E \to E/h$ is an isomorphism on fibers, so this map and p give an isomorphism $E \approx q^*(E/h)$. □

$$\begin{array}{ccc} E & \longrightarrow & E/h \\ {\scriptstyle p}\downarrow & & \downarrow \\ X & \xrightarrow{\;q\;} & X/A \end{array}$$

There is an easy way to extend the exact sequence $\widetilde{K}(X/A) \to \widetilde{K}(X) \to \widetilde{K}(A)$ to the left, using the following diagram, where C and S denote cone and suspension:

$$\begin{array}{ccccccccc} A & \hookrightarrow & X & \hookrightarrow & X \cup CA & \hookrightarrow & (X \cup CA) \cup CX & \hookrightarrow & \big((X \cup CA) \cup CX\big) \cup C(X \cup CA) \\ & & & & \downarrow & & \downarrow & & \downarrow \\ & & & & X/A & & SA & & SX \end{array}$$

In the first row, each space is obtained from its predecessor by attaching a cone on the subspace two steps back in the sequence. The vertical maps are the quotient maps obtained by collapsing the most recently attached cone to a point. In many cases the quotient map collapsing a contractible subspace to a point is a homotopy equivalence, hence induces an isomorphism on $\widetilde{K}$. This conclusion holds generally, in fact:

Lemma 2.10. *If A is contractible, the quotient map $q : X \to X/A$ induces a bijection $q^* : \mathrm{Vect}^n(X/A) \to \mathrm{Vect}^n(X)$ for all n.*

Proof: A vector bundle $E \to X$ must be trivial over A since A is contractible. A trivialization h gives a vector bundle $E/h \to X/A$ as in the proof of the previous proposition. We assert that the isomorphism class of E/h does not depend on h. This can be seen as follows. Given two trivializations h_0 and h_1, by writing $h_1 = (h_1 h_0^{-1}) h_0$ we see that h_0 and h_1 differ by an element of $g_x \in GL_n(\mathbb{C})$ over each point $x \in A$. The resulting map $g: A \to GL_n(\mathbb{C})$ is homotopic to a constant map $x \mapsto \alpha \in GL_n(\mathbb{C})$ since A is contractible. Writing now $h_1 = (h_1 h_0^{-1} \alpha^{-1})(\alpha h_0)$, we see that by composing h_0 with α in each fiber, which does not change E/h_0, we may assume that α is the identity. Then the homotopy from g to the identity gives a homotopy H from h_0 to h_1. In the same way that we constructed E/h we construct a vector bundle $(E \times I)/H \to (X/A) \times I$ restricting to E/h_0 over one end and to E/h_1 over the other end, hence $E/h_0 \approx E/h_1$.

Thus we have a well-defined map $\mathrm{Vect}^n(X) \to \mathrm{Vect}^n(X/A)$, $E \mapsto E/h$. This is an inverse to q^* since $q^*(E/h) \approx E$ as we noted in the preceding proposition, and for a bundle $E \to X/A$ we have $q^*(E)/h \approx E$ for the evident trivialization h of $q^*(E)$ over A $\square$

From this lemma and the preceding proposition it follows that we have a long exact sequence of $\widetilde{K}$ groups

$$\cdots \to \widetilde{K}(SX) \to \widetilde{K}(SA) \to \widetilde{K}(X/A) \to \widetilde{K}(X) \to \widetilde{K}(A)$$

For example, if X is the wedge sum $A \vee B$ then $X/A = B$ and the sequence breaks up into split short exact sequences, which implies that the map $\widetilde{K}(X) \to \widetilde{K}(A) \oplus \widetilde{K}(B)$ obtained by restriction to A and B is an isomorphism.

Deducing Periodicity from the Product Theorem

We can use the exact sequence displayed above to obtain a reduced version of the external product, a ring homomorphism $\widetilde{K}(X) \otimes \widetilde{K}(Y) \to \widetilde{K}(X \wedge Y)$ where $X \wedge Y = X \times Y / X \vee Y$ and $X \vee Y = X \times \{y_0\} \cup \{x_0\} \times Y \subset X \times Y$ for chosen basepoints $x_0 \in X$ and $y_0 \in Y$. The space $X \wedge Y$ is called the *smash product* of X and Y. To define the reduced product, consider the long exact sequence for the pair $(X \times Y, X \vee Y)$:

$$\begin{array}{ccccccccc} \widetilde{K}(S(X \times Y)) & \longrightarrow & \widetilde{K}(S(X \vee Y)) & \longrightarrow & \widetilde{K}(X \wedge Y) & \longrightarrow & \widetilde{K}(X \times Y) & \longrightarrow & \widetilde{K}(X \vee Y) \\ & & \wr\wr & & & & & & \wr\wr \\ & & \widetilde{K}(SX) \oplus \widetilde{K}(SY) & & & & & & \widetilde{K}(X) \oplus \widetilde{K}(Y) \end{array}$$

The second of the two vertical isomorphisms here was noted earlier. The first vertical isomorphism arises in similar fashion by using a reduced version of the suspension operator which associates to a space Z with basepoint z_0 the quotient space ΣZ of SZ obtained by collapsing the segment $\{z_0\} \times I$ to a point. The quotient map $SZ \to \Sigma Z$ induces an isomorphism $\widetilde{K}(SZ) \approx \widetilde{K}(\Sigma Z)$ by the preceding lemma. For reduced suspension we have $\Sigma(X \vee Y) = \Sigma X \vee \Sigma Y$, which gives the first isomorphism in the diagram. The last horizontal map in the sequence is a split surjection, with splitting $\widetilde{K}(X) \oplus \widetilde{K}(Y) \to \widetilde{K}(X \times Y)$, $(a, b) \mapsto p_1^*(a) + p_2^*(b)$ where p_1 and p_2 are the projections of $X \times Y$ onto X and Y. Similarly, the first map splits via $(Sp_1)^* + (Sp_2)^*$. So we get a splitting $\widetilde{K}(X \times Y) \approx \widetilde{K}(X \wedge Y) \oplus \widetilde{K}(X) \oplus \widetilde{K}(Y)$.

For $a \in \widetilde{K}(X) = \mathrm{Ker}(K(X) \to K(x_0))$ and $b \in \widetilde{K}(Y) = \mathrm{Ker}(K(Y) \to K(y_0))$ the external product $a * b = p_1^*(a)p_2^*(b) \in K(X \times Y)$ has $p_1^*(a)$ restricting to zero in $K(Y)$ and $p_2^*(b)$ restricting to zero in $K(X)$, so $p_1^*(a)p_2^*(b)$ restricts to zero in both $K(X)$ and $K(Y)$, hence in $K(X \vee Y)$. In particular, $a * b$ lies in $\widetilde{K}(X \times Y)$, and from the short exact sequence above, $a * b$ pulls back to a unique element of $\widetilde{K}(X \wedge Y)$. This defines the reduced external product $\widetilde{K}(X) \otimes \widetilde{K}(Y) \to \widetilde{K}(X \wedge Y)$. It is essentially a restriction of the unreduced external product, as shown in the following diagram,

$$\begin{array}{ccccccc} K(X) \otimes K(Y) & \approx & (\widetilde{K}(X) \otimes \widetilde{K}(Y)) & \oplus & \widetilde{K}(X) \oplus & \widetilde{K}(Y) \oplus & \mathbb{Z} \\ \downarrow & & \downarrow & & \| & \| & \| \\ K(X \times Y) & \approx & \widetilde{K}(X \wedge Y) & \oplus & \widetilde{K}(X) \oplus & \widetilde{K}(Y) \oplus & \mathbb{Z} \end{array}$$

so the reduced external product is also a ring homomorphism, and we shall use the same notation $a * b$ for both reduced and unreduced external product, leaving the reader to determine from context which is meant.

Since $S^n \wedge X$ is the n-fold iterated reduced suspension $\Sigma^n X$, which is a quotient of the ordinary n-fold suspension $S^n X$ obtained by collapsing an n-disk in $S^n X$ to a point, the quotient map $S^n X \to S^n \wedge X$ induces an isomorphism on $\widetilde{K}$ by Lemma 2.10. Then the reduced external product gives rise to a homomorphism

$$\beta : \widetilde{K}(X) \to \widetilde{K}(S^2 X), \qquad \beta(a) = (H - 1) * a$$

where H is the canonical line bundle over $S^2 = \mathbb{C}\mathrm{P}^1$. Here at last is the Bott Periodicity Theorem:

Theorem 2.11. *The homomorphism $\beta : \widetilde{K}(X) \to \widetilde{K}(S^2 X)$, $\beta(a) = (H - 1) * a$, is an isomorphism for all compact Hausdorff spaces X.*

By the diagram displayed above this is equivalent to the Product Theorem proved in the previous section.

Corollary 2.12. *$\widetilde{K}(S^{2n+1}) = 0$ and $\widetilde{K}(S^{2n}) \approx \mathbb{Z}$, generated by the n-fold reduced external product $(H - 1) * \cdots * (H - 1)$.*

Extending to a Cohomology Theory

As we saw earlier, a pair (X, A) of compact Hausdorff spaces gives rise to an exact sequence of $\widetilde{K}$ groups, the first row in the following diagram:

$$\begin{array}{ccccccccccccccc} \widetilde{K}(S^2X) & \to & \widetilde{K}(S^2A) & \to & \widetilde{K}(S(X/A)) & \to & \widetilde{K}(SX) & \to & \widetilde{K}(SA) & \to & \widetilde{K}(X/A) & \to & \widetilde{K}(X) & \to & \widetilde{K}(A) \\ \| & & \| & & \| & & \| & & \| & & \| & & \| & & \| \\ \widetilde{K}^{-2}(X) & \to & \widetilde{K}^{-2}(A) & \to & \widetilde{K}^{-1}(X,A) & \to & \widetilde{K}^{-1}(X) & \to & \widetilde{K}^{-1}(A) & \to & \widetilde{K}^{0}(X,A) & \to & \widetilde{K}^{0}(X) & \to & \widetilde{K}^{0}(A) \\ \beta \uparrow \approx & & \beta \uparrow \approx & & & & & & & & & & & & \\ \widetilde{K}^{0}(X) & \to & \widetilde{K}^{0}(A) & & & & & & & & & & & & \end{array}$$

If we set $\widetilde{K}^{-n}(X) = \widetilde{K}(S^n X)$ and $\widetilde{K}^{-n}(X,A) = \widetilde{K}(S^n(X/A))$, this sequence can be written as in the second row. Negative indices are chosen here so that the 'coboundary' maps in this sequence increase dimension, as in ordinary cohomology. The lower left corner of the diagram containing the Bott periodicity isomorphisms β commutes since external tensor product with $H-1$ commutes with maps between spaces. So the long exact sequence in the second row can be rolled up into a six-term periodic exact sequence. It is reasonable to extend the definition of $\widetilde{K}^n$ to positive n by setting $\widetilde{K}^{2i}(X) = \widetilde{K}(X)$ and $\widetilde{K}^{2i+1}(X) = \widetilde{K}(SX)$. Then the six-term exact sequence can be written

$$\begin{array}{ccccc} \widetilde{K}^0(X,A) & \longrightarrow & \widetilde{K}^0(X) & \longrightarrow & \widetilde{K}^0(A) \\ \uparrow & & & & \downarrow \\ \widetilde{K}^1(A) & \longleftarrow & \widetilde{K}^1(X) & \longleftarrow & \widetilde{K}^1(X,A) \end{array}$$

A product $\widetilde{K}^i(X) \otimes \widetilde{K}^j(Y) \to \widetilde{K}^{i+j}(X \wedge Y)$ is obtained from the external product $\widetilde{K}(X) \otimes \widetilde{K}(Y) \to \widetilde{K}(X \wedge Y)$ by replacing X and Y by $S^i X$ and $S^j Y$. If we define $\widetilde{K}^*(X) = \widetilde{K}^0(X) \oplus \widetilde{K}^1(X)$, then this gives a product $\widetilde{K}^*(X) \otimes \widetilde{K}^*(Y) \to \widetilde{K}^*(X \wedge Y)$. The relative form of this is a product $\widetilde{K}^*(X,A) \otimes \widetilde{K}^*(Y,B) \to \widetilde{K}^*(X \times Y, X \times B \cup A \times Y)$, coming from the products $\widetilde{K}(\Sigma^i(X/A)) \otimes \widetilde{K}(\Sigma^j(Y/B)) \longrightarrow \widetilde{K}(\Sigma^{i+j}(X/A \wedge Y/B))$ using the natural identification $(X \times Y)/(X \times B \cup A \times Y) = X/A \wedge Y/B$.

If we compose the external product $\widetilde{K}^*(X) \otimes \widetilde{K}^*(X) \to \widetilde{K}^*(X \wedge X)$ with the map $\widetilde{K}^*(X \wedge X) \to \widetilde{K}^*(X)$ induced by the diagonal map $X \to X \wedge X$, $x \mapsto (x,x)$, then we obtain a multiplication on $\widetilde{K}^*(X)$ making it into a ring, and it is not hard to check that this extends the previously defined ring structure on $\widetilde{K}^0(X)$. The general relative form of this product on $\widetilde{K}^*(X)$ is a product $\widetilde{K}^*(X,A) \otimes \widetilde{K}^*(X,B) \to \widetilde{K}^*(X, A \cup B)$ which is induced by the relativized diagonal map $X/(A \cup B) \to X/A \wedge X/B$.

Example 2.13. Suppose that $X = A \cup B$ where A and B are compact contractible subspaces of X containing the basepoint. Then the product $\widetilde{K}^*(X) \otimes \widetilde{K}^*(X) \to \widetilde{K}^*(X)$ is identically zero since it is equivalent to the composition

$$\widetilde{K}^*(X,A) \otimes \widetilde{K}^*(X,B) \to \widetilde{K}^*(X, A \cup B) \to \widetilde{K}^*(X)$$

and $\widetilde{K}^*(X, A \cup B) = 0$ since $X = A \cup B$. For example if X is a suspension we can take A and B to be its two cones, with a basepoint in their intersection. As a particular case we see that the product in $\widetilde{K}^*(S^n) \approx \mathbb{Z}$ is trivial for $n > 0$. For $n = 0$ the multiplication in $\widetilde{K}^*(S^0) \approx \mathbb{Z}$ is just the usual multiplication of integers since $\mathbb{R}^m \otimes \mathbb{R}^n \approx \mathbb{R}^{mn}$. This illustrates the necessity of the condition that A and B both contain the basepoint of X, since without this condition we could take A and B to be the two points of S^0.

More generally, if X is the union of compact contractible subspaces $A_1, \cdots, A_n$ containing the basepoint then the n-fold product

$$\widetilde{K}^*(X, A_1) \otimes \cdots \otimes \widetilde{K}^*(X, A_n) \to \widetilde{K}^*(X, A_1 \cup \cdots \cup A_n)$$

is trivial, so all n-fold products in $\widetilde{K}^*(X)$ are trivial. In particular all elements of $\widetilde{K}^*(X)$ are nilpotent since their n^{th} power is zero. This applies to all compact manifolds for example since they are covered by finitely many closed balls, and the condition that each A_i contain the basepoint can be achieved by adjoining to each ball an arc to a fixed basepoint. In a similar fashion one can see that this observation applies to all finite cell complexes, by induction on the number of cells.

Whereas multiplication in $\widetilde{K}(X)$ is commutative, in $\widetilde{K}^*(X)$ this is only true up to sign:

Proposition 2.14. *$\alpha\beta = (-1)^{ij}\beta\alpha$ for $\alpha \in \widetilde{K}^i(X)$ and $\beta \in \widetilde{K}^j(X)$.*

Proof: The product is the composition

$$\widetilde{K}(S^i \wedge X) \otimes \widetilde{K}(S^j \wedge X) \longrightarrow \widetilde{K}(S^i \wedge S^j \wedge X \wedge X) \longrightarrow \widetilde{K}(S^i \wedge S^j \wedge X)$$

where the first map is external product and the second is induced by the diagonal map on the X factors. Replacing the product $\alpha\beta$ by the product $\beta\alpha$ amounts to switching the two factors in the first term $\widetilde{K}(S^i \wedge X) \otimes \widetilde{K}(S^j \wedge X)$, and this corresponds to switching the S^i and S^j factors in the third term $\widetilde{K}(S^i \wedge S^j \wedge X)$. Viewing $S^i \wedge S^j$ as the smash product of $i+j$ copies of S^1, then switching S^i and S^j in $S^i \wedge S^j$ is a product of ij transpositions of adjacent factors. Transposing the two factors of $S^1 \wedge S^1$ is equivalent to reflection of S^2 across an equator. Thus it suffices to see that switching the two ends of a suspension SY induces multiplication by -1 in $\widetilde{K}(SY)$. If we view $\widetilde{K}(SY)$ as $\langle Y, U \rangle$, then switching ends of SY corresponds to the map $U \to U$ sending a matrix to its inverse. We noted in the proof of Proposition 2.2 that the group operation in $K(SY)$ is the same as the operation induced by the product in U, so the result follows. ☐

Proposition 2.15. *The exact sequence at the right is an exact sequence of $\widetilde{K}^*(X)$-modules, with the maps homomorphisms of $\widetilde{K}^*(X)$-modules.*

$$\begin{array}{ccc} \widetilde{K}^*(X,A) & \longrightarrow & \widetilde{K}^*(X) \\ \nwarrow & & \swarrow \\ & \widetilde{K}^*(A) & \end{array}$$

The $\widetilde{K}^*(X)$-module structure on $\widetilde{K}^*(A)$ is defined by $\xi \cdot \alpha = i^*(\xi)\alpha$ where i is the inclusion $A \hookrightarrow X$ and the product on the right side of the equation is multiplication in the ring $\widetilde{K}^*(A)$. To define the module structure on $\widetilde{K}^*(X,A)$, observe that the diagonal map $X \to X \wedge X$ induces a well-defined quotient map $X/A \to X \wedge X/A$, and this leads to a product $\widetilde{K}^*(X) \otimes \widetilde{K}^*(X,A) \to \widetilde{K}^*(X,A)$.

Proof: To see that the maps in the exact sequence are module homomorphisms we look at the diagram

$$\begin{array}{ccccccc}
\widetilde{K}(S^jSA) & \longrightarrow & \widetilde{K}(S^j(X/A)) & \longrightarrow & \widetilde{K}(S^jX) & \longrightarrow & \widetilde{K}(S^jA) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\widetilde{K}(S^iX\wedge S^jSA) & \longrightarrow & \widetilde{K}(S^iX\wedge S^j(X/A)) & \longrightarrow & \widetilde{K}(S^iX\wedge S^jX) & \longrightarrow & \widetilde{K}(S^iX\wedge S^jA) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\widetilde{K}(S^{i+j}SA) & \longrightarrow & \widetilde{K}(S^{i+j}(X/A)) & \longrightarrow & \widetilde{K}(S^{i+j}X) & \longrightarrow & \widetilde{K}(S^{i+j}A)
\end{array}$$

where the vertical maps between the first two rows are external product with a fixed element of $\widetilde{K}(S^iX)$ and the vertical maps between the second and third rows are induced by diagonal maps. What we must show is that the diagram commutes. For the upper two rows this follows from naturality of external product since the horizontal maps are induced by maps between spaces. The lower two rows are induced from suspensions of maps between spaces,

$$\begin{array}{ccccccc}
X\wedge SA & \longleftarrow & X\wedge X/A & \longleftarrow & X\wedge X & \longleftarrow & X\wedge A \\
\uparrow & & \uparrow & & \uparrow & & \uparrow \\
SA & \longleftarrow & X/A & \longleftarrow & X & \longleftarrow & A
\end{array}$$

so it suffices to show this diagram commutes up to homotopy. This is obvious for the middle and right squares. The left square can be rewritten

$$\begin{array}{ccc}
X\wedge SA & \longleftarrow & X\wedge(X\cup CA) \\
\uparrow & & \uparrow \\
SA & \longleftarrow & X\cup CA
\end{array}$$

where the horizontal maps collapse the copy of X in $X\cup CA$ to a point, the left vertical map sends $(a,s)\in SA$ to $(a,a,s)\in X\wedge SA$, and the right vertical map sends $x\in X$ to $(x,x)\in X\cup CA$ and $(a,s)\in CA$ to $(a,a,s)\in X\wedge CA$. Commutativity is then obvious. □

It is often convenient to have an unreduced version of the groups $\widetilde{K}^n(X)$, and this can easily be done by the simple device of defining $K^n(X)$ to be $\widetilde{K}^n(X_+)$ where X_+ is X with a disjoint basepoint labeled '+' adjoined. For $n=0$ this is consistent with the relation between K and $\widetilde{K}$ since $K^0(X)=\widetilde{K}^0(X_+)=\widetilde{K}(X_+)=\mathrm{Ker}(K(X_+)\to K(+))=K(X)$. For $n=1$ this definition yields $K^1(X)=\widetilde{K}^1(X)$ since $S(X_+)\simeq SX\vee S^1$ and $\widetilde{K}(SX\vee S^1)\approx\widetilde{K}(SX)\oplus\widetilde{K}(S^1)\approx\widetilde{K}(SX)$ since $\widetilde{K}(S^1)=0$. For a pair (X,A) with $A\neq\varnothing$ one defines $K^n(X,A)=\widetilde{K}^n(X,A)$, and then the six-term long exact sequence is valid also for unreduced groups. When $A=\varnothing$ this remains valid if we interpret $X/\varnothing$ as X_+.

Since $X_+\wedge Y_+=(X\times Y)_+$, the external product $\widetilde{K}^*(X)\otimes\widetilde{K}^*(Y)\to\widetilde{K}^*(X\wedge Y)$ gives a product $K^*(X)\otimes K^*(Y)\to K^*(X\times Y)$. Taking $X=Y$ and composing with the map $K^*(X\times X)\to K^*(X)$ induced by the diagonal map $X\to X\times X$, $x\mapsto(x,x)$, we get a product $K^*(X)\otimes K^*(X)\to K^*(X)$ which makes $K^*(X)$ into a ring.

There is a relative product $K^i(X,A)\otimes K^j(Y,B)\to K^{i+j}(X\times Y, X\times B\cup A\times Y)$ defined as the external product $\widetilde{K}(\Sigma^i(X/A))\otimes\widetilde{K}(\Sigma^j(Y/B))\longrightarrow\widetilde{K}(\Sigma^{i+j}(X/A\wedge Y/B))$, using the natural identification $(X\times Y)/(X\times B\cup A\times Y) = X/A\wedge Y/B$. This works when $A=\varnothing$ since we interpret $X/\varnothing$ as X_+, and similarly if $Y=\varnothing$. Via the diagonal map we obtain also a product $K^i(X,A)\otimes K^j(X,B)\to K^{i+j}(X,A\cup B)$.

With these definitions the preceding two propositions are valid also for unreduced K–groups.

Elementary Applications

With the calculation $\widetilde{K}^*(S^n)\approx\mathbb{Z}$ completed, it would be possible to derive many of the same applications that follow from the corresponding calculation for ordinary homology or cohomology, as in [AT]. For example:

— There is no retraction of D^n onto its boundary S^{n-1}, since this would mean that the identity map of $\widetilde{K}^*(S^{n-1})$ factored as $\widetilde{K}^*(S^{n-1})\to\widetilde{K}^*(D^n)\to\widetilde{K}^*(S^{n-1})$, but the middle group is trivial.

— The Brouwer fixed point theorem, that for every map $f:D^n\to D^n$ there is a point $x\in D^n$ with $f(x)=x$. For if not then it is easy to construct a retraction of D^n onto S^{n-1}.

— The notion of degree for maps $f:S^n\to S^n$, namely the integer $d(f)$ such that the induced homomorphism $f^*:\widetilde{K}^*(S^n)\to\widetilde{K}^*(S^n)$ is multiplication by $d(f)$. Reasoning as in Proposition 2.2, one sees that d is a homomorphism $\pi_n(S^n)\to\mathbb{Z}$. In particular a reflection has degree -1 and hence the antipodal map of S^n, which is the composition of $n+1$ reflections, has degree $(-1)^{n+1}$ since $d(fg)=d(f)d(g)$. Consequences of this include the fact that an even-dimensional sphere has no nonvanishing vector fields.

However there are some things homology can do that would be harder using K–theory since $\widetilde{K}^*(S^n)$ only distinguishes even-dimensional spheres from odd-dimensional spheres. Also, since we have so far only defined K–theory for compact spaces, it would take more work to derive some of the other classical applications of homology such as Brouwer's theorems on invariance of dimension and invariance of domain, or the Jordan curve theorem and its higher-dimensional analogs.

2.3. Division Algebras and Parallelizable Spheres

With the hard work of proving Bott Periodicity now behind us, the goal of this section is to prove Adams' theorem on the Hopf invariant, with its famous applications including the nonexistence of division algebras beyond the Cayley octonions:

Theorem 2.16. *The following statements are true only for $n = 1, 2, 4$, and 8:*
(a) $\mathbb{R}^n$ *is a division algebra.*
(b) S^{n-1} *is parallelizable, i.e., there exist $n-1$ tangent vector fields to S^{n-1} which are linearly independent at each point, or in other words, the tangent bundle to S^{n-1} is trivial.*

A division algebra structure on $\mathbb{R}^n$ is a multiplication map $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ such that the maps $x \mapsto ax$ and $x \mapsto xa$ are linear for each $a \in \mathbb{R}^n$ and invertible if $a \neq 0$. Since we are dealing with linear maps $\mathbb{R}^n \to \mathbb{R}^n$, invertibility is equivalent to having trivial kernel, which translates into the statement that the multiplication has no zero divisors. An identity element is not assumed, but the multiplication can be modified to produce an identity in the following way. Choose a unit vector $e \in \mathbb{R}^n$. After composing the multiplication with an invertible linear map $\mathbb{R}^n \to \mathbb{R}^n$ taking e^2 to e we may assume that $e^2 = e$. Let α be the map $x \mapsto xe$ and β the map $x \mapsto ex$. The new product $(x, y) \mapsto \alpha^{-1}(x)\beta^{-1}(y)$ then sends (x, e) to $\alpha^{-1}(x)\beta^{-1}(e) = \alpha^{-1}(x)e = x$, and similarly it sends (e, y) to y. Since the maps $x \mapsto ax$ and $x \mapsto xa$ are surjective for each $a \neq 0$, the equations $ax = e$ and $xa = e$ are solvable, so nonzero elements of the division algebra have multiplicative inverses on the left and right.

H-Spaces

The first step in the proof of the theorem is to reduce it to showing when the sphere S^{n-1} is an H-space.

To say that S^{n-1} is an H-space means there is a continuous multiplication map $S^{n-1} \times S^{n-1} \to S^{n-1}$ having a two-sided identity element $e \in S^{n-1}$. This is weaker than being a topological group since associativity and inverses are not assumed. For example, S^1, S^3, and S^7 are H-spaces by restricting the multiplication of complex numbers, quaternions, and Cayley octonions to the respective unit spheres, but only S^1 and S^3 are topological groups since the multiplication of octonions is nonassociative.

Lemma 2.17. *If $\mathbb{R}^n$ is a division algebra, or if S^{n-1} is parallelizable, then S^{n-1} is an H-space.*

Proof: Having a division algebra structure on $\mathbb{R}^n$ with two-sided identity, an H-space structure on S^{n-1} is given by $(x, y) \mapsto xy/|xy|$, which is well-defined since a division algebra has no zero divisors.

Now suppose that S^{n-1} is parallelizable, with tangent vector fields $v_1, \cdots, v_{n-1}$ which are linearly independent at each point of S^{n-1}. By the Gram-Schmidt process we may make the vectors $x, v_1(x), \cdots, v_{n-1}(x)$ orthonormal for all $x \in S^{n-1}$. We may assume also that at the first standard basis vector e_1, the vectors $v_1(e_1), \cdots, v_{n-1}(e_1)$ are the standard basis vectors $e_2, \cdots, e_n$, by changing the sign of v_{n-1} if necessary to get orientations right, then deforming the vector fields near e_1. Let $\alpha_x \in SO(n)$ send the standard basis to $x, v_1(x), \cdots, v_{n-1}(x)$. Then the map $(x,y) \mapsto \alpha_x(y)$ defines an H-space structure on S^{n-1} with identity element the vector e_1 since α_{e_1} is the identity map and $\alpha_x(e_1) = x$ for all x. □

Before proceeding further let us list a few easy consequences of Bott periodicity which will be needed.

(1) We have already seen that $\widetilde{K}(S^n)$ is $\mathbb{Z}$ for n even and 0 for n odd. This comes from repeated application of the periodicity isomorphism $\widetilde{K}(X) \approx \widetilde{K}(S^2X)$, $\alpha \mapsto \alpha * (H-1)$, the external product with the generator $H-1$ of $\widetilde{K}(S^2)$, where H is the canonical line bundle over $S^2 = \mathbb{C}\mathrm{P}^1$. In particular we see that a generator of $\widetilde{K}(S^{2k})$ is the k-fold external product $(H-1) * \cdots * (H-1)$. We note also that the multiplication in $\widetilde{K}(S^{2k})$ is trivial since this ring is the k-fold tensor product of the ring $\widetilde{K}(S^2)$, which has trivial multiplication by Corollary 2.3.

(2) The external product $\widetilde{K}(S^{2k}) \otimes \widetilde{K}(X) \to \widetilde{K}(S^{2k} \wedge X)$ is an isomorphism since it is an iterate of the periodicity isomorphism.

(3) The external product $K(S^{2k}) \otimes K(X) \to K(S^{2k} \times X)$ is an isomorphism. This follows from (2) by the same reasoning which showed the equivalence of the reduced and unreduced forms of Bott periodicity. Since external product is a ring homomorphism, the isomorphism $\widetilde{K}(S^{2k} \wedge X) \approx \widetilde{K}(S^{2k}) \otimes \widetilde{K}(X)$ is a ring isomorphism. For example, since $K(S^{2k})$ can be described as the quotient ring $\mathbb{Z}[\alpha]/(\alpha^2)$, we can deduce that $K(S^{2k} \times S^{2\ell})$ is $\mathbb{Z}[\alpha,\beta]/(\alpha^2,\beta^2)$ where α and β are the pullbacks of generators of $\widetilde{K}(S^{2k})$ and $\widetilde{K}(S^{2\ell})$ under the projections of $S^{2k} \times S^{2\ell}$ onto its two factors. An additive basis for $K(S^{2k} \times S^{2\ell})$ is thus $\{1, \alpha, \beta, \alpha\beta\}$.

We can apply the last calculation to show that S^{2k} is not an H-space if $k > 0$. Suppose $\mu : S^{2k} \times S^{2k} \to S^{2k}$ is an H-space multiplication. The induced homomorphism of K-rings then has the form $\mu^* : \mathbb{Z}[\gamma]/(\gamma^2) \to \mathbb{Z}[\alpha,\beta]/(\alpha^2,\beta^2)$. We claim that $\mu^*(\gamma) = \alpha + \beta + m\alpha\beta$ for some integer m. The composition $S^{2k} \xrightarrow{i} S^{2k} \times S^{2k} \xrightarrow{\mu} S^{2k}$ is the identity, where i is the inclusion onto either of the subspaces $S^{2k} \times \{e\}$ or $\{e\} \times S^{2k}$, with e the identity element of the H-space structure. The map i^* for i the inclusion onto the first factor sends α to γ and β to 0, so the coefficient of α in $\mu^*(\gamma)$ must be 1. Similarly the coefficient of β must be 1, proving the claim. However, this leads to a contradiction since it implies that $\mu^*(\gamma^2) = (\alpha + \beta + m\alpha\beta)^2 = 2\alpha\beta \neq 0$, which is impossible since $\gamma^2 = 0$.

There remains the much more difficult problem of showing that S^{n-1} is not an H-space when n is even and different from 2, 4, and 8. The first step is a simple construction which associates to a map $g : S^{n-1} \times S^{n-1} \to S^{n-1}$ a map $\widehat{g} : S^{2n-1} \to S^n$.

To define this, we regard S^{2n-1} as $\partial(D^n \times D^n) = \partial D^n \times D^n \cup D^n \times \partial D^n$, and S^n we take as the union of two disks D^n_+ and D^n_- with their boundaries identified. Then $\widehat{g}$ is defined on $\partial D^n \times D^n$ by $\widehat{g}(x,y) = |y|g(x,y/|y|) \in D^n_+$ and on $D^n \times \partial D^n$ by $\widehat{g}(x,y) = |x|g(x/|x|,y) \in D^n_-$. Note that $\widehat{g}$ is well-defined and continuous, even when $|x|$ or $|y|$ is zero, and $\widehat{g}$ agrees with g on $S^{n-1} \times S^{n-1}$.

Now we specialize to the case that n is even, or in other words, we replace n by $2n$. For a map $f : S^{4n-1} \to S^{2n}$, let C_f be S^{2n} with a cell e^{4n} attached by f. The quotient C_f/S^{2n} is then S^{4n}, and since $\widetilde{K}^1(S^{4n}) = \widetilde{K}^1(S^{2n}) = 0$, the exact sequence of the pair (C_f, S^{2n}) becomes a short exact sequence

$$0 \longrightarrow \widetilde{K}(S^{4n}) \longrightarrow \widetilde{K}(C_f) \longrightarrow \widetilde{K}(S^{2n}) \longrightarrow 0$$

Let $\alpha \in \widetilde{K}(C_f)$ be the image of the generator $(H-1) * \cdots * (H-1)$ of $\widetilde{K}(S^{4n})$ and let $\beta \in \widetilde{K}(C_f)$ map to the generator $(H-1) * \cdots * (H-1)$ of $\widetilde{K}(S^{2n})$. The element β^2 maps to 0 in $\widetilde{K}(S^{2n})$ since the square of any element of $\widetilde{K}(S^{2n})$ is zero. So by exactness we have $\beta^2 = h\alpha$ for some integer h. The mod 2 value of h depends only on f, not on the choice of β, since β is unique up to adding an integer multiple of α, and $(\beta + m\alpha)^2 = \beta^2 + 2m\alpha\beta$ since $\alpha^2 = 0$. The value of h mod 2 is called the *mod* 2 *Hopf invariant* of f. In fact $\alpha\beta = 0$ so h is well-defined in $\mathbb{Z}$ not just $\mathbb{Z}_2$, as we will see in §4.1, but for our present purposes the mod 2 value of h suffices.

Lemma 2.18. *If $g : S^{2n-1} \times S^{2n-1} \to S^{2n-1}$ is an H-space multiplication, then the associated map $\widehat{g} : S^{4n-1} \to S^{2n}$ has Hopf invariant ± 1.*

Proof: Let $e \in S^{2n-1}$ be the identity element for the H-space multiplication, and let $f = \widehat{g}$. In view of the definition of f it is natural to view the characteristic map Φ of the $4n$-cell of C_f as a map $(D^{2n} \times D^{2n}, \partial(D^{2n} \times D^{2n})) \to (C_f, S^{2n})$. In the following commutative diagram the horizontal maps are the product maps. The diagonal map is external product, equivalent to the external product $\widetilde{K}(S^{2n}) \otimes \widetilde{K}(S^{2n}) \to \widetilde{K}(S^{4n})$, which is an isomorphism since it is an iterate of the Bott periodicity isomorphism.

$$\begin{array}{ccc}
\widetilde{K}(C_f) \otimes \widetilde{K}(C_f) & \longrightarrow & \widetilde{K}(C_f) \\
\uparrow \approx & & \uparrow \\
\widetilde{K}(C_f, D^{2n}_-) \otimes \widetilde{K}(C_f, D^{2n}_+) & \longrightarrow & \widetilde{K}(C_f, S^{2n}) \\
\Phi^* \otimes \Phi^* \downarrow & & \Phi^* \downarrow \approx \\
\widetilde{K}(D^{2n} \times D^{2n}, \partial D^{2n} \times D^{2n}) \otimes \widetilde{K}(D^{2n} \times D^{2n}, D^{2n} \times \partial D^{2n}) & \longrightarrow & \widetilde{K}(D^{2n} \times D^{2n}, \partial(D^{2n} \times D^{2n})) \\
\downarrow \approx & \nearrow \approx & \\
\widetilde{K}(D^{2n} \times \{e\}, \partial D^{2n} \times \{e\}) \otimes \widetilde{K}(\{e\} \times D^{2n}, \{e\} \times \partial D^{2n}) & &
\end{array}$$

By the definition of an H-space and the definition of f, the map Φ restricts to a homeomorphism from $D^{2n} \times \{e\}$ onto D^{2n}_+ and from $\{e\} \times D^{2n}$ onto D^{2n}_-. It follows that the element $\beta \otimes \beta$ in the upper left group maps to a generator of the group in the bottom row of the diagram, since β restricts to a generator of $\widetilde{K}(S^{2n})$ by definition.

Therefore by commutativity of the diagram, the product map in the top row sends $\beta \otimes \beta$ to $\pm\alpha$ since α was defined to be the image of a generator of $\widetilde{K}(C_f, S^{2n})$. Thus we have $\beta^2 = \pm\alpha$, which says that the Hopf invariant of f is ± 1. □

In view of this lemma, Theorem 2.16 becomes a consequence of the following theorem of Adams:

Theorem 2.19. *If $f : S^{4n-1} \to S^{2n}$ is a map whose mod 2 Hopf invariant is 1, then $n = 1, 2,$ or 4.*

The proof of this will occupy the rest of this section.

Adams Operations

The Hopf invariant is defined in terms of the ring structure in K–theory, but in order to prove Adams' theorem, more structure is needed, namely certain ring homomorphisms $\psi^k : K(X) \to K(X)$. Here are their basic properties:

Theorem 2.20. *There exist ring homomorphisms $\psi^k : K(X) \to K(X)$, defined for all compact Hausdorff spaces X and all integers $k \geq 0$, and satisfying:*

(1) $\psi^k f^* = f^* \psi^k$ *for all maps $f : X \to Y$. (Naturality)*
(2) $\psi^k(L) = L^k$ *if L is a line bundle.*
(3) $\psi^k \circ \psi^\ell = \psi^{k\ell}$.
(4) $\psi^p(\alpha) \equiv \alpha^p \bmod p$ *for p prime.*

This last statement means that $\psi^p(\alpha) - \alpha^p = p\beta$ for some $\beta \in K(X)$.

In the special case of a vector bundle E which is a sum of line bundles L_i, properties (2) and (3) give the formula $\psi^k(L_1 \oplus \cdots \oplus L_n) = L_1^k + \cdots + L_n^k$. We would like a general definition of $\psi^k(E)$ which specializes to this formula when E is a sum of line bundles. The idea is to use the exterior powers $\lambda^k(E)$. From the corresponding properties for vector spaces we have:

(i) $\lambda^k(E_1 \oplus E_2) \approx \bigoplus_i \big(\lambda^i(E_1) \otimes \lambda^{k-i}(E_2)\big)$.
(ii) $\lambda^0(E) = 1$, the trivial line bundle.
(iii) $\lambda^1(E) = E$.
(iv) $\lambda^k(E) = 0$ for k greater than the maximum dimension of the fibers of E.

Recall that we want $\psi^k(E)$ to be $L_1^k + \cdots + L_n^k$ when $E = L_1 \oplus \cdots \oplus L_n$ for line bundles L_i. We will show in this case that there is a polynomial s_k with integer coefficients such that $L_1^k + \cdots + L_n^k = s_k(\lambda^1(E), \cdots, \lambda^k(E))$. This will lead us to define $\psi^k(E) = s_k(\lambda^1(E), \cdots, \lambda^k(E))$ for an arbitrary vector bundle E.

To see what the polynomial s_k should be, we first use the exterior powers $\lambda^i(E)$ to define a polynomial $\lambda_t(E) = \sum_i \lambda^i(E) t^i \in K(X)[t]$. This is a finite sum by property (iv), and property (i) says that $\lambda_t(E_1 \oplus E_2) = \lambda_t(E_1)\lambda_t(E_2)$. When $E = L_1 \oplus \cdots \oplus L_n$ this implies that $\lambda_t(E) = \prod_i \lambda_t(L_i)$, which equals $\prod_i (1 + L_i t)$ by (ii), (iii), and (iv).

The coefficient $\lambda^j(E)$ of t^j in $\lambda_t(E) = \prod_i(1 + L_i t)$ is the j^{th} elementary symmetric function σ_j of the L_i's, the sum of all products of j distinct L_i's. Thus we have

$$(*) \qquad \lambda^j(E) = \sigma_j(L_1, \cdots, L_n) \quad \text{if } E = L_1 \oplus \cdots \oplus L_n$$

To make the discussion completely algebraic, let us introduce the variable t_i for L_i. Thus $(1+t_1)\cdots(1+t_n) = 1+\sigma_1+\cdots+\sigma_n$, where σ_j is the j^{th} elementary symmetric polynomial in the t_i's. The fundamental theorem on symmetric polynomials, proved for example in [Lang, p. 134] or [van der Waerden, §26], asserts that every degree k symmetric polynomial in $t_1, \cdots, t_n$ can be expressed as a unique polynomial in $\sigma_1, \cdots, \sigma_k$. In particular, $t_1^k + \cdots + t_n^k$ is a polynomial $s_k(\sigma_1, \cdots, \sigma_k)$, called a *Newton polynomial*. This polynomial s_k is independent of n since we can pass from n to $n-1$ by setting $t_n = 0$. A recursive formula for s_k is

$$s_k = \sigma_1 s_{k-1} - \sigma_2 s_{k-2} + \cdots + (-1)^{k-2}\sigma_{k-1}s_1 + (-1)^{k-1}k\sigma_k$$

To derive this we may take $n = k$, and then if we substitute $x = -t_i$ in the identity $(x+t_1)\cdots(x+t_k) = x^k + \sigma_1 x^{k-1} + \cdots + \sigma_k$ we get $t_i^k = \sigma_1 t_i^{k-1} - \cdots + (-1)^{k-1}\sigma_k$. Summing over i then gives the recursion relation. The recursion relation easily yields for example

$$s_1 = \sigma_1 \qquad s_2 = \sigma_1^2 - 2\sigma_2 \qquad s_3 = \sigma_1^3 - 3\sigma_1\sigma_2 + 3\sigma_3$$
$$s_4 = \sigma_1^4 - 4\sigma_1^2\sigma_2 + 4\sigma_1\sigma_3 + 2\sigma_2^2 - 4\sigma_4$$

Summarizing, if we define $\psi^k(E) = s_k(\lambda^1(E), \cdots, \lambda^k(E))$, then in the case that E is a sum of line bundles $L_1 \oplus \cdots \oplus L_n$ we have

$$\begin{aligned}\psi^k(E) &= s_k(\lambda^1(E), \cdots, \lambda^k(E)) \\ &= s_k(\sigma_1(L_1, \cdots, L_n), \cdots, \sigma_k(L_1, \cdots, L_n)) \quad \text{by } (*) \\ &= L_1^k + \cdots + L_n^k\end{aligned}$$

Verifying that the definition $\psi^k(E) = s_k(\lambda^1(E), \cdots, \lambda^k(E))$ gives operations on $K(X)$ satisfying the properties listed in the theorem will be rather easy if we make use of the following general result:

The Splitting Principle. *Given a vector bundle $E \to X$ with X compact Hausdorff, there is a compact Hausdorff space $F(E)$ and a map $p : F(E) \to X$ such that the induced map $p^* : K^*(X) \to K^*(F(E))$ is injective and $p^*(E)$ splits as a sum of line bundles.*

This will be proved later in this section, but for the moment let us assume it and proceed with the proof of Theorem 2.20 and Adams' theorem.

Proof of Theorem 2.20: Property (1) for vector bundles, $f^*(\psi^k(E)) = \psi^k(f^*(E))$, follows immediately from the relation $f^*(\lambda^i(E)) = \lambda^i(f^*(E))$. Additivity of ψ^k for vector bundles, $\psi^k(E_1 \oplus E_2) = \psi^k(E_1) + \psi^k(E_2)$, follows from the splitting principle since we can first pull back to split E_1 then do a further pullback to split E_2, and the formula $\psi^k(L_1 \oplus \cdots \oplus L_n) = L_1^k + \cdots + L_n^k$ preserves sums. Since ψ^k is additive on vector bundles, it induces an additive operation on $K(X)$ defined by $\psi^k(E_1 - E_2) = \psi^k(E_1) - \psi^k(E_2)$.

For this extended ψ^k the properties (1) and (2) are clear. Multiplicativity is also easy from the splitting principle: If E is the sum of line bundles L_i and E' is the sum of line bundles L'_j, then $E \otimes E'$ is the sum of the line bundles $L_i \otimes L'_j$, so $\psi^k(E \otimes E') = \sum_{i,j} \psi^k(L_i \otimes L'_j) = \sum_{i,j}(L_i \otimes L'_j)^k = \sum_{i,j} L_i^k \otimes {L'_j}^k = \sum_i L_i^k \sum_j {L'_j}^k = \psi^k(E)\psi^k(E')$. Thus ψ^k is multiplicative for vector bundles, and it follows formally that it is multiplicative on elements of $K(X)$. For property (3) the splitting principle and additivity reduce us to the case of line bundles, where $\psi^k(\psi^\ell(L)) = L^{k\ell} = \psi^{k\ell}(L)$. Likewise for (4), if $E = L_1 + \cdots + L_n$, then $\psi^p(E) = L_1^p + \cdots + L_n^p \equiv (L_1 + \cdots + L_n)^p = E^p \bmod p$. $\square$

By the naturality property (1), ψ^k restricts to an operation $\psi^k : \widetilde{K}(X) \to \widetilde{K}(X)$ since $\widetilde{K}(X)$ is the kernel of the homomorphism $K(X) \to K(x_0)$ for $x_0 \in X$. For the external product $\widetilde{K}(X) \otimes \widetilde{K}(Y) \to \widetilde{K}(X \wedge Y)$, we have the formula $\psi^k(\alpha * \beta) = \psi^k(\alpha) * \psi^k(\beta)$ since if one looks back at the definition of $\alpha * \beta$, one sees this was defined as $p_1^*(\alpha)p_2^*(\beta)$, hence

$$\begin{aligned}\psi^k(\alpha * \beta) &= \psi^k(p_1^*(\alpha)p_2^*(\beta)) \\ &= \psi^k(p_1^*(\alpha))\psi^k(p_2^*(\beta)) \\ &= p_1^*(\psi^k(\alpha))p_2^*(\psi^k(\beta)) \\ &= \psi^k(\alpha) * \psi^k(\beta).\end{aligned}$$

This will allow us to compute ψ^k on $\widetilde{K}(S^{2n}) \approx \mathbb{Z}$. In this case ψ^k must be multiplication by some integer since it is an additive homomorphism of $\mathbb{Z}$.

Proposition 2.21. *$\psi^k : \widetilde{K}(S^{2n}) \to \widetilde{K}(S^{2n})$ is multiplication by k^n.*

Proof: Consider first the case $n = 1$. Since ψ^k is additive, it will suffice to show $\psi^k(\alpha) = k\alpha$ for α a generator of $\widetilde{K}(S^2)$. We can take $\alpha = H - 1$ for H the canonical line bundle over $S^2 = \mathbb{C}\mathrm{P}^1$. Then

$$\begin{aligned}\psi^k(\alpha) = \psi^k(H-1) &= H^k - 1 \quad \text{by property (2)} \\ &= (1+\alpha)^k - 1 \\ &= 1 + k\alpha - 1 \quad \text{since } \alpha^i = (H-1)^i = 0 \text{ for } i \ge 2 \\ &= k\alpha\end{aligned}$$

When $n > 1$ we use the external product $\widetilde{K}(S^2) \otimes \widetilde{K}(S^{2n-2}) \to \widetilde{K}(S^{2n})$, which is an isomorphism, and argue by induction. Assuming the desired formula holds in $\widetilde{K}(S^{2n-2})$, we have $\psi^k(\alpha * \beta) = \psi^k(\alpha) * \psi^k(\beta) = k\alpha * k^{n-1}\beta = k^n(\alpha * \beta)$, and we are done. $\square$

Now we can use the operations ψ^2 and ψ^3 and the relation $\psi^2\psi^3 = \psi^6 = \psi^3\psi^2$ to prove Adams' theorem.

Proof of Theorem 2.19: The definition of the Hopf invariant of a map $f:S^{4n-1}\to S^{2n}$ involved elements $\alpha,\beta \in \widetilde{K}(C_f)$. By Proposition 2.21, $\psi^k(\alpha) = k^{2n}\alpha$ since α is the image of a generator of $\widetilde{K}(S^{4n})$. Similarly, $\psi^k(\beta) = k^n\beta + \mu_k\alpha$ for some $\mu_k \in \mathbb{Z}$. Therefore

$$\psi^k\psi^\ell(\beta) = \psi^k(\ell^n\beta + \mu_\ell\alpha) = k^n\ell^n\beta + (k^{2n}\mu_\ell + \ell^n\mu_k)\alpha$$

Since $\psi^k\psi^\ell = \psi^{k\ell} = \psi^\ell\psi^k$, the coefficient $k^{2n}\mu_\ell + \ell^n\mu_k$ of α is unchanged when k and ℓ are switched. This gives the relation

$$k^{2n}\mu_\ell + \ell^n\mu_k = \ell^{2n}\mu_k + k^n\mu_\ell, \qquad \text{or} \qquad (k^{2n} - k^n)\mu_\ell = (\ell^{2n} - \ell^n)\mu_k$$

By property (6) of ψ^2, we have $\psi^2(\beta) \equiv \beta^2 \bmod 2$. Since $\beta^2 = h\alpha$ with h the Hopf invariant of f, the formula $\psi^2(\beta) = 2^n\beta + \mu_2\alpha$ implies that $\mu_2 \equiv h \bmod 2$, so μ_2 is odd if we assume $h = \pm 1$. By the preceding displayed formula we have $(2^{2n} - 2^n)\mu_3 = (3^{2n} - 3^n)\mu_2$, or $2^n(2^n - 1)\mu_3 = 3^n(3^n - 1)\mu_2$, so 2^n divides $3^n(3^n - 1)\mu_2$. Since 3^n and μ_2 are odd, 2^n must then divide $3^n - 1$. The proof is completed by the following elementary number theory fact. □

Lemma 2.22. *If 2^n divides $3^n - 1$ then $n = 1, 2$, or 4.*

Proof: Write $n = 2^\ell m$ with m odd. We will show that the highest power of 2 dividing $3^n - 1$ is 2 for $\ell = 0$ and $2^{\ell+2}$ for $\ell > 0$. This implies the lemma since if 2^n divides $3^n - 1$, then by this fact, $n \le \ell + 2$, hence $2^\ell \le 2^\ell m = n \le \ell + 2$, which implies $\ell \le 2$ and $n \le 4$. The cases $n = 1, 2, 3, 4$ can then be checked individually.

We find the highest power of 2 dividing $3^n - 1$ by induction on ℓ. For $\ell = 0$ we have $3^n - 1 = 3^m - 1 \equiv 2 \bmod 4$ since $3 \equiv -1 \bmod 4$ and m is odd. In the next case $\ell = 1$ we have $3^n - 1 = 3^{2m} - 1 = (3^m - 1)(3^m + 1)$. The highest power of 2 dividing the first factor is 2 as we just showed, and the highest power of 2 dividing the second factor is 4 since $3^m + 1 \equiv 4 \bmod 8$ because $3^2 \equiv 1 \bmod 8$ and m is odd. So the highest power of 2 dividing the product $(3^m - 1)(3^m + 1)$ is 8. For the inductive step of passing from ℓ to $\ell + 1$ with $\ell \ge 1$, or in other words from n to $2n$ with n even, write $3^{2n} - 1 = (3^n - 1)(3^n + 1)$. Then $3^n + 1 \equiv 2 \bmod 4$ since n is even, so the highest power of 2 dividing $3^n + 1$ is 2. Thus the highest power of 2 dividing $3^{2n} - 1$ is twice the highest power of 2 dividing $3^n - 1$. □

The Splitting Principle

The splitting principle will be a fairly easy consequence of a general result about the K–theory of fiber bundles called the Leray-Hirsch theorem, together with a calculation of the ring structure of $K^*(\mathbb{C}\mathrm{P}^n)$. The following proposition will allow us to compute at least the additive structure of $K^*(\mathbb{C}\mathrm{P}^n)$.

Proposition 2.23. *If X is a finite cell complex with n cells, then $K^*(X)$ is a finitely generated group with at most n generators. If all the cells of X have even dimension then $K^1(X) = 0$ and $K^0(X)$ is free abelian with one basis element for each cell.*

The phrase 'finite cell complex' would normally mean 'finite CW complex' but we can take it to be something slightly more general: a space built from a finite discrete set by attaching a finite number of cells in succession, with no conditions on the dimensions of these cells, so cells are not required to attach only to cells of lower dimension. Finite cell complexes are always homotopy equivalent to finite CW complexes (by deforming each successive attaching map to be cellular) so the only advantages of finite cell complexes are technical. In particular, it is easy to see that a space is a finite cell complex if it is a fiber bundle over a finite cell complex with fibers that are also finite cell complexes. This is shown in Proposition 2.28 in a brief appendix to this section. It implies that the splitting principle can be applied staying within the realm of finite cell complexes.

Proof: We show this by induction on the number of cells. The complex X is obtained from a subcomplex A by attaching a k-cell, for some k. For the pair (X,A) we have an exact sequence $\widetilde{K}^*(X/A) \longrightarrow \widetilde{K}^*(X) \longrightarrow \widetilde{K}^*(A)$. Since $X/A = S^k$, we have $\widetilde{K}^*(X/A) \approx \mathbb{Z}$, and exactness implies that $\widetilde{K}^*(X)$ requires at most one more generator than $\widetilde{K}^*(A)$.

The first term of the exact sequence $K^1(X/A) \to K^1(X) \to K^1(A)$ is zero if all cells of X are of even dimension, so induction on the number of cells implies that $K^1(X) = 0$. Then there is a short exact sequence $0 \to \widetilde{K}^0(X/A) \to \widetilde{K}^0(X) \to \widetilde{K}^0(A) \to 0$ with $\widetilde{K}^0(X/A) \approx \mathbb{Z}$. By induction $\widetilde{K}(A)$ is free, so this sequence splits, hence $K^0(X) \approx \mathbb{Z} \oplus K^0(A)$ and the final statement of the proposition follows. □

This proposition applies in particular to $\mathbb{C}\mathrm{P}^n$, which has a cell structure with one cell in each dimension $0, 2, 4, \cdots, 2n$, so $K^1(\mathbb{C}\mathrm{P}^n) = 0$ and $K^0(\mathbb{C}\mathrm{P}^n) \approx \mathbb{Z}^{n+1}$. The ring structure is as simple as one could hope for:

Proposition 2.24. *$K(\mathbb{C}\mathrm{P}^n)$ is the quotient ring $\mathbb{Z}[L]/(L-1)^{n+1}$ where L is the canonical line bundle over $\mathbb{C}\mathrm{P}^n$.*

Thus by the change of variable $x = L-1$ we see that $K(\mathbb{C}\mathrm{P}^n)$ is the truncated polynomial ring $\mathbb{Z}[x]/(x^{n+1})$, with additive basis $1, x, \cdots, x^n$. It follows that $1, L, \cdots, L^n$ is also an additive basis.

Proof: The exact sequence for the pair $(\mathbb{C}\mathrm{P}^n, \mathbb{C}\mathrm{P}^{n-1})$ gives a short exact sequence

$$0 \longrightarrow K(\mathbb{C}\mathrm{P}^n, \mathbb{C}\mathrm{P}^{n-1}) \longrightarrow K(\mathbb{C}\mathrm{P}^n) \xrightarrow{\rho} K(\mathbb{C}\mathrm{P}^{n-1}) \longrightarrow 0$$

We shall prove:

(a_n) $\quad (L-1)^n$ generates the kernel of the restriction map ρ.

Hence if we assume inductively that $K(\mathbb{C}\mathrm{P}^{n-1}) = \mathbb{Z}[L]/(L-1)^n$, then (a_n) and the preceding exact sequence imply that $\{1, L-1, \cdots, (L-1)^n\}$ is an additive basis for $K(\mathbb{C}\mathrm{P}^n)$. Since $(L-1)^{n+1} = 0$ in $K(\mathbb{C}\mathrm{P}^n)$ by (a_{n+1}), it follows that $K(CP^n)$ is the quotient ring $\mathbb{Z}[L]/(L-1)^{n+1}$, completing the induction.

A reason one might expect (a_n) to be true is that the kernel of ρ can be identified with $K(\mathbb{C}\mathrm{P}^n, \mathbb{C}\mathrm{P}^{n-1}) = \widetilde{K}(S^{2n})$ by the short exact sequence, and Bott periodicity implies that the n-fold reduced external product of the generator $L-1$ of $\widetilde{K}(S^2)$ with itself generates $\widetilde{K}(S^{2n})$. To make this rough argument into a proof we will have to relate the external product $\widetilde{K}(S^2)\otimes\cdots\otimes\widetilde{K}(S^2)\to\widetilde{K}(S^{2n})$ to the 'internal' product $K(\mathbb{C}\mathrm{P}^n)\otimes\cdots\otimes K(\mathbb{C}\mathrm{P}^n)\to K(\mathbb{C}\mathrm{P}^n)$.

The space $\mathbb{C}\mathrm{P}^n$ is the quotient of the unit sphere S^{2n+1} in $\mathbb{C}^{n+1}$ under multiplication by scalars in $S^1 \subset \mathbb{C}$. Instead of S^{2n+1} we could equally well take the boundary of the product $D_0^2\times\cdots\times D_n^2$ where D_i^2 is the unit disk in the i^{th} coordinate of $\mathbb{C}^{n+1}$, and we start the count with $i = 0$ for convenience. Then we have

$$\partial(D_0^2\times\cdots\times D_n^2) = \textstyle\bigcup_i (D_0^2\times\cdots\times\partial D_i^2\times\cdots\times D_n^2)$$

The action of S^1 by scalar multiplication respects this decomposition. The orbit space of $D_0^2\times\cdots\times\partial D_i^2\times\cdots\times D_n^2$ under the action is a subspace $C_i \subset \mathbb{C}\mathrm{P}^n$ homeomorphic to the product $D_0^2\times\cdots\times D_n^2$ with the factor D_i^2 deleted. Thus we have a decomposition $\mathbb{C}\mathrm{P}^n = \bigcup_i C_i$ with each C_i homeomorphic to D^{2n} and with $C_i\cap C_j = \partial C_i \cap \partial C_j$ for $i \neq j$.

Consider now $C_0 = D_1^2\times\cdots\times D_n^2$. Its boundary is decomposed into the pieces $\partial_i C_0 = D_1^2\times\cdots\times\partial D_i^2\times\cdots\times D_n^2$. The inclusions $(D_i^2, \partial D_i^2) \subset (C_0, \partial_i C_0) \subset (\mathbb{C}\mathrm{P}^n, C_i)$ give rise to a commutative diagram

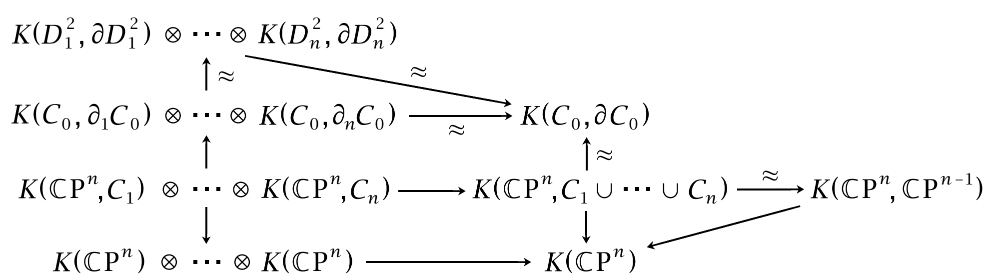

where the maps from the first column to the second are the n-fold products. The upper map in the middle column is an isomorphism because the inclusion $C_0 \hookrightarrow \mathbb{C}\mathrm{P}^n$ induces a homeomorphism $C_0/\partial C_0 \approx \mathbb{C}\mathrm{P}^n/(C_1\cup\cdots\cup C_n)$. The $\mathbb{C}\mathrm{P}^{n-1}$ at the right side of the diagram sits in $\mathbb{C}\mathrm{P}^n$ in the last n coordinates of $\mathbb{C}^{n+1}$, so is disjoint from C_0, hence the quotient map $\mathbb{C}\mathrm{P}^n/\mathbb{C}\mathrm{P}^{n-1}\to\mathbb{C}\mathrm{P}^n/(C_1\cup\cdots\cup C_n)$ is a homotopy equivalence.

The element $x_i \in K(\mathbb{C}\mathrm{P}^n, C_i)$ mapping downward to $L-1 \in K(\mathbb{C}\mathrm{P}^n)$ maps upward to a generator of $K(C_0, \partial_i C_0) \approx K(D_i^2, \partial D_i^2)$. By commutativity of the diagram, the product $x_1\cdots x_n$ then generates $K(\mathbb{C}\mathrm{P}^n, C_1\cup\cdots\cup C_n)$. This means that $(L-1)^n$

generates the image of the map $K(\mathbb{C}P^n,\mathbb{C}P^{n-1})\to K(\mathbb{C}P^n)$, which equals the kernel of ρ, proving (a_n). ☐

Since $\mathbb{C}P^n$ is the union of the $n+1$ balls C_i, Example 2.13 shows that all products of $n+1$ elements of $\widetilde{K}(\mathbb{C}P^n)$ must be zero, in particular $(L-1)^{n+1}=0$. But as we have just seen, $(L-1)^n$ is nonzero, so the result in Example 2.13 is best possible in terms of the degree of nilpotency.

Now we come to the Leray-Hirsch theorem for K–theory, which will be the major theoretical ingredient in the proof of the splitting principle:

Theorem 2.25. *Let $p:E\to B$ 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 $c_1,\cdots,c_k\in K^*(E)$ that restrict to a basis for $K^*(F)$ in each fiber F. If either*
(a) *B is a finite cell complex, or*
(b) *F is a finite cell complex having all cells of even dimension,*
then $K^*(E)$, as a module over $K^*(B)$, is free with basis $\{c_1,\cdots,c_k\}$.

Here the $K^*(B)$-module structure on $K^*(E)$ is defined by $\beta\cdot\gamma=p^*(\beta)\gamma$ for $\beta\in K^*(B)$ and $\gamma\in K^*(E)$. Another way to state the conclusion of the theorem is to say that the map $\Phi:K^*(B)\otimes K^*(F)\to K^*(E)$, $\Phi(\sum_i b_i\otimes i^*(c_i))=\sum_i p^*(b_i)c_i$ for i the inclusion $F\hookrightarrow E$, is an isomorphism.

In the case of the product bundle $E=F\times B$ the classes c_i can be chosen to be the pullbacks under the projection $E\to F$ of a basis for $K^*(F)$. The theorem then asserts that the external product $K^*(F)\otimes K^*(B)\to K^*(F\times B)$ is an isomorphism.

For most of our applications of the theorem either case (a) or case (b) will suffice. The proof of (a) is somewhat simpler than (b), and we include (b) mainly to obtain the splitting principle for vector bundles over arbitrary compact Hausdorff base spaces.

Proof: For a subspace $B'\subset B$ let $E'=p^{-1}(B')$. Then we have a diagram

$$(*)\qquad \begin{array}{ccccccccc}
\longrightarrow & K^*(B,B')\otimes K^*(F) & \longrightarrow & K^*(B)\otimes K^*(F) & \longrightarrow & K^*(B')\otimes K^*(F) & \longrightarrow \\
 & \downarrow\Phi & & \downarrow\Phi & & \downarrow\Phi & \\
\longrightarrow & K^*(E,E') & \longrightarrow & K^*(E) & \longrightarrow & K^*(E') & \longrightarrow
\end{array}$$

where the left-hand Φ is defined by the same formula $\Phi(\sum_i b_i\otimes i^*(c_i))=\sum_i p^*(b_i)c_i$, but with $p^*(b_i)c_i$ referring now to the relative product $K^*(E,E')\times K^*(E)\to K^*(E,E')$. The right-hand Φ is defined using the restrictions of the c_i's to the subspace E'. To see that the diagram $(*)$ commutes, we can interpolate between its two rows the row

$$\longrightarrow K^*(E,E')\otimes K^*(F)\longrightarrow K^*(E)\otimes K^*(F)\longrightarrow K^*(E')\otimes K^*(F)\longrightarrow$$

by factoring Φ as the composition $\sum_i b_i\otimes i^*(c_i)\mapsto\sum_i p^*(b_i)\otimes i^*(c_i)\mapsto\sum_i p^*(b_i)c_i$. The upper squares of the enlarged diagram then commute trivially, and the lower squares commute by Proposition 2.15. The lower row of the diagram is of course exact. The upper row is also exact since we assume $K^*(F)$ is free, and tensoring an exact sequence with a free abelian group preserves exactness, the result of the tensoring operation being simply to replace the given exact sequence by the direct sum of a number of copies of itself.

The proof in case (a) will be by a double induction, first on the dimension of B, then within a given dimension, on the number of cells. The induction starts with the trivial case that B is zero-dimensional, hence a finite discrete set. For the induction step, suppose B is obtained from a subcomplex B' by attaching a cell e^n, and let $E' = p^{-1}(B')$ as above. By induction on the number of cells of B we may assume the right-hand Φ in $(*)$ is an isomorphism. If the left-hand Φ is also an isomorphism, then the five-lemma will imply that the middle Φ is an isomorphism, finishing the induction step.

Let $\varphi:(D^n,S^{n-1})\to(B,B')$ be a characteristic map for the attached n-cell. Since D^n is contractible, the pullback bundle $\varphi^*(E)$ is a product, and so we have a commutative diagram

$$\begin{array}{ccccc} K^*(B,B')\otimes K^*(F) & \xrightarrow{\approx} & K^*(D^n,S^{n-1})\otimes K^*(F) & & \\ \downarrow\Phi & & \downarrow\Phi & \searrow^{\Phi} & \\ K^*(E,E') & \xrightarrow{\approx} & K^*(\varphi^*(E),\varphi^*(E')) & \approx & K^*(D^n\times F,S^{n-1}\times F) \end{array}$$

The two horizontal maps are isomorphisms since φ restricts to a homeomorphism on the interior of D^n, hence induces homeomorphisms $B/B'\approx D^n/S^{n-1}$ and $E/E'\approx \varphi^*(E)/\varphi^*(E')$. Thus the diagram reduces the proof to showing that the right-hand Φ, involving the product bundle $D^n\times F\to D^n$, is an isomorphism.

Consider the diagram $(*)$ with (B,B') replaced by (D^n,S^{n-1}). We may assume the right-hand Φ in $(*)$ is an isomorphism since S^{n-1} has smaller dimension than the original cell complex B. The middle Φ is an isomorphism by the case of zero-dimensional B since D^n deformation retracts to a point. Therefore by the five-lemma the left-hand Φ in $(*)$ is an isomorphism for $(B,B')=(D^n,S^{n-1})$. This finishes the proof in case (a).

In case (b) let us first prove the result for a product bundle $E=F\times B$. In this case Ψ is just the external product, so we are free to interchange the roles of F and B. Thus we may use the diagram $(*)$ with F an arbitrary compact Hausdorff space and B a finite cell complex having all cells of even dimension, obtained by attaching a cell e^n to a subcomplex B'. The upper row of $(*)$ is then an exact sequence since it is obtained from the split short exact sequence $0\to K^*(B,B')\to K^*(B)\to K^*(B')\to 0$ by tensoring with the fixed group $K^*(F)$. If we can show that the left-hand Φ in $(*)$ is an isomorphism, then by induction on the number of cells of B we may assume the right-hand Φ is an isomorphism, so the five-lemma will imply that the middle Φ is also an isomorphism.

To show the left-hand Φ is an isomorphism, note first that $B/B'=S^n$ so we may as well take the pair (B, B') to be (D^n, S^{n-1}). Then the middle Φ in $(*)$ is obviously an isomorphism, so the left-hand Φ will be an isomorphism iff the right-hand Φ is an isomorphism. When the sphere S^{n-1} is even-dimensional we have already shown that Φ is an isomorphism in the remarks following the proof of Lemma 2.17, and the same argument applies also when the sphere is odd-dimensional, since K^1 of an odd-dimensional sphere is K^0 of an even-dimensional sphere.

Now we turn to case (b) for nonproducts. The proof will once again be inductive, but this time we need a more subtle inductive statement since B is just a compact Hausdorff space, not a cell complex. Consider the following condition on a compact subspace $U \subset B$:

> For all compact $V \subset U$ the map $\Phi : K^*(V) \otimes K^*(F) \to K^*(p^{-1}(V))$ is an isomorphism.

If this is satisfied, let us call U *good*. By the special case already proved, each point of B has a compact neighborhood U that is good. Since B is compact, a finite number of these neighborhoods cover B, so by induction it will be enough to show that if U_1 and U_2 are good, then so is $U_1 \cup U_2$.

A compact $V \subset U_1 \cup U_2$ is the union of $V_1 = V \cap U_1$ and $V_2 = V \cap U_2$. Consider the diagram like $(*)$ for the pair (V, V_2). Since $K^*(F)$ is free, the upper row of this diagram is exact. Assuming U_2 is good, the map Φ is an isomorphism for V_2, so Φ will be an isomorphism for V if it is an isomorphism for (V, V_2). The quotient V/V_2 is homeomorphic to $V_1/(V_1 \cap V_2)$ so Φ will be an isomorphism for (V, V_2) if it is an isomorphism for $(V_1, V_1 \cap V_2)$. Now look at the diagram like $(*)$ for $(V_1, V_1 \cap V_2)$. Assuming U_1 is good, the maps Φ are isomorphisms for V_1 and $V_1 \cap V_2$. Hence Φ is an isomorphism for $(V_1, V_1 \cap V_2)$, and the induction step is finished. □

Example 2.26. Let $E \to X$ be a vector bundle with fibers $\mathbb{C}^n$ and compact base X. Then we have an associated projective bundle $p : P(E) \to X$ with fibers $\mathbb{C}P^{n-1}$, where $P(E)$ is the space of lines in E, that is, one-dimensional linear subspaces of fibers of E. Over $P(E)$ there is the canonical line bundle $L \to P(E)$ consisting of the vectors in the lines of $P(E)$. In each fiber $\mathbb{C}P^{n-1}$ of $P(E)$ the classes $1, L, \cdots, L^{n-1}$ in $K^*(P(E))$ restrict to a basis for $K^*(\mathbb{C}P^{n-1})$ by Proposition 2.24. From the Leray-Hirsch theorem we deduce that $K^*(P(E))$ is a free $K^*(X)$-module with basis $1, L, \cdots, L^{n-1}$.

Proof of the Splitting Principle: In the preceding example, the fact that 1 is among the basis elements implies that $p^* : K^*(X) \to K^*(P(E))$ is injective. The pullback bundle $p^*(E) \to P(E)$ contains the line bundle L as a subbundle, hence splits as $L \oplus E'$ for $E' \to P(E)$ the subbundle of $p^*(E)$ orthogonal to L with respect to some choice of inner product. Now repeat the process by forming $P(E')$, splitting off another line bundle from the pullback of E' over $P(E')$. Note that $P(E')$ is the space of pairs of orthogonal lines in fibers of E. After a finite number of repetitions we obtain the flag bundle $F(E) \to X$ described at the end of §1.1, whose points are n-tuples of orthogonal lines in fibers of E, where n is the dimension of E. (If the fibers of E have different dimensions over different components of X, we do the construction for each component separately.) The pullback of E over $F(E)$ splits as a sum of line bundles, and the map $F(E) \to X$ induces an injection on K^* since it is a composition of maps with this property. $\square$

In the preceding Example 2.26 we saw that $K^*(P(E))$ is free as a $K^*(X)$-module, with basis $1, L, \cdots, L^{n-1}$. In order to describe the multiplication in $K^*(P(E))$ one therefore needs only a relation expressing L^n in terms of lower powers of L. Such a relation can be found as follows. The pullback of E over $P(E)$ splits as $L \oplus E'$ for some bundle E' of dimension $n-1$, and the desired relation will be $\lambda^n(E') = 0$. To compute $\lambda^n(E') = 0$ we use the formula $\lambda_t(E) = \lambda_t(L)\lambda_t(E')$ in $K^*(P(E))[t]$, where to simplify notation we let 'E' also denote the pullback of E over $P(E)$. The equation $\lambda_t(E) = \lambda_t(L)\lambda_t(E')$ can be rewritten as $\lambda_t(E') = \lambda_t(E)\lambda_t(L)^{-1}$ where $\lambda_t(L)^{-1} = \sum_i (-1)^i L^i t^i$ since $\lambda_t(L) = 1 + Lt$. Equating coefficients of t^n in the two sides of $\lambda_t(E') = \lambda_t(E)\lambda_t(L)^{-1}$, we get $\lambda^n(E') = \sum_i (-1)^{n-i}\lambda^i(E)L^{n-i}$. The relation $\lambda^n(E') = 0$ can be written as $\sum_i (-1)^i \lambda^i(E) L^{n-i} = 0$, with the coefficient of L^n equal to 1, as desired. The result can be stated in the following form:

Proposition 2.27. *For an n-dimensional vector bundle $E \to X$ the ring $K(P(E))$ is isomorphic to the quotient ring $K^*(X)[L]/\left(\sum_i (-1)^i \lambda^i(E) L^{n-i}\right)$.* $\square$

For example when X is a point we have $P(E) = \mathbb{C}\mathrm{P}^{n-1}$ and $\lambda^i(E) = \mathbb{C}^k$ for $k = \binom{n}{i}$, so the polynomial $\sum_i (-1)^i \lambda^i(E) L^{n-i}$ becomes $(L-1)^n$ and we see that the proposition generalizes the isomorphism $K^*(\mathbb{C}\mathrm{P}^{n-1}) \approx \mathbb{Z}[L]/(L-1)^n)$.

Appendix: Finite Cell Complexes

As we mentioned in the remarks following Proposition 2.23 it is convenient for purposes of the splitting principle to work with spaces slightly more general than finite CW complexes. By a *finite cell complex* we mean a space which has a finite filtration $X_0 \subset X_1 \subset \cdots \subset X_k = X$ where X_0 is a finite discrete set and X_{i+1} is obtained from X_i by attaching a cell e^{n_i} via a map $\varphi_i : S^{n_i - 1} \to X_i$. Thus X_{i+1} is the quotient space of the disjoint union of X_i and a disk D^{n_i} under the identifications $x \sim \varphi_i(x)$ for $x \in \partial D^{n_i} = S^{n_i - 1}$.

Proposition 2.28. *If $p : E \to B$ is a fiber bundle whose fiber F and base B are both finite cell complexes, then E is also a finite cell complex, whose cells are products of cells in B with cells in F.*

Proof: Suppose B is obtained from a subcomplex B' by attaching a cell e^n. By induction on the number of cells of B we may assume that $p^{-1}(B')$ is a finite cell complex. If $\Phi : D^n \to B$ is a characteristic map for e^n then the pullback bundle $\Phi^*(E) \to D^n$ is a product since D^n is contractible. Since F is a finite cell complex, this means that we may obtain $\Phi^*(E)$ from its restriction over S^{n-1} by attaching cells. Hence we may obtain E from $p^{-1}(B')$ by attaching cells. □

Chapter 3
Characteristic Classes

Characteristic classes are cohomology classes in $H^*(B;R)$ associated to vector bundles $E \to B$ by some general rule which applies to all base spaces B. The four classical types of characteristic classes are:

1. Stiefel-Whitney classes $w_i(E) \in H^i(B;\mathbb{Z}_2)$ for a real vector bundle E.
2. Chern classes $c_i(E) \in H^{2i}(B;\mathbb{Z})$ for a complex vector bundle E.
3. Pontryagin classes $p_i(E) \in H^{4i}(B;\mathbb{Z})$ for a real vector bundle E.
4. The Euler class $e(E) \in H^n(B;\mathbb{Z})$ when E is an oriented n-dimensional real vector bundle.

The Stiefel-Whitney and Chern classes are formally quite similar. Pontryagin classes can be regarded as a refinement of Stiefel-Whitney classes when one takes $\mathbb{Z}$ rather than $\mathbb{Z}_2$ coefficients, and the Euler class is a further refinement in the orientable case.

Stiefel-Whitney and Chern classes lend themselves well to axiomatization since in most applications it is the formal properties encoded in the axioms which one uses rather than any particular construction of these classes. The construction we give, using the Leray-Hirsch theorem (proved in §4.D of [AT]), has the virtues of simplicity and elegance, though perhaps at the expense of geometric intuition into what properties of vector bundles these characteristic classes are measuring. There is another definition via obstruction theory which does provide some geometric insights, and this will be described in the Appendix to this chapter.

3.1. Stiefel-Whitney and Chern Classes

Stiefel-Whitney classes are defined for real vector bundles, Chern classes for complex vector bundles. The two cases are quite similar, but for concreteness we shall emphasize the real case, with occasional comments on the minor modifications needed to treat the complex case.

A technical point before we begin: We shall assume without further mention that all base spaces of vector bundles are paracompact, so that the fundamental results of Chapter 1 apply. For the study of characteristic classes this is not an essential restriction since one can always pass to pullbacks over a CW approximation to a given base space, and CW complexes are paracompact.

Axioms and Construction

Here is the main result giving axioms for *Stiefel-Whitney classes*:

Theorem 3.1. *There is a unique sequence of functions* $w_1, w_2, \cdots$ *assigning to each real vector bundle* $E \to B$ *a class* $w_i(E) \in H^i(B;\mathbb{Z}_2)$*, depending only on the isomorphism type of* E*, such that*

(a) $w_i(f^*(E)) = f^*(w_i(E))$ *for a pullback* $f^*(E)$*.*

(b) $w(E_1 \oplus E_2) = w(E_1) \smile w(E_2)$ *for* $w = 1 + w_1 + w_2 + \cdots \in H^*(B;\mathbb{Z}_2)$*.*

(c) $w_i(E) = 0$ *if* $i > \dim E$*.*

(d) *For the canonical line bundle* $E \to \mathbb{R}\mathrm{P}^\infty$*,* $w_1(E)$ *is a generator of* $H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$*.*

The sum $w(E) = 1 + w_1(E) + w_2(E) + \cdots$ is the *total Stiefel-Whitney class*. Note that (c) implies that the sum $1 + w_1(E) + w_2(E) + \cdots$ has only finitely many nonzero terms, so this sum does indeed lie in $H^*(B;\mathbb{Z}_2)$, the direct sum of the groups $H^i(B;\mathbb{Z}_2)$. From the formal identity

$$(1 + w_1 + w_2 + \cdots)(1 + w_1' + w_2' + \cdots) = 1 + (w_1 + w_1') + (w_2 + w_1w_1' + w_2') + \cdots$$

it follows that the formula $w(E_1 \oplus E_2) = w(E_1) \smile w(E_2)$ is just a compact way of writing the relations $w_n(E_1 \oplus E_2) = \sum_{i+j=n} w_i(E_1) \smile w_j(E_2)$, where $w_0 = 1$. This relation is sometimes called the *Whitney sum formula*.

For complex vector bundles there are analogous *Chern classes*:

Theorem 3.2. *There is a unique sequence of functions* $c_1, c_2, \cdots$ *assigning to each complex vector bundle* $E \to B$ *a class* $c_i(E) \in H^{2i}(B;\mathbb{Z})$*, depending only on the isomorphism type of* E*, such that*

(a) $c_i(f^*(E)) = f^*(c_i(E))$ *for a pullback* $f^*(E)$*.*

(b) $c(E_1 \oplus E_2) = c(E_1) \smile c(E_2)$ *for* $c = 1 + c_1 + c_2 + \cdots \in H^*(B;\mathbb{Z})$*.*

(c) $c_i(E) = 0$ *if* $i > \dim E$*.*

(d) *For the canonical line bundle* $E \to \mathbb{C}\mathrm{P}^\infty$*,* $c_1(E)$ *is a generator of* $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ *specified in advance.*

As in the real case, the formula in (b) for the total Chern classes can be rewritten in the form $c_n(E_1 \oplus E_2) = \sum_{i+j=n} c_i(E_1) \smile c_j(E_2)$, where $c_0 = 1$.

Proof of 3.1 and 3.2: Associated to a vector bundle $\pi : E \to B$ with fiber $\mathbb{R}^n$ is the projective bundle $P(\pi) : P(E) \to B$, where $P(E)$ is the space of all lines through the origin in all the fibers of E, and $P(\pi)$ is the natural projection sending each line in $\pi^{-1}(b)$ to $b \in B$. We topologize $P(E)$ as a quotient of the complement of the zero section of E, the quotient obtained by factoring out scalar multiplication in each fiber. Over a neighborhood U in B where E is a product $U \times \mathbb{R}^n$, this quotient is $U \times \mathbb{R}\mathrm{P}^{n-1}$, so $P(E)$ is a fiber bundle over B with fiber $\mathbb{R}\mathrm{P}^{n-1}$.

We would like to apply the Leray-Hirsch theorem for cohomology with $\mathbb{Z}_2$ coefficients to this bundle $P(E) \to B$. To do this we need classes $x_i \in H^i(P(E); \mathbb{Z}_2)$ restricting to generators of $H^i(\mathbb{R}\mathrm{P}^{n-1}; \mathbb{Z}_2)$ in each fiber $\mathbb{R}\mathrm{P}^{n-1}$ for $i = 0, \cdots, n-1$. Recall from the proof of Theorem 1.16 that there is a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber. Projectivizing the map g by deleting zero vectors and then factoring out scalar multiplication produces a map $P(g) : P(E) \to \mathbb{R}\mathrm{P}^\infty$. Let α be a generator of $H^1(\mathbb{R}\mathrm{P}^\infty; \mathbb{Z}_2)$ and let $x = P(g)^*(\alpha) \in H^1(P(E); \mathbb{Z}_2)$. Then the powers x^i for $i = 0, \cdots, n-1$ are the desired classes x_i since a linear injection $\mathbb{R}^n \to \mathbb{R}^\infty$ induces an embedding $\mathbb{R}\mathrm{P}^{n-1} \hookrightarrow \mathbb{R}\mathrm{P}^\infty$ for which α pulls back to a generator of $H^1(\mathbb{R}\mathrm{P}^{n-1}; \mathbb{Z}_2)$, hence α^i pulls back to a generator of $H^i(\mathbb{R}\mathrm{P}^{n-1}; \mathbb{Z}_2)$. Note that any two linear injections $\mathbb{R}^n \to \mathbb{R}^\infty$ are homotopic through linear injections, so the induced embeddings $\mathbb{R}\mathrm{P}^{n-1} \hookrightarrow \mathbb{R}\mathrm{P}^\infty$ of different fibers of $P(E)$ are all homotopic. We showed in the proof of Theorem 1.16 that any two choices of g are homotopic through maps that are linear injections on fibers, so the classes x^i are independent of the choice of g.

The Leray-Hirsch theorem then says that $H^*(P(E); \mathbb{Z}_2)$ is a free $H^*(B; \mathbb{Z}_2)$-module with basis $1, x, \cdots, x^{n-1}$. Consequently, x^n can be expressed uniquely as a linear combination of these basis elements with coefficients in $H^*(B; \mathbb{Z}_2)$. Thus there is a unique relation of the form

$$x^n + w_1(E)x^{n-1} + \cdots + w_n(E) \cdot 1 = 0$$

for certain classes $w_i(E) \in H^i(B; \mathbb{Z}_2)$. Here $w_i(E)x^i$ means $P(\pi)^*(w_i(E)) \smile x^i$, by the definition of the $H^*(B; \mathbb{Z}_2)$-module structure on $H^*(P(E); \mathbb{Z}_2)$. For completeness we define $w_i(E) = 0$ for $i > n$ and $w_0(E) = 1$.

To prove property (a), consider a pullback $f^*(E) = E'$, fitting into the diagram at the right. If $g : E \to \mathbb{R}^\infty$ is a linear injection on fibers then so is $g\widetilde{f}$, and it follows that $P(\widetilde{f})^*$ takes the canonical class $x = x(E)$ for $P(E)$ to the canonical class $x(E')$ for $P(E')$. Then

$$\begin{array}{ccc} E' & \xrightarrow{\ \widetilde{f}\ } & E \\ \big\downarrow{\scriptstyle \pi'} & & \big\downarrow{\scriptstyle \pi} \\ B' & \xrightarrow{\ f\ } & B \end{array}$$

$$\begin{aligned} P(\widetilde{f})^*\Big(\sum_i P(\pi)^*(w_i(E)) \smile x(E)^{n-i}\Big) &= \sum_i P(\widetilde{f})^* P(\pi)^*(w_i(E)) \smile P(\widetilde{f})^*(x(E)^{n-i}) \\ &= \sum_i P(\pi')^* f^*(w_i(E)) \smile x(E')^{n-i} \end{aligned}$$

so the relation $x(E)^n + w_1(E)x(E)^{n-1} + \cdots + w_n(E)\cdot 1 = 0$ defining $w_i(E)$ pulls back to the relation $x(E')^n + f^*(w_1(E))x(E')^{n-1} + \cdots + f^*(w_n(E))\cdot 1 = 0$ defining $w_i(E')$. By the uniqueness of this relation, $w_i(E') = f^*(w_i(E))$.

Proceeding to property (b), the inclusions of E_1 and E_2 into $E_1 \oplus E_2$ give inclusions of $P(E_1)$ and $P(E_2)$ into $P(E_1 \oplus E_2)$ with $P(E_1) \cap P(E_2) = \varnothing$. Let $U_1 = P(E_1 \oplus E_2) - P(E_1)$ and $U_2 = P(E_1 \oplus E_2) - P(E_2)$. These are open sets in $P(E_1 \oplus E_2)$ that deformation retract onto $P(E_2)$ and $P(E_1)$, respectively. A map $g : E_1 \oplus E_2 \to \mathbb{R}^\infty$ which is a linear injection on fibers restricts to such a map on E_1 and E_2, so the canonical class $x \in H^1(P(E_1 \oplus E_2);\mathbb{Z}_2)$ for $E_1 \oplus E_2$ restricts to the canonical classes for E_1 and E_2. If E_1 and E_2 have dimensions m and n, consider the classes $\omega_1 = \sum_j w_j(E_1)x^{m-j}$ and $\omega_2 = \sum_j w_j(E_2)x^{n-j}$ in $H^*(P(E_1 \oplus E_2);\mathbb{Z}_2)$, with cup product $\omega_1\omega_2 = \sum_j [\sum_{r+s=j} w_r(E_1)w_s(E_2)]x^{m+n-j}$. By the definition of the classes $w_j(E_1)$, the class ω_1 restricts to zero in $H^m(P(E_1);\mathbb{Z}_2)$, hence ω_1 pulls back to a class in the relative group $H^m(P(E_1 \oplus E_2), P(E_1);\mathbb{Z}_2) \approx H^m(P(E_1 \oplus E_2), U_2;\mathbb{Z}_2)$, and similarly for ω_2. The following commutative diagram, with $\mathbb{Z}_2$ coefficients understood, then shows that $\omega_1\omega_2 = 0$:

$$\begin{array}{ccc}
H^m(P(E_1\oplus E_2),U_2) \times H^n(P(E_1\oplus E_2),U_1) & \xrightarrow{\smile} & H^{m+n}(P(E_1\oplus E_2),U_1\cup U_2) = 0 \\
\downarrow & & \downarrow \\
H^m(P(E_1\oplus E_2)) \times H^n(P(E_1\oplus E_2)) & \xrightarrow{\smile} & H^{m+n}(P(E_1\oplus E_2))
\end{array}$$

Thus $\omega_1\omega_2 = \sum_j [\sum_{r+s=j} w_r(E_1)w_s(E_2)]x^{m+n-j} = 0$ is the defining relation for the Stiefel-Whitney classes of $E_1 \oplus E_2$, and so $w_j(E_1 \oplus E_2) = \sum_{r+s=j} w_r(E_1)w_s(E_2)$.

Property (c) holds by definition. For (d), recall that the canonical line bundle is $E = \{\,(\ell, v) \in \mathbb{R}\mathrm{P}^\infty \times \mathbb{R}^\infty \mid v \in \ell\,\}$. The map $P(\pi)$ in this case is the identity. The map $g : E \to \mathbb{R}^\infty$ which is a linear injection on fibers can be taken to be $g(\ell, v) = v$. So $P(g)$ is also the identity, hence $x(E)$ is a generator of $H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$. The defining relation $x(E) + w_1(E)\cdot 1 = 0$ then says that $w_1(E)$ is a generator of $H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$.

The proof of uniqueness of the classes w_i will use a general property of vector bundles called the *splitting principle*:

Proposition 3.3. *For each vector bundle $\pi : E \to B$ there is a space $F(E)$ and a map $p : F(E) \to B$ such that the pullback $p^*(E) \to F(E)$ splits as a direct sum of line bundles, and $p^* : H^*(B;\mathbb{Z}_2) \to H^*(F(E);\mathbb{Z}_2)$ is injective.*

Proof: Consider the pullback $P(\pi)^*(E)$ of E via the map $P(\pi) : P(E) \to B$. This pullback contains a natural one-dimensional subbundle $L = \{\,(\ell, v) \in P(E) \times E \mid v \in \ell\,\}$. An inner product on E pulls back to an inner product on the pullback bundle, so we have a splitting of the pullback as a sum $L \oplus L^\perp$ with the orthogonal bundle $L^\perp$ having dimension one less than E. As we have seen, the Leray-Hirsch theorem applies to $P(E) \to B$, so $H^*(P(E);\mathbb{Z}_2)$ is the free $H^*(B;\mathbb{Z}_2)$-module with basis $1, x, \cdots, x^{n-1}$

and in particular the induced map $H^*(B;\mathbb{Z}_2)\to H^*(P(E);\mathbb{Z}_2)$ is injective since one of the basis elements is 1.

This construction can be repeated with $L^\perp\to P(E)$ in place of $E\to B$. After finitely many repetitions we obtain the desired result. □

Looking at this construction a little more closely, $L^\perp$ consists of pairs $(\ell, v)\in P(E)\times E$ with $v\perp\ell$. At the next stage we form $P(L^\perp)$, whose points are pairs (ℓ,ℓ') where ℓ and ℓ' are orthogonal lines in E. Continuing in this way, we see that the final base space $F(E)$ is the space of all orthogonal splittings $\ell_1\oplus\cdots\oplus\ell_n$ of fibers of E as sums of lines, and the vector bundle over $F(E)$ consists of all n-tuples of vectors in these lines. Alternatively, $F(E)$ can be described as the space of all chains $V_1\subset\cdots\subset V_n$ of linear subspaces of fibers of E with $\dim V_i = i$. Such chains are called *flags*, and $F(E)\to B$ is the *flag bundle* associated to E. Note that the description of points of $F(E)$ as flags does not depend on a choice of inner product in E.

Now we can finish the proof of Theorem 3.1. Property (d) determines $w_1(E)$ for the canonical line bundle $E\to\mathbb{R}P^\infty$. Property (c) then determines all the w_i's for this bundle. Since the canonical line bundle is the universal line bundle, property (a) therefore determines the classes w_i for all line bundles. Property (b) extends this to sums of line bundles, and finally the splitting principle implies that the w_i's are determined for all bundles.

For complex vector bundles we can use the same proof, but with $\mathbb{Z}$ coefficients since $H^*(\mathbb{C}P^\infty;\mathbb{Z})\approx\mathbb{Z}[\alpha]$, with α now two-dimensional. The defining relation for the $c_i(E)$'s is modified to be

$$x^n - c_1(E)x^{n-1}+\cdots+(-1)^n c_n(E)\cdot 1 = 0$$

with alternating signs. This is equivalent to changing the sign of α, so it does not affect the proofs of properties (a)–(c), but it has the advantage that the canonical line bundle $E\to\mathbb{C}P^\infty$ has $c_1(E)=\alpha$ rather than $-\alpha$, since the defining relation in this case is $x(E)-c_1(E)\cdot 1 = 0$ and $x(E)=\alpha$. □

Note that in property (d) for Stiefel-Whitney classes we could just as well use the canonical line bundle over $\mathbb{R}P^1$ instead of $\mathbb{R}P^\infty$ since the inclusion $\mathbb{R}P^1\hookrightarrow\mathbb{R}P^\infty$ induces an isomorphism $H^1(\mathbb{R}P^\infty;\mathbb{Z}_2)\approx H^1(\mathbb{R}P^1;\mathbb{Z}_2)$. The analogous remark for Chern classes is valid as well.

Example 3.4. Property (a), the naturality of Stiefel-Whitney classes, implies that a product bundle $E=B\times\mathbb{R}^n$ has $w_i(E)=0$ for $i>0$ since a product is the pullback of a bundle over a point, which must have $w_i=0$ for $i>0$ since a point has trivial cohomology in positive dimensions.

Example 3.5: Stability. Property (b) implies that taking the direct sum of a bundle with a product bundle does not change its Stiefel-Whitney classes. In this sense Stiefel-Whitney classes are *stable*. For example, the tangent bundle TS^n to S^n is stably trivial since its direct sum with the normal bundle to S^n in $\mathbb{R}^{n+1}$, which is a trivial line bundle, produces a trivial bundle. Hence the Stiefel-Whitney classes $w_i(TS^n)$ are zero for $i > 0$.

From the identity

$$(1+w_1+w_2+\cdots)(1+w_1'+w_2'+\cdots) = 1+(w_1+w_1')+(w_2+w_1w_1'+w_2')+\cdots$$

we see that $w(E_1)$ and $w(E_1 \oplus E_2)$ determine $w(E_2)$ since the equations

$$\begin{aligned} &w_1 + w_1' = a_1 \\ &w_2 + w_1w_1' + w_2' = a_2 \\ &\cdots \\ &\textstyle\sum_i w_{n-i}w_i' = a_n \end{aligned}$$

can be solved successively for the w_i''s in terms of the w_i's and a_i's. In particular, if $E_1 \oplus E_2$ is the trivial bundle, then we have the case that $a_i = 0$ for $i > 0$ and so $w(E_1)$ determines $w(E_2)$ uniquely by explicit formulas that can be worked out. For example, $w_1' = -w_1$ and $w_2' = -w_1w_1' - w_2 = w_1^2 - w_2$. Of course for $\mathbb{Z}_2$ coefficients the signs do not matter, but the same reasoning applies to Chern classes, with $\mathbb{Z}$ coefficients.

Example 3.6. Let us illustrate this principle by showing that there is no bundle $E \to \mathbb{R}P^\infty$ whose sum with the canonical line bundle $E_1(\mathbb{R}^\infty)$ is trivial. For we have $w(E_1(\mathbb{R}^\infty)) = 1 + \omega$ where ω is a generator of $H^1(\mathbb{R}P^\infty;\mathbb{Z}_2)$, and hence $w(E)$ must be $(1+\omega)^{-1} = 1+\omega+\omega^2+\cdots$ since we are using $\mathbb{Z}_2$ coefficients. Thus $w_i(E) = \omega^i$, which is nonzero in $H^*(\mathbb{R}P^\infty;\mathbb{Z}_2)$ for all i. However, this contradicts the fact that $w_i(E) = 0$ for $i > \dim E$.

This shows the necessity of the compactness assumption in Proposition 1.4. To further delineate the question, note that Proposition 1.4 says that the restriction $E_1(\mathbb{R}^{n+1})$ of the canonical line bundle to the subspace $\mathbb{R}P^n \subset \mathbb{R}P^\infty$ does have an 'inverse' bundle. In fact, the bundle $E_1^\perp(\mathbb{R}^{n+1})$ consisting of pairs (ℓ, v) where ℓ is a line through the origin in $\mathbb{R}^{n+1}$ and v is a vector orthogonal to ℓ is such an inverse. But for any bundle $E \to \mathbb{R}P^n$ whose sum with $E_1(\mathbb{R}^{n+1})$ is trivial we must have $w(E) = 1+\omega+\cdots+\omega^n$, and since $w_n(E) = \omega^n \neq 0$, E must be at least n-dimensional. So we see there is no chance of choosing such bundles E for varying n so that they fit together to form a single bundle over $\mathbb{R}P^\infty$.

Example 3.7. Let us describe an n-dimensional vector bundle $E \to B$ with $w_i(E)$ nonzero for each $i \le n$. This will be the n-fold Cartesian product $(E_1)^n \to (G_1)^n$ of the canonical line bundle over $G_1 = \mathbb{R}P^\infty$ with itself. This vector bundle is the direct sum $\pi_1^*(E_1) \oplus \cdots \oplus \pi_n^*(E_1)$ where $\pi_i : (G_1)^n \to G_1$ is projection onto the i^{th} factor, so $w((E_1)^n) = \prod_i (1+\alpha_i) \in \mathbb{Z}_2[\alpha_1, \cdots, \alpha_n] \approx H^*((\mathbb{R}P^\infty)^n;\mathbb{Z}_2)$. Hence $w_i((E_1)^n)$ is the i^{th} elementary symmetric polynomial σ_i in the α_j's, the sum of all the products of i different α_j's. For example, if $n = 3$ then $\sigma_1 = \alpha_1+\alpha_2+\alpha_3$, $\sigma_2 = \alpha_1\alpha_2+\alpha_1\alpha_3+\alpha_2\alpha_3$, and $\sigma_3 = \alpha_1\alpha_2\alpha_3$. Since each σ_i with $i \le n$ is nonzero in $\mathbb{Z}_2[\alpha_1, \cdots, \alpha_n]$, we have an n-dimensional bundle whose first n Stiefel-Whitney classes are all nonzero.

The same reasoning applies in the complex case to show that the n-fold Cartesian product of the canonical line bundle over $\mathbb{C}\mathrm{P}^\infty$ has its first n Chern classes nonzero.

In this example we see that the w_i's and c_i's can be identified with elementary symmetric functions, and in fact this can be done in general using the splitting principle. Given an n-dimensional vector bundle $E \to B$ we know that the pullback to $F(E)$ splits as a sum $L_1 \oplus \cdots \oplus L_n \to F(E)$. Letting $\alpha_i = w_1(L_i)$, we see that $w(E)$ pulls back to $w(L_1 \oplus \cdots \oplus L_n) = (1+\alpha_1)\cdots(1+\alpha_n) = 1 + \sigma_1 + \cdots + \sigma_n$, so $w_i(E)$ pulls back to σ_i. Thus we have embedded $H^*(B;\mathbb{Z}_2)$ in a larger ring $H^*(F(E);\mathbb{Z}_2)$ such that $w_i(E)$ becomes the i^{th} elementary symmetric polynomial in the elements $\alpha_1, \cdots, \alpha_n$ of $H^*(F(E);\mathbb{Z}_2)$.

Besides the evident formal similarity between Stiefel-Whitney and Chern classes there is also a direct relation:

Proposition 3.8. *Regarding an n-dimensional complex vector bundle $E \to B$ as a $2n$-dimensional real vector bundle, then $w_{2i+1}(E) = 0$ and $w_{2i}(E)$ is the image of $c_i(E)$ under the coefficient homomorphism $H^{2i}(B;\mathbb{Z}) \to H^{2i}(B;\mathbb{Z}_2)$.*

For example, since the canonical complex line bundle over $\mathbb{C}\mathrm{P}^\infty$ has c_1 a generator of $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$, the same is true for its restriction over $S^2 = \mathbb{C}\mathrm{P}^1$, so by the proposition this 2-dimensional real vector bundle $E \to S^2$ has $w_2(E) \ne 0$.

Proof: The bundle E has two projectivizations $\mathbb{R}\mathrm{P}(E)$ and $\mathbb{C}\mathrm{P}(E)$, consisting of all the real and all the complex lines in fibers of E, respectively. There is a natural projection $p : \mathbb{R}\mathrm{P}(E) \to \mathbb{C}\mathrm{P}(E)$ sending each real line to the complex line containing it, since a real line is all the real scalar multiples of any nonzero vector in it and a complex line is all the complex scalar multiples. This projection p fits into a commutative diagram

$$\begin{array}{ccccc} \mathbb{R}\mathrm{P}^{2n-1} & \longrightarrow & \mathbb{R}\mathrm{P}(E) & \xrightarrow{\mathbb{R}\mathrm{P}(g)} & \mathbb{R}\mathrm{P}^\infty \\ \downarrow & & \downarrow p & & \downarrow \\ \mathbb{C}\mathrm{P}^{n-1} & \longrightarrow & \mathbb{C}\mathrm{P}(E) & \xrightarrow{\mathbb{C}\mathrm{P}(g)} & \mathbb{C}\mathrm{P}^\infty \end{array}$$

where the left column is the restriction of p to a fiber of E and the maps $\mathbb{R}\mathrm{P}(g)$ and $\mathbb{C}\mathrm{P}(g)$ are obtained by projectivizing, over $\mathbb{R}$ and $\mathbb{C}$, a map $g : E \to \mathbb{C}^\infty$ which is a $\mathbb{C}$-linear injection on fibers. It is easy to see that all three vertical maps in this diagram are fiber bundles with fiber $\mathbb{R}\mathrm{P}^1$, the real lines in a complex line. The Leray-Hirsch theorem applies to the bundle $\mathbb{R}\mathrm{P}^\infty \to \mathbb{C}\mathrm{P}^\infty$, with $\mathbb{Z}_2$ coefficients, so if β is the standard generator of $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$, the $\mathbb{Z}_2$-reduction $\overline{\beta} \in H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}_2)$ pulls back to a generator of $H^2(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$, namely the square α^2 of the generator $\alpha \in H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$. Hence the $\mathbb{Z}_2$-reduction $\overline{x}_{\mathbb{C}}(E) = \mathbb{C}\mathrm{P}(g)^*(\overline{\beta}) \in H^2(\mathbb{C}\mathrm{P}(E);\mathbb{Z}_2)$ of the basic class $x_{\mathbb{C}}(E) = \mathbb{C}\mathrm{P}(g)^*(\beta)$ pulls back to the square of the basic class $x_{\mathbb{R}}(E) =$

$\mathbb{R}P(g)^*(\alpha) \in H^1(\mathbb{R}P(E);\mathbb{Z}_2)$. Consequently the $\mathbb{Z}_2$-reduction of the defining relation for the Chern classes of E, which is $\overline{x}_{\mathbb{C}}(E)^n + \overline{c}_1(E)\overline{x}_{\mathbb{C}}(E)^{n-1} + \cdots + \overline{c}_n(E)\cdot 1 = 0$, pulls back to the relation $x_{\mathbb{R}}(E)^{2n} + \overline{c}_1(E)x_{\mathbb{R}}(E)^{2n-2} + \cdots + \overline{c}_n(E)\cdot 1 = 0$, which is the defining relation for the Stiefel-Whitney classes of E. This means that $w_{2i+1}(E) = 0$ and $w_{2i}(E) = \overline{c}_i(E)$. □

Cohomology of Grassmannians

From Example 3.7 and naturality it follows that the universal bundle $E_n \to G_n$ must also have all its Stiefel-Whitney classes $w_1(E_n), \cdots, w_n(E_n)$ nonzero. In fact a much stronger statement is true. Let $f:(\mathbb{R}P^\infty)^n \to G_n$ be the classifying map for the n-fold Cartesion product $(E_1)^n$ of the canonical line bundle E_1, and for notational simplicity let $w_i = w_i(E_n)$. Then the composition

$$\mathbb{Z}_2[w_1,\cdots,w_n] \longrightarrow H^*(G_n;\mathbb{Z}_2) \xrightarrow{f^*} H^*\big((\mathbb{R}P^\infty)^n;\mathbb{Z}_2\big) \approx \mathbb{Z}_2[\alpha_1,\cdots,\alpha_n]$$

sends w_i to σ_i, the i^{th} elementary symmetric polynomial. It is a classical algebraic result that the polynomials σ_i are algebraically independent in $\mathbb{Z}_2[\alpha_1,\cdots,\alpha_n]$. Proofs of this can be found in [van der Waerden, §26] or [Lang, p. 134] for example. Thus the composition $\mathbb{Z}_2[w_1,\cdots,w_n] \to \mathbb{Z}_2[\alpha_1,\cdots,\alpha_n]$ is injective, hence also the map $\mathbb{Z}_2[w_1,\cdots,w_n] \to H^*(G_n;\mathbb{Z}_2)$. In other words, the classes $w_i(E_n)$ generate a polynomial subalgebra $\mathbb{Z}_2[w_1,\cdots,w_n] \subset H^*(G_n;\mathbb{Z}_2)$. This subalgebra is in fact equal to $H^*(G_n;\mathbb{Z}_2)$, and the corresponding statement for Chern classes holds as well:

Theorem 3.9. *$H^*(G_n;\mathbb{Z}_2)$ is the polynomial ring $\mathbb{Z}_2[w_1,\cdots,w_n]$ on the Stiefel-Whitney classes $w_i = w_i(E_n)$ of the universal bundle $E_n \to G_n$. Similarly, in the complex case $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}) \approx \mathbb{Z}[c_1,\cdots,c_n]$ where $c_i = c_i(E_n(\mathbb{C}^\infty))$ for the universal bundle $E_n(\mathbb{C}^\infty) \to G_n(\mathbb{C}^\infty)$.*

The proof we give here for this basic result will be a fairly quick application of the CW structure on G_n constructed at the end of §1.2. A different proof will be given in §3.2 where we also compute the cohomology of G_n with $\mathbb{Z}$ coefficients, which is somewhat more subtle.

Proof: Consider a map $f:(\mathbb{R}P^\infty)^n \to G_n$ which pulls E_n back to the bundle $(E_1)^n$ considered above. We have noted that the image of f^* contains the symmetric polynomials in $\mathbb{Z}_2[\alpha_1,\cdots,\alpha_n] \approx H^*((\mathbb{R}P^\infty)^n;\mathbb{Z}_2)$. The opposite inclusion holds as well, since if $\pi:(\mathbb{R}P^\infty)^n \to (\mathbb{R}P^\infty)^n$ is an arbitrary permutation of the factors, then π pulls $(E_1)^n$ back to itself, so $f\pi \simeq f$, which means that $f^* = \pi^* f^*$, so the image of f^* is invariant under $\pi^*: H^*((\mathbb{R}P^\infty)^n;\mathbb{Z}_2) \to H^*((\mathbb{R}P^\infty)^n;\mathbb{Z}_2)$, but the latter map is just the same permutation of the variables α_i.

To finish the proof in the real case it remains to see that f^* is injective. It suffices to find a CW structure on G_n in which the r-cells are in one-to-one correspondence with monomials $w_1^{r_1}\cdots w_n^{r_n}$ of dimension $r = r_1 + 2r_2 + \cdots + nr_n$, since the number of r-cells in a CW complex X is an upper bound on the dimension of $H^r(X;\mathbb{Z}_2)$ as a $\mathbb{Z}_2$ vector space, and a surjective linear map between finite-dimensional vector spaces is injective if the dimension of the domain is not greater than the dimension of the range.

Monomials $w_1^{r_1}\cdots w_n^{r_n}$ of dimension r correspond to n-tuples $(r_1,\cdots,r_n)$ with $r = r_1 + 2r_2 + \cdots + nr_n$. Such n-tuples in turn correspond to partitions of r into at most n integers, via the correspondence

$$(r_1,\cdots,r_n) \quad\longleftrightarrow\quad r_n \le r_n + r_{n-1} \le \cdots \le r_n + r_{n-1} + \cdots + r_1.$$

Such a partition becomes the sequence $\sigma_1 - 1 \le \sigma_2 - 2 \le \cdots \le \sigma_n - n$, corresponding to the strictly increasing sequence $0 < \sigma_1 < \sigma_2 < \cdots < \sigma_n$. For example, when $n = 3$ we have:

	(r_1, r_2, r_3)	$(\sigma_1 - 1, \sigma_2 - 2, \sigma_3 - 3)$	$(\sigma_1, \sigma_2, \sigma_3)$	dimension
1	0 0 0	0 0 0	1 2 3	0
w_1	1 0 0	0 0 1	1 2 4	1
w_2	0 1 0	0 1 1	1 3 4	2
w_1^2	2 0 0	0 0 2	1 2 5	2
w_3	0 0 1	1 1 1	2 3 4	3
w_1w_2	1 1 0	0 1 2	1 3 5	3
w_1^3	3 0 0	0 0 3	1 2 6	3

The cell structure on G_n constructed in §1.2 has one cell of dimension $(\sigma_1 - 1) + (\sigma_2 - 2) + \cdots + (\sigma_n - n)$ for each increasing sequence $0 < \sigma_1 < \sigma_2 < \cdots < \sigma_n$. So we are done in the real case.

The complex case is entirely similar, keeping in mind that c_i has dimension $2i$ rather than i. The CW structure on $G_n(\mathbb{C}^\infty)$ described in §1.2 also has these extra factors of 2 in the dimensions of its cells. In particular, the cells are all even-dimensional, so the cellular boundary maps for $G_n(\mathbb{C}^\infty)$ are all trivial and the cohomology with $\mathbb{Z}$ coefficients consists of a $\mathbb{Z}$ summand for each cell. Injectivity of f^* then follows from the algebraic fact that a surjective homomorphism between free abelian groups of finite rank is injective if the rank of the domain is not greater than the rank of the range. ☐

One might guess that the monomial $w_1^{r_1}\cdots w_n^{r_n}$ corresponding to a given cell of G_n in the way described above was the cohomology class dual to this cell, represented by the cellular cochain assigning the value 1 to the cell and 0 to all the other cells. This is true for the classes w_i themselves, but unfortunately it is not true in general. For example the monomial w_1^i corresponds to the cell whose associated partition is the trivial partition $i = i$, but the cohomology class dual to this cell is w_i' where $1 + w_1' + w_2' + \cdots$ is the multiplicative inverse of $1 + w_1 + w_2 + \cdots$. If one replaces the basis of monomials by the more geometric basis of cohomology classes dual to cells, the formulas for multiplying these dual classes become rather complicated. In the parallel situation of Chern classes this question has very classical roots in algebraic geometry, and the rules for multiplying cohomology classes dual to cells are part of the so-called Schubert calculus. Accessible expositions of this subject from a modern viewpoint can be found in [Fulton] and [Hiller].

Applications of w_1 and c_1

We saw in §1.1 that the set $\mathrm{Vect}^1(X)$ of isomorphism classes of line bundles over X forms a group with respect to tensor product. We know also that $\mathrm{Vect}^1(X) = [X, G_1(\mathbb{R}^\infty)]$, and $G_1(\mathbb{R}^\infty)$ is just $\mathbb{R}\mathrm{P}^\infty$, an Eilenberg-MacLane space $K(\mathbb{Z}_2, 1)$. It is a basic fact in algebraic topology that $[X, K(G,n)] \approx H^n(X;G)$ when X has the homotopy type of a CW complex; see Theorem 4.56 of [AT], for example. Thus one might ask whether the groups $\mathrm{Vect}^1(X)$ and $H^1(X;\mathbb{Z}_2)$ are isomorphic. For complex line bundles we have $G_1(\mathbb{C}^\infty) = \mathbb{C}\mathrm{P}^\infty$, and this is a $K(\mathbb{Z},2)$, so the corresponding question is whether $\mathrm{Vect}^1_{\mathbb{C}}(X)$ is isomorphic to $H^2(X;\mathbb{Z})$.

Proposition 3.10. *The function $w_1 : \mathrm{Vect}^1(X) \to H^1(X;\mathbb{Z}_2)$ is a homomorphism, and is an isomorphism if X has the homotopy type of a CW complex. The same is also true for $c_1 : \mathrm{Vect}^1_{\mathbb{C}}(X) \to H^2(X;\mathbb{Z})$.*

Proof: The argument is the same in both the real and complex cases, so for definiteness let us describe the complex case. To show that $c_1 : \mathrm{Vect}^1_{\mathbb{C}}(X) \to H^2(X)$ is a homomorphism, we first prove that $c_1(L_1 \otimes L_2) = c_1(L_1) + c_1(L_2)$ for the bundle $L_1 \otimes L_2 \to G_1 \times G_1$ where L_1 and L_2 are the pullbacks of the canonical line bundle $L \to G_1 = \mathbb{C}\mathrm{P}^\infty$ under the projections $p_1, p_2 : G_1 \times G_1 \to G_1$ onto the two factors. Since $c_1(L)$ is the generator α of $H^2(\mathbb{C}\mathrm{P}^\infty)$, we know that $H^*(G_1 \times G_1) \approx \mathbb{Z}[\alpha_1, \alpha_2]$ where $\alpha_i = p_i^*(\alpha) = c_1(L_i)$. The inclusion $G_1 \vee G_1 \subset G_1 \times G_1$ induces an isomorphism on H^2, so to compute $c_1(L_1 \otimes L_2)$ it suffices to restrict to $G_1 \vee G_1$. Over the first G_1 the bundle L_2 is the trivial line bundle, so the restriction of $L_1 \otimes L_2$ over this G_1 is $L_1 \otimes 1 \approx L_1$. Similarly, $L_1 \otimes L_2$ restricts to L_2 over the second G_1. So $c_1(L_1 \otimes L_2)$ restricted to $G_1 \vee G_1$ is $\alpha_1 + \alpha_2$ restricted to $G_1 \vee G_1$. Hence $c_1(L_1 \otimes L_2) = \alpha_1 + \alpha_2 = c_1(L_1) + c_1(L_2)$.

The general case of the formula $c_1(E_1 \otimes E_2) = c_1(E_1) + c_1(E_2)$ for line bundles E_1 and E_2 now follows by naturality: We have $E_1 \approx f_1^*(L)$ and $E_2 \approx f_2^*(L)$ for maps $f_1, f_2 : X \to G_1$. For the map $F = (f_1, f_2) : X \to G_1 \times G_1$ we have $F^*(L_i) = f_i^*(L) \approx E_i$, so

$$\begin{aligned} c_1(E_1 \otimes E_2) &= c_1(F^*(L_1) \otimes F^*(L_2)) = c_1(F^*(L_1 \otimes L_2)) = F^*(c_1(L_1 \otimes L_2)) \\ &= F^*(c_1(L_1) + c_1(L_2)) = F^*(c_1(L_1)) + F^*(c_1(L_2)) \\ &= c_1(F^*(L_1)) + c_1(F^*(L_2)) = c_1(E_1) + c_1(E_2). \end{aligned}$$

As noted above, if X is a CW complex, there is a bijection $[X, \mathbb{C}\mathrm{P}^\infty] \approx H^2(X;\mathbb{Z})$, and the more precise statement is that this bijection is given by the map $[f] \mapsto f^*(u)$ for some class $u \in H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$. The class u must be a generator, otherwise the map would not always be surjective. Which of the two generators we choose for u is not important, so we may take it to be the class α. The map $[f] \mapsto f^*(\alpha)$ factors as the composition $[X, \mathbb{C}P^\infty] \to \mathrm{Vect}^1_{\mathbb{C}}(X) \to H^2(X;\mathbb{Z})$, $[f] \mapsto f^*(L) \mapsto c_1(f^*(L)) = f^*(c_1(L)) = f^*(\alpha)$. The first map in this composition is a bijection, so since the composition is a bijection, the second map c_1 must be a bijection also. ◻

The first Stiefel-Whitney class w_1 is closely related to orientability:

Proposition 3.11. *A vector bundle $E \to X$ is orientable iff $w_1(E) = 0$, assuming that X is homotopy equivalent to a CW complex.*

Thus w_1 can be viewed as the obstruction to orientability of vector bundles. An interpretation of the other classes w_i as obstructions will be given in the Appendix to this chapter.

Proof: Without loss we may assume X is a CW complex. By restricting to path-components we may further assume X is connected. There are natural isomorphisms

$$(*) \qquad H^1(X;\mathbb{Z}_2) \xrightarrow{\approx} \mathrm{Hom}(H_1(X), \mathbb{Z}_2) \xrightarrow{\approx} \mathrm{Hom}(\pi_1(X), \mathbb{Z}_2)$$

from the universal coefficient theorem and the fact that $H_1(X)$ is the abelianization of $\pi_1(X)$. When $X = G_n$ we have $\pi_1(G_n) \approx \mathbb{Z}_2$, and $w_1(E_n) \in H^1(G_n;\mathbb{Z}_2)$ corresponds via $(*)$ to this isomorphism $\pi_1(G_n) \approx \mathbb{Z}_2$ since $w_1(E_n)$ is the unique nontrivial element of $H^1(G_n;\mathbb{Z}_2)$. By naturality of $(*)$ it follows that for any map $f : X \to G_n$, $f^*(w_1(E_n))$ corresponds under $(*)$ to the homomorphism $f_* : \pi_1(X) \to \pi_1(G_n) \approx \mathbb{Z}_2$. Thus if we choose f so that $f^*(E_n)$ is a given vector bundle E, we have $w_1(E)$ corresponding under $(*)$ to the induced map $f_* : \pi_1(X) \to \pi_1(G_n) \approx \mathbb{Z}_2$. Hence $w_1(E) = 0$ iff this f_* is trivial, which is exactly the condition for lifting f to the universal cover $\widetilde{G}_n$, i.e., orientability of E. ◻

3.2. Euler and Pontryagin Classes

A *characteristic class* can be defined to be a function associating to each vector bundle $E \to B$ of dimension n a class $x(E) \in H^k(B;G)$, for some fixed n and k, such that the naturality property $x(f^*(E)) = f^*(x(E))$ is satisfied. In particular, for the universal bundle $E_n \to G_n$ there is the class $x = x(E_n) \in H^k(G_n;G)$. Conversely, any element $x \in H^k(G_n;G)$ defines a characteristic class by the rule $x(E) = f^*(x)$ where $E \approx f^*(E_n)$ for $f:B \to G_n$. Since f is unique up to homotopy, $x(E)$ is well-defined, and it is clear that the naturality property is satisfied. Thus characteristic classes correspond bijectively with cohomology classes of G_n.

With $\mathbb{Z}_2$ coefficients all characteristic classes are simply polynomials in the Stiefel-Whitney classes since we showed in Theorem 3.9 that $H^*(G_n;\mathbb{Z}_2)$ is the polynomial ring $\mathbb{Z}_2[w_1,\cdots,w_n]$. Similarly for complex vector bundles all characteristic classes with $\mathbb{Z}$ coefficients are polynomials in the Chern classes since $H^*(G_n(\mathbb{C});\mathbb{Z}) \approx \mathbb{Z}[c_1,\cdots,c_n]$. Our goal in this section is to describe the more refined characteristic classes for real vector bundles that arise when we take cohomology with integer coefficients rather than $\mathbb{Z}_2$ coefficients.

The main tool we will use will be the Gysin exact sequence associated to an n-dimensional real vector bundle $p:E \to B$. This is an easy consequence of the Thom isomorphism $\Phi:H^i(B) \to H^{i+n}(D(E),S(E))$ defined by $\Phi(b) = p^*(b) \smile c$ for a Thom class $c \in H^n(D(E),S(E))$, a class whose restriction to each fiber is a generator of $H^n(D^n,S^{n-1})$. The map Φ is an isomorphism whenever a Thom class exists, as shown in Corollary 4D.9 of [AT]. Thom classes with $\mathbb{Z}$ coefficients exist for all orientable real vector bundles, and with $\mathbb{Z}_2$ coefficients they exist for all vector bundles. This is shown in Theorem 4D.10 of [AT].

Once one has the Thom isomorphism, this gives the Gysin sequence as the lower row of the following commutative diagram, whose upper row is the exact sequence for the pair $(D(E),S(E))$:

$$\begin{array}{ccccccccc}
\cdots \longrightarrow & H^i(D(E),S(E)) & \xrightarrow{j^*} & H^i(D(E)) & \longrightarrow & H^i(S(E)) & \longrightarrow & H^{i+1}(D(E),S(E)) & \longrightarrow \cdots \\
& \approx\uparrow\Phi & & \approx\uparrow p^* & & \| & & \approx\uparrow\Phi & \\
\cdots \longrightarrow & H^{i-n}(B) & \xrightarrow{\smile e} & H^i(B) & \xrightarrow{p^*} & H^i(S(E)) & \longrightarrow & H^{i-n+1}(B) & \longrightarrow \cdots
\end{array}$$

The vertical map p^* is an isomorphism since p is a homotopy equivalence from $D(E)$ to B. The *Euler class* $e \in H^n(B)$ is defined to be $(p^*)^{-1}j^*(c)$, or in other words the restriction of the Thom class to the zero section of E. The square containing the map $\smile e$ commutes since for $b \in H^{i-n}(B)$ we have $j^*\Phi(b) = j^*(p^*(b) \smile c) = p^*(b) \smile j^*(c)$, which equals $p^*(b \smile e) = p^*(b) \smile p^*(e)$ since $p^*(e) = j^*(c)$. The Euler class can also be defined as the class corresponding to $c \smile c$ under the Thom isomorphism, since $\Phi(e) = p^*(e) \smile c = j^*(c) \smile c = c \smile c$.

As a warm-up application of the Gysin sequence let us use it to give a different proof of Theorem 3.9 computing $H^*(G_n;\mathbb{Z}_2)$ and $H^*(G_n(\mathbb{C});\mathbb{Z})$. Consider first the real case. The proof will be by induction on n using the Gysin sequence for the universal bundle $E_n \xrightarrow{\pi} G_n$. The sphere bundle $S(E_n)$ is the space of pairs (v,ℓ) where ℓ is an n-dimensional linear subspace of $\mathbb{R}^\infty$ and v is a unit vector in ℓ. There is a natural map $p:S(E_n)\to G_{n-1}$ sending (v,ℓ) to the $(n-1)$-dimensional linear subspace $v^\perp\subset\ell$ orthogonal to v. It is an exercise to check that p is a fiber bundle. Its fiber is S^∞, all the unit vectors in $\mathbb{R}^\infty$ orthogonal to a given $(n-1)$-dimensional subspace. Since S^∞ is contractible, p induces an isomorphism on all homotopy groups, hence also on all cohomology groups. Using this isomorphism p^* the Gysin sequence, with $\mathbb{Z}_2$ coefficients, has the form

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

where $e\in H^n(G_n;\mathbb{Z}_2)$ is the $\mathbb{Z}_2$ Euler class.

We show first that $\eta(w_j(E_n)) = w_j(E_{n-1})$. By definition the map η is the composition $H^*(G_n)\to H^*(S(E_n)) \xleftarrow{\approx} H^*(G_{n-1})$ induced from $G_{n-1} \xleftarrow{p} S(E_n) \xrightarrow{\pi} G_n$. The pullback $\pi^*(E_n)$ consists of triples (v,w,ℓ) where $\ell\in G_n$ and $v,w\in\ell$ with v a unit vector. This pullback splits naturally as a sum $L\oplus p^*(E_{n-1})$ where L is the subbundle of triples (v,tv,ℓ), $t\in\mathbb{R}$, and $p^*(E_{n-1})$ consists of the triples (v,w,ℓ) with $w\in v^\perp$. The line bundle L is trivial, having the section (v,v,ℓ). Thus the cohomology homomorphism π^* takes $w_j(E_n))$ to $w_j(L\oplus p^*(E_{n-1})) = w_j(p^*(E_{n-1})) = p^*(w_j(E_{n-1}))$, so $\eta(w_j(E_n)) = w_j(E_{n-1})$.

By induction on n, $H^*(G_{n-1})$ is the polynomial ring on the classes $w_j(E_{n-1})$, $j<n$. The induction can start with the case $n=1$, where $G_1=\mathbb{R}\mathrm{P}^\infty$ and $H^*(\mathbb{R}\mathrm{P}^\infty)\approx \mathbb{Z}_2[w_1]$ since $w_1(E_1)$ is a generator of $H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$. Or we could start with the trivial case $n=0$. Since $\eta(w_j(E_n)) = w_j(E_{n-1})$, the maps η are surjective and the Gysin sequence splits into short exact sequences

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

The image of $\smile e:H^0(G_n)\to H^n(G_n)$ is a $\mathbb{Z}_2$ generated by e. By exactness, this $\mathbb{Z}_2$ is the kernel of $\eta:H^n(G_n)\to H^n(G_{n-1})$. The class $w_n(E_n)$ lies in this kernel since $w_n(E_{n-1})=0$. Moreover, $w_n(E_n)\neq 0$, since if $w_n(E_n)=0$ then w_n is zero for all n-dimensional vector bundles, but the bundle $E\to\mathbb{R}\mathrm{P}^\infty$ which is the direct sum of n copies of the canonical line bundle has total Stiefel-Whitney class $w(E)=(1+\alpha)^n$, where α generates $H^1(\mathbb{R}\mathrm{P}^\infty)$, hence $w_n(E)=\alpha^n\neq 0$. Thus e and $w_n(E_n)$ generate the same $\mathbb{Z}_2$, so $e=w_n(E_n)$.

Now we argue that each element $\xi\in H^k(G_n)$ can be expressed as a unique polynomial in the classes $w_i=w_i(E_n)$, by induction on k. First, $\eta(\xi)$ is a unique polynomial f in the $w_i(E_{n-1})$'s by the basic induction on n. Then $\xi - f(w_1,\cdots,w_{n-1})$ is in $\operatorname{Ker}\eta = \operatorname{Im}(\smile w_n)$, hence has the form $\zeta\smile w_n$ for $\zeta\in H^{k-n}(G_n)$ which is unique since $\smile w_n$ is injective. By induction on k, ζ is a unique polynomial g in the w_i's. Thus we have ξ expressed uniquely as a polynomial $f(w_1,\cdots,w_{n-1})+w_n g(w_1,\cdots,w_n)$.

Since every polynomial in $w_1, \cdots, w_n$ has a unique expression in this form, the theorem follows in the real case.

Virtually the same argument works in the complex case. We noted earlier that complex vector bundles always have a Gysin sequence with $\mathbb{Z}$ coefficients. The only elaboration needed to extend the preceding proof to the complex case is at the step where we showed the $\mathbb{Z}_2$ Euler class is w_n. The argument from the real case shows that c_n is a multiple me for some $m \in \mathbb{Z}$, e being now the $\mathbb{Z}$ Euler class. Then for the bundle $E \to \mathbb{C}\mathrm{P}^\infty$ which is the direct sum of n copies of the canonical line bundle, classified by $f : \mathbb{C}\mathrm{P}^\infty \to G_n(\mathbb{C}^\infty)$, we have $\alpha^n = c_n(E) = f^*(c_n) = mf^*(e)$ in $H^{2n}(\mathbb{C}\mathrm{P}^\infty; \mathbb{Z}) \approx \mathbb{Z}$, with α^n a generator, hence $m = \pm 1$ and $e = \pm c_n$. The rest of the proof goes through without change.

We can also compute $H^*(\widetilde{G}_n; \mathbb{Z}_2)$ where $\widetilde{G}_n$ is the oriented Grassmannian. To state the result, let $\pi : \widetilde{G}_n \to G_n$ be the covering projection, so $\widetilde{E}_n = \pi^*(E_n)$, and let $\widetilde{w}_i = w_i(\widetilde{E}_n) = \pi^*(w_i) \in H^i(\widetilde{G}_n; \mathbb{Z}_2)$, where $w_i = w_i(E_n)$.

Proposition 3.12. *$\pi^* : H^*(G_n; \mathbb{Z}_2) \to H^*(\widetilde{G}_n; \mathbb{Z}_2)$ is surjective with kernel the ideal generated by w_1, hence $H^*(\widetilde{G}_n; \mathbb{Z}_2) \approx \mathbb{Z}_2[\widetilde{w}_2, \cdots, \widetilde{w}_n]$.*

This is just the answer one would hope for. Since $\widetilde{G}_n$ is simply-connected, $\widetilde{w}_1$ has to be zero, so the isomorphism $H^*(\widetilde{G}_n; \mathbb{Z}_2) \approx \mathbb{Z}_2[\widetilde{w}_2, \cdots, \widetilde{w}_n]$ is the simplest thing that could happen.

Proof: The 2-sheeted covering $\pi : \widetilde{G}_n \to G_n$ can be regarded as the unit sphere bundle of a 1-dimensional vector bundle, so we have a Gysin sequence beginning

$$0 \longrightarrow H^0(G_n; \mathbb{Z}_2) \longrightarrow H^0(\widetilde{G}_n; \mathbb{Z}_2) \longrightarrow H^0(G_n; \mathbb{Z}_2) \xrightarrow{\smile x} H^1(G_n; \mathbb{Z}_2)$$

where $x \in H^1(G_n; \mathbb{Z}_2)$ is the $\mathbb{Z}_2$-Euler class. Since $\widetilde{G}_n$ is connected, $H^0(\widetilde{G}_n; \mathbb{Z}_2) \approx \mathbb{Z}_2$ and so the map $\smile x$ is injective, hence $x = w_1$, the only nonzero element of $H^1(G_n; \mathbb{Z}_2)$. Since $H^*(G_n; \mathbb{Z}_2) \approx \mathbb{Z}_2[w_1, \cdots, w_n]$, the map $\smile w_1$ is injective in all dimensions, so the Gysin sequence breaks up into short exact sequences

$$0 \longrightarrow H^i(G_n; \mathbb{Z}_2) \xrightarrow{\smile w_1} H^i(G_n; \mathbb{Z}_2) \xrightarrow{\pi^*} H^i(\widetilde{G}_n; \mathbb{Z}_2) \longrightarrow 0$$

from which the conclusion is immediate. □

The goal for the rest of this section is to determine $H^*(G_n; \mathbb{Z})$ and $H^*(\widetilde{G}_n; \mathbb{Z})$, or in other words, to find all characteristic classes for real vector bundles with $\mathbb{Z}$ coefficients, rather than the $\mathbb{Z}_2$ coefficients used for Stiefel-Whitney classes. It turns out that $H^*(G_n; \mathbb{Z})$, modulo elements of order 2 which are just certain polynomials in Stiefel-Whitney classes, is a polynomial ring $\mathbb{Z}[p_1, p_2, \cdots]$ on certain classes p_i of dimension $4i$, called Pontryagin classes. There is a similar statement for $H^*(\widetilde{G}_n; \mathbb{Z})$, but with one of the Pontryagin classes replaced by an Euler class when n is even.

The Euler Class

Recall that the Euler class $e(E) \in H^n(B;\mathbb{Z})$ of an orientable n-dimensional vector bundle $E \to B$ is the restriction of a Thom class $c \in H^n(D(E), S(E);\mathbb{Z})$ to the zero section, that is, the image of c under the composition

$$H^n(D(E), S(E);\mathbb{Z}) \to H^n(D(E);\mathbb{Z}) \to H^n(B;\mathbb{Z})$$

where the first map is the usual passage from relative to absolute cohomology and the second map is induced by the inclusion $B \hookrightarrow D(E)$ as the zero section. By its definition, $e(E)$ depends on the choice of c. However, the assertion (*) in the construction of a Thom class in Theorem 4D.10 of [AT] implies that c is determined by its restriction to each fiber, and the restriction of c to each fiber is in turn determined by an orientation of the bundle, so in fact $e(E)$ depends only on the choice of an orientation of E. Choosing the opposite orientation changes the sign of c. There are exactly two choices of orientation for each path-component of B.

Here are the basic properties of Euler classes $e(E) \in H^n(B;\mathbb{Z})$ associated to oriented n-dimensional vector bundles $E \to B$:

Proposition 3.13.

(a) *An orientation of a vector bundle $E \to B$ induces an orientation of a pullback bundle $f^*(E)$ such that $e(f^*(E)) = f^*(e(E))$.*

(b) *Orientations of vector bundles $E_1 \to B$ and $E_2 \to B$ determine an orientation of the sum $E_1 \oplus E_2$ such that $e(E_1 \oplus E_2) = e(E_1) \smile e(E_2)$.*

(c) *For an orientable n-dimensional real vector bundle E, the coefficient homomorphism $H^n(B;\mathbb{Z}) \to H^n(B;\mathbb{Z}_2)$ carries $e(E)$ to $w_n(E)$. For an n-dimensional complex vector bundle E there is the relation $e(E) = c_n(E) \in H^{2n}(B;\mathbb{Z})$, for a suitable choice of orientation of E.*

(d) *$e(E) = -e(E)$ if the fibers of E have odd dimension.*

(e) *$e(E) = 0$ if E has a nowhere-zero section.*

Proof: (a) For an n-dimensional vector bundle E, let $E' \subset E$ be the complement of the zero section. A Thom class for E can be viewed as an element of $H^n(E, E';\mathbb{Z})$ which restricts to a generator of $H^n(\mathbb{R}^n, \mathbb{R}^n - \{0\};\mathbb{Z})$ in each fiber $\mathbb{R}^n$. For a pullback $f^*(E)$, we have a map $\tilde{f}: f^*(E) \to E$ which is a linear isomorphism in each fiber, so $\tilde{f}^*(c(E))$ restricts to a generator of $H^n(\mathbb{R}^n, \mathbb{R}^n - \{0\};\mathbb{Z})$ in each fiber $\mathbb{R}^n$ of $f^*(E)$. Thus $\tilde{f}^*(c(E)) = c(f^*(E))$. Passing from relative to absolute cohomology classes and then restricting to zero sections, we get $e(f^*(E)) = f^*(e(E))$.

(b) There is a natural projection $p_1: E_1 \oplus E_2 \to E_1$ which is linear in each fiber, and likewise we have $p_2: E_1 \oplus E_2 \to E_2$. If E_1 is m-dimensional we can view a Thom class $c(E_1)$ as lying in $H^m(E_1, E_1')$ where E_1' is the complement of the zero section in E_1. Similarly we have a Thom class $c(E_2) \in H^n(E_2, E_2')$ if E_2 has dimension n. Then the product $p_1^*(c(E_1)) \smile p_2^*(c(E_2))$ is a Thom class for $E_1 \oplus E_2$ since in each fiber

$\mathbb{R}^m \times \mathbb{R}^n = \mathbb{R}^{m+n}$ we have the cup product

$$H^m(\mathbb{R}^{m+n}, \mathbb{R}^{m+n} - \mathbb{R}^n) \times H^n(\mathbb{R}^{m+n}, \mathbb{R}^{m+n} - \mathbb{R}^m) z \to H^{m+n}(\mathbb{R}^{m+n}, \mathbb{R}^{m+n} - \{0\})$$

which takes generator cross generator to generator by the calculations in Example 3.11 of [AT]. Passing from relative to absolute cohomology and restricting to the zero section, we get the relation $e(E_1 \oplus E_2) = e(E_1) \smile e(E_2)$.

(c) We showed this for the universal bundle in the calculation of the cohomology of Grassmannians a couple pages back, so by the naturality property in (a) it holds for all bundles.

(d) When we defined the Euler class we observed that it could also be described as the element of $H^n(B;\mathbb{Z})$ corresponding to $c \smile c \in H^{2n}(D(E), S(E), \mathbb{Z})$ under the Thom isomorphism. If n is odd, the basic commutativity relation for cup products gives $c \smile c = -c \smile c$, so $e(E) = -e(E)$.

(e) A nowhere-zero section of E gives rise to a section $s : B \to S(E)$ by normalizing vectors to have unit length. Then in the exact sequence

$$H^n(D(E), S(E); \mathbb{Z}) \xrightarrow{j^*} H^n(D(E); \mathbb{Z}) \xrightarrow{i^*} H^n(S(E); \mathbb{Z})$$

the map i^* is injective since the composition $D(E) \longrightarrow B \xrightarrow{s} S(E) \xrightarrow{i} D(E)$ is homotopic to the identity. Since i^* is injective, the map j^* is zero by exactness, and hence $e(E) = 0$ from the definition of the Euler class. □

Consider the tangent bundle TS^n to S^n. This bundle is orientable since its base S^n is simply-connected if $n > 1$, while if $n = 1$, TS^1 is just the product $S^1 \times \mathbb{R}$. When n is odd, $e(TS^n) = 0$ either by part (d) of the proposition since $H^*(S^n;\mathbb{Z})$ has no elements of order two, or by part (e) since there is a nonzero tangent vector field to S^n when n is odd, namely $s(x_1, \cdots, x_{n+1}) = (-x_2, x_1, \cdots, -x_{n+1}, x_n)$. However, when n is even $e(TS^n)$ is nonzero:

Proposition 3.14. *For even n, $e(TS^n)$ is twice a generator of $H^n(S^n;\mathbb{Z})$.*

Proof: Let $E' \subset E = TS^n$ be the complement of the zero section. Under the Thom isomorphism the Euler class $e(TS^n)$ corresponds to the square of a Thom class $c \in H^n(E, E')$, so it suffices to show that c^2 is twice a generator of $H^{2n}(E, E')$. Let $A \subset S^n \times S^n$ consist of the antipodal pairs $(x, -x)$. Define a homeomorphism $f : S^n \times S^n - A \to E$ sending a pair $(x, y) \in S^n \times S^n - A$ to the vector from x to the point of intersection of the line through $-x$ and y with the tangent plane at x. The diagonal $D = \{(x, x)\}$ corresponds under f to the zero section of E. Excision then gives the first of the following isomorphisms:

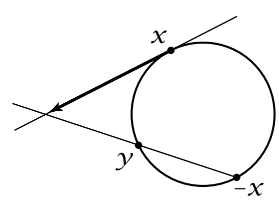

$$H^*(E, E') \approx H^*(S^n \times S^n, S^n \times S^n - D) \approx H^*(S^n \times S^n, A) \approx H^*(S^n \times S^n, D),$$

The second isomorphism holds since $S^n \times S^n - D$ deformation retracts onto A by sliding a point $y \neq \pm x$ along the great circle through x and y to $-x$, and the third comes from the homeomorphism $(x,y) \mapsto (x,-y)$ of $S^n \times S^n$ interchanging D and A. From the long exact sequence of the pair $(S^n \times S^n, D)$ we extract a short exact sequence

$$0 \to H^n(S^n \times S^n, D) \to H^n(S^n \times S^n) \to H^n(D) \to 0$$

The middle group $H^n(S^n \times S^n)$ has generators α, β which are pullbacks of a generator of $H^n(S^n)$ under the two projections $S^n \times S^n \to S^n$. Both α and β restrict to the same generator of $H^n(D)$ since the two projections $S^n \times S^n \to S^n$ restrict to the same homeomorphism $D \approx S^n$, so $\alpha - \beta$ generates $H^n(S^n \times S^n, D)$, the kernel of the restriction map $H^n(S^n \times S^n) \to H^n(D)$. Thus $\alpha - \beta$ corresponds to the Thom class and $(\alpha - \beta)^2 = -\alpha\beta - \beta\alpha$, which equals $-2\alpha\beta$ if n is even. This is twice a generator of $H^{2n}(S^n \times S^n, D) \approx H^{2n}(S^n \times S^n)$. □

It is a fairly elementary theorem in differential topology that the Euler class of the unit tangent bundle of a closed, connected, orientable smooth manifold M^n is $|\chi(M)|$ times a generator of $H^n(M)$, where $\chi(M)$ is the Euler characteristic of M; see for example [Milnor-Stasheff]. This agrees with what we have just seen in the case $M = S^n$, and is the reason for the name 'Euler class.'

One might ask which elements of $H^n(S^n)$ can occur as Euler classes of vector bundles $E \to S^n$ in the nontrivial case that n is even. If we form the pullback of the tangent bundle TS^n by a map $S^n \to S^n$ of degree d, we realize $2d$ times a generator, by part (a) of the preceding proposition, so all even multiples of a generator of $H^n(S^n)$ are realizable. To investigate odd multiples, consider the Thom space $T(E)$. This has integral cohomology consisting of $\mathbb{Z}$'s in dimensions 0, n, and $2n$ by the Thom isomorphism, which also says that the Thom class c is a generator of $H^n(T(E))$. We know that the Euler class corresponds under the Thom isomorphism to $c \smile c$, so $e(E)$ is k times a generator of $H^n(S^n)$ iff $c \smile c$ is k times a generator of $H^{2n}(T(E))$. This is precisely the context of the Hopf invariant, and the solution of the Hopf invariant one problem in Chapter 2 shows that $c \smile c$ can be an odd multiple of a generator only if $n = 2$, 4, or 8. In these three cases there is a bundle $E \to S^n$ for which $c \smile c$ is a generator of $H^{2n}(T(E))$, namely the vector bundle whose unit sphere bundle is the complex, quaternionic, or octonionic Hopf bundle, and whose Thom space, the mapping cone of the sphere bundle, is the associated projective plane $\mathbb{C}\mathrm{P}^2$, $\mathbb{H}\mathrm{P}^2$, or $\mathbb{O}\mathrm{P}^2$. Since we can realize a generator of $H^n(S^n)$ as an Euler class in these three cases, we can realize any multiple of a generator by taking pullbacks as before.

Pontryagin Classes

The easiest definition of the Pontryagin classes $p_i(E) \in H^{4i}(B;\mathbb{Z})$ associated to a real vector bundle $E \to B$ is in terms of Chern classes. For a real vector bundle $E \to B$, its complexification is the complex vector bundle $E^{\mathbb{C}} \to B$ obtained from the real vector bundle $E \oplus E$ by defining scalar multiplication by the complex number i in each fiber $\mathbb{R}^n \oplus \mathbb{R}^n$ via the familiar rule $i(x,y) = (-y,x)$. Thus each fiber $\mathbb{R}^n$ of E becomes a fiber $\mathbb{C}^n$ of $E^{\mathbb{C}}$. The Pontryagin class $p_i(E)$ is then defined to be $(-1)^i c_{2i}(E^{\mathbb{C}}) \in H^{4i}(B;\mathbb{Z})$. The sign $(-1)^i$ is introduced in order to avoid a sign in the formula in (b) of the next proposition. The reason for restricting attention to the even Chern classes $c_{2i}(E^{\mathbb{C}})$ is that the odd classes $c_{2i+1}(E^{\mathbb{C}})$ turn out to be expressible in terms of Stiefel-Whitney classes, and hence give nothing new. The exercises at the end of the section give an explicit formula.

Here is how Pontryagin classes are related to Stiefel-Whitney and Euler classes:

Proposition 3.15.
(a) *For a real vector bundle $E \to B$, $p_i(E)$ maps to $w_{2i}(E)^2$ under the coefficient homomorphism $H^{4i}(B;\mathbb{Z}) \to H^{4i}(B;\mathbb{Z}_2)$.*
(b) *For an orientable real $2n$-dimensional vector bundle with Euler class $e(E) \in H^{2n}(B;\mathbb{Z})$, $p_n(E) = e(E)^2$.*

Note that statement (b) is independent of the choice of orientation of E since the Euler class is squared.

Proof: (a) By Proposition 3.4, $c_{2i}(E^{\mathbb{C}})$ reduces mod 2 to $w_{4i}(E \oplus E)$, which equals $w_{2i}(E)^2$ since $w(E \oplus E) = w(E)^2$ and squaring is an additive homomorphism mod 2.
(b) First we need to determine the relationship between the two orientations of $E^{\mathbb{C}} \approx E \oplus E$, one coming from the canonical orientation of the complex bundle $E^{\mathbb{C}}$, the other coming from the orientation of $E \oplus E$ determined by an orientation of E. If $v_1, \cdots, v_{2n}$ is a basis for a fiber of E agreeing with the given orientation, then $E^{\mathbb{C}}$ is oriented by the ordered basis $v_1, iv_1, \cdots, v_{2n}, iv_{2n}$, while $E \oplus E$ is oriented by $v_1, \cdots, v_{2n}, iv_1, \cdots, iv_{2n}$. To make these two orderings agree requires $(2n-1) + (2n-2) + \cdots + 1 = 2n(2n-1)/2 = n(2n-1)$ transpositions, so the two orientations differ by a sign $(-1)^{n(2n-1)} = (-1)^n$. Thus we have $p_n(E) = (-1)^n c_{2n}(E^{\mathbb{C}}) = (-1)^n e(E^{\mathbb{C}}) = e(E \oplus E) = e(E)^2$. □

Pontryagin classes can be used to describe the cohomology of G_n and $\widetilde{G}_n$ with $\mathbb{Z}$ coefficients. First let us remark that since G_n has a CW structure with finitely many cells in each dimension, so does $\widetilde{G}_n$, hence the homology and cohomology groups of G_n and $\widetilde{G}_n$ are finitely generated. For the universal bundles $E_n \to G_n$ and $\widetilde{E}_n \to \widetilde{G}_n$ let $p_i = p_i(E_n)$, $\widetilde{p}_i = p_i(\widetilde{E}_n)$, and $e = e(\widetilde{E}_n)$, the Euler class being defined via the canonical orientation of $\widetilde{E}_n$.

Theorem 3.16.
(a) *All torsion in* $H^*(G_n;\mathbb{Z})$ *consists of elements of order* 2*, and* $H^*(G_n;\mathbb{Z})/torsion$ *is the polynomial ring* $\mathbb{Z}[p_1,\cdots,p_k]$ *for* $n=2k$ *or* $2k+1$*.*
(b) *All torsion in* $H^*(\widetilde{G}_n;\mathbb{Z})$ *consists of elements of order* 2*, and* $H^*(\widetilde{G}_n;\mathbb{Z})/torsion$ *is* $\mathbb{Z}[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ *for* $n=2k+1$ *and* $\mathbb{Z}[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1},e]$ *for* $n=2k$*, with* $e^2=\widetilde{p}_k$ *in the latter case.*

The torsion subgroup of $H^*(G_n;\mathbb{Z})$ therefore maps injectively to $H^*(G_n;\mathbb{Z}_2)$, with image the image of the Bockstein $\beta: H^*(G_n;\mathbb{Z}_2)\to H^*(G_n;\mathbb{Z}_2)$, which we shall compute in the course of proving the theorem; for the definition and basic properties of Bockstein homomorphisms see §3.E of [AT]. The same remarks apply to $H^*(\widetilde{G}_n;\mathbb{Z})$. The theorem implies that Stiefel-Whitney and Pontryagin classes determine all characteristic classes for unoriented real vector bundles, while for oriented bundles the only additional class needed is the Euler class.

Proof: We shall work on (b) first since for orientable bundles there is a Gysin sequence with $\mathbb{Z}$ coefficients. As a first step we compute $H^*(\widetilde{G}_n;R)$ where $R=\mathbb{Z}[1/2]\subset\mathbb{Q}$, the rational numbers with denominator a power of 2. Since we are dealing with finitely generated integer homology groups, changing from $\mathbb{Z}$ coefficients to R coefficients eliminates any 2-torsion in the homology, that is, elements of order a power of 2, and Z summands of homology become R summands. The assertion to be proved is that $H^*(\widetilde{G}_n;R)$ is $R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ for $n=2k+1$ and $R[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1},e]$ for $n=2k$. This implies that $H^*(\widetilde{G}_n;\mathbb{Z})$ has no odd-order torsion and that $H^*(\widetilde{G}_n;\mathbb{Z})/torsion$ is as stated in the theorem. Then it will remain only to show that all 2-torsion in $H^*(\widetilde{G}_n;\mathbb{Z})$ consists of elements of order 2.

As in the calculation of $H^*(G_n;\mathbb{Z}_2)$ via the Gysin sequence, consider the sphere bundle $S^{n-1}\longrightarrow S(\widetilde{E}_n)\xrightarrow{\pi}\widetilde{G}_n$, where $S(\widetilde{E}_n)$ is the space of pairs (v,ℓ) where ℓ is an oriented n-dimensional linear subspace of $\mathbb{R}^\infty$ and v is a unit vector in ℓ. The orthogonal complement $v^\perp\subset\ell$ of v is then naturally oriented, so we get a projection $p:S(\widetilde{E}_n)\to\widetilde{G}_{n-1}$. The Gysin sequence with coefficients in R has the form

$$\cdots\longrightarrow H^i(\widetilde{G}_n)\xrightarrow{\smile e}H^{i+n}(\widetilde{G}_n)\xrightarrow{\eta}H^{i+n}(\widetilde{G}_{n-1})\longrightarrow H^{i+1}(\widetilde{G}_n)\longrightarrow\cdots$$

where η takes $\widetilde{p}_i(\widetilde{E}_n)$ to $\widetilde{p}_i(\widetilde{E}_{n-1})$.

If $n=2k$, then by induction $H^*(\widetilde{G}_{n-1})\approx R[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1}]$, so η is surjective and the sequence splits into short exact sequences. The proof in this case then follows the $H^*(G_n;\mathbb{Z}_2)$ model.

If $n=2k+1$, then e is zero in $H^n(\widetilde{G}_n;R)$ since with $\mathbb{Z}$ coefficients it has order 2. The Gysin sequence now splits into short exact sequences

$$0\longrightarrow H^{i+n}(\widetilde{G}_n)\xrightarrow{\eta}H^{i+n}(\widetilde{G}_{n-1})\longrightarrow H^{i+1}(\widetilde{G}_n)\longrightarrow 0$$

Thus η injects $H^*(\widetilde{G}_n)$ as a subring of $H^*(\widetilde{G}_{n-1})\approx R[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1},e]$, where e now means $e(\widetilde{E}_{n-1})$. The subring $\operatorname{Im}\eta$ contains $R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ and is torsionfree, so we can show it equals $R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ by comparing ranks of these R-modules in each dimension. Let r_j be the rank of $R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ in dimension j and r'_j the rank of $H^j(\widetilde{G}_n)$. Since $R[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1},e]$ is a free module over $R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ with basis $\{1,e\}$, the rank of $H^*(\widetilde{G}_{n-1})\approx R[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1},e]$ in dimension j is r_j+r_{j-2k}, the class $e=e(\widetilde{E}_{n-1})$ having dimension $2k$. On the other hand, the exact sequence above says this rank also equals $r'_j+r'_{j-2k}$. Since $r'_m\ge r_m$ for each m, we get $r'_j=r_j$, and so $H^*(\widetilde{G}_n)=R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$, completing the induction step. The induction can start with the case $n=1$, with $\widetilde{G}_1\approx S^\infty$.

Before studying the remaining 2-torsion question let us extend what we have just done to $H^*(G_n;\mathbb{Z})$, to show that for $R=\mathbb{Z}[1/2]$, $H^*(G_n;R)$ is $R[p_1,\cdots,p_k]$, where $n=2k$ or $2k+1$. For the 2-sheeted covering $\pi:\widetilde{G}_n\to G_n$ consider the transfer homomorphism $\pi_*:H^*(\widetilde{G}_n;R)\to H^*(G_n;R)$ defined in §3.G of [AT]. The main feature of π_* is that the composition $\pi_*\pi^*:H^*(G_n;R)\to H^*(\widetilde{G}_n;R)\to H^*(G_n;R)$ is multiplication by 2, the number of sheets in the covering space. This is an isomorphism for $R=\mathbb{Z}[1/2]$, so π^* is injective. The image of π^* contains $R[\widetilde{p}_1,\cdots,\widetilde{p}_k]$ since $\pi^*(p_i)=\widetilde{p}_i$. So when n is odd, π^* is an isomorphism and we are done. When n is even, observe that the image of π^* is invariant under the map τ^* induced by the deck transformation $\tau:\widetilde{G}_n\to\widetilde{G}_n$ interchanging sheets of the covering, since $\pi\tau=\pi$ implies $\tau^*\pi^*=\pi^*$. The map τ reverses orientation in each fiber of $\widetilde{E}_n\to\widetilde{G}_n$, so τ^* takes e to $-e$. The subring of $H^*(\widetilde{G}_n;R)\approx R[\widetilde{p}_1,\cdots,\widetilde{p}_{k-1},e]$ invariant under τ^* is then exactly $R[\widetilde{p}_1,\cdots,\widetilde{p}_{[n/2]}]$, finishing the proof that $H^*(G_n;R)=R[p_1,\cdots,p_k]$.

To show that all 2-torsion in $H^*(G_n;\mathbb{Z})$ and $H^*(\widetilde{G}_n;\mathbb{Z})$ has order 2 we use the Bockstein homomorphism β associated to the short exact sequence of coefficient groups $0\to\mathbb{Z}_2\to\mathbb{Z}_4\to\mathbb{Z}_2\to 0$. The goal is to show that $\operatorname{Ker}\beta/\operatorname{Im}\beta$ consists exactly of the mod 2 reductions of nontorsion classes in $H^*(G_n;\mathbb{Z})$ and $H^*(\widetilde{G}_n;\mathbb{Z})$, that is, polynomials in the classes w_{2i}^2 in the case of G_n and $\widetilde{G}_{2k+1}$, and for $\widetilde{G}_{2k}$, polynomials in the w_{2i}^2's for $i<k$ together with w_{2k}, the mod 2 reduction of the Euler class. By general properties of Bockstein homomorphisms proved in §3.E of [AT] this will finish the proof.

Lemma 3.17. $\beta w_{2i+1}=w_1w_{2i+1}$ *and* $\beta w_{2i}=w_{2i+1}+w_1w_{2i}$.

Proof: By naturality it suffices to prove this for the universal bundle $E_n\to G_n$ with $w_i=w_i(E_n)$. As observed in §3.1, we can view w_k as the k^{th} elementary symmetric polynomial σ_k in the polynomial algebra $\mathbb{Z}_2[\alpha_1,\cdots,\alpha_n]\approx H^*((\mathbb{R}\mathrm{P}^\infty)^n;\mathbb{Z}_2)$. Thus to compute βw_k we can compute $\beta\sigma_k$. Using the derivation property $\beta(x\smile y)=\beta x\smile y+x\smile\beta y$ and the fact that $\beta\alpha_i=\alpha_i^2$, we see that $\beta\sigma_k$ is the sum of all products $\alpha_{i1}\cdots\alpha_{ij}^2\cdots\alpha_{ik}$ for $i_1<\cdots<i_k$ and $j=1,\cdots,k$. On the other hand, multiplying $\sigma_1\sigma_k$ out, one obtains $\beta\sigma_k+(k+1)\sigma_{k+1}$. □

Now for the calculation of $\operatorname{Ker}\beta/\operatorname{Im}\beta$. First consider the case of G_{2k+1}. The ring $\mathbb{Z}_2[w_1,\cdots,w_{2k+1}]$ is also the polynomial ring $\mathbb{Z}_2[w_1,w_2,\beta w_2,\cdots,w_{2k},\beta w_{2k}]$ since the substitution $w_1 \mapsto w_1, w_{2i} \mapsto w_{2i}, w_{2i+1} \mapsto w_{2i+1} + w_1 w_{2i} = \beta w_{2i}$ for $i > 0$ is invertible, being its own inverse in fact. Thus $\mathbb{Z}_2[w_1, \cdots, w_{2k+1}]$ splits as the tensor product of the polynomial rings $\mathbb{Z}_2[w_1]$ and $\mathbb{Z}_2[w_{2i}, \beta w_{2i}]$, each of which is invariant under β. Moreover, viewing $\mathbb{Z}_2[w_1, \cdots, w_{2k+1}]$ as a chain complex with boundary map β, this tensor product is a tensor product of chain complexes. According to the algebraic Künneth theorem, the homology of $\mathbb{Z}_2[w_1, \cdots, w_{2k+1}]$ with respect to the boundary map β is therefore the tensor product of the homologies of the chain complexes $\mathbb{Z}_2[w_1]$ and $\mathbb{Z}_2[w_{2i}, \beta w_{2i}]$.

For $\mathbb{Z}_2[w_1]$ we have $\beta(w_1^\ell) = \ell w_1^{\ell+1}$, so $\operatorname{Ker}\beta$ is generated by the even powers of w_1, all of which are also in $\operatorname{Im}\beta$, and hence the β-homology of $\mathbb{Z}_2[w_1]$ is trivial in positive dimensions; we might remark that this had to be true since the calculation is the same as for $\mathbb{R}\mathrm{P}^\infty$. For $\mathbb{Z}_2[w_{2i}, \beta w_{2i}]$ we have $\beta(w_{2i}^\ell(\beta w_{2i})^m) = \ell w_{2i}^{\ell-1}(\beta w_{2i})^{m+1}$, so $\operatorname{Ker}\beta$ is generated by the monomials $w_{2i}^\ell(\beta w_{2i})^m$ with ℓ even, and such monomials with $m > 0$ are in $\operatorname{Im}\beta$. Hence $\operatorname{Ker}\beta/\operatorname{Im}\beta = \mathbb{Z}_2[w_{2i}^2]$.

For $n = 2k$, $\mathbb{Z}_2[w_1, \cdots, w_{2k}]$ is the tensor product of the $\mathbb{Z}_2[w_{2i}, \beta w_{2i}]$'s for $i < k$ and $\mathbb{Z}_2[w_1, w_{2k}]$, with $\beta(w_{2k}) = w_1 w_{2k}$. We then have the formula $\beta(w_1^\ell w_{2k}^m) = \ell w_1^{\ell+1} w_{2k}^m + m w_1^{\ell+1} w_{2k}^m = (\ell + m) w_1^{\ell+1} w_{2k}^m$. For $w_1^\ell w_{2k}^m$ to be in $\operatorname{Ker}\beta$ we must have $\ell + m$ even, and to be in $\operatorname{Im}\beta$ we must have in addition $\ell > 0$. So $\operatorname{Ker}\beta/\operatorname{Im}\beta = \mathbb{Z}_2[w_{2k}^2]$.

Thus the homology of $\mathbb{Z}_2[w_1, \cdots, w_n]$ with respect to β is the polynomial ring in the classes w_{2i}^2, the mod 2 reductions of the Pontryagin classes. By general properties of Bocksteins this finishes the proof of part (a) of the theorem.

The case of $\widetilde{G}_n$ is simpler since $w_1 = 0$, hence $\beta w_{2i} = w_{2i+1}$ and $\beta w_{2i+1} = 0$. Then we can break $\mathbb{Z}_2[w_2, \cdots, w_n]$ up as the tensor product of the chain complexes $\mathbb{Z}_2[w_{2i}, w_{2i+1}]$, plus $\mathbb{Z}_2[w_{2k}]$ when $n = 2k$. The calculations are quite similar to those we have just done, so further details will be left as an exercise. □

Exercises

1. Show that every class in $H^{2k}(\mathbb{C}\mathrm{P}^\infty)$ can be realized as the Euler class of some vector bundle over $\mathbb{C}\mathrm{P}^\infty$ that is a sum of complex line bundles.

2. Show that $c_{2i+1}(E^{\mathbb{C}}) = \beta(w_{2i}(E) w_{2i+1}(E))$.

3. For an oriented $(2k+1)$-dimensional vector bundle E show that $e(E) = \beta w_{2k}(E)$.

Chapter 4

The J-Homomorphism

Homotopy groups of spheres are notoriously difficult to compute, but some partial information can be gleaned from certain naturally defined homomorphisms

$$J:\pi_i(O(n))\rightarrow\pi_{n+i}(S^n)$$

The goal of this chapter is to determine these J-homomorphisms in the stable dimension range $n >> i$ where both domain and range are independent of n, according to Proposition 1.14 for $O(n)$ and the Freudenthal suspension theorem [AT] for S^n. The real form of Bott periodicity proved in Chapter 2 implies that the domain of the stable J-homomorphism $\pi_i(O)\rightarrow\pi_i^s$ is nonzero only for $i = 4n-1$ when $\pi_i(O)$ is $\mathbb{Z}$ and for $i = 8n$ and $8n+1$ when $\pi_i(O)$ is $\mathbb{Z}_2$. In the latter two cases we will show that J is injective. When $i = 4n-1$ the image of J is a finite cyclic group of some order a_n since π_i^s is a finite group for $i > 0$ by a theorem of Serre proved in [SSAT].

The values of a_n have been computed in terms of Bernoulli numbers. Here is a table for small values of n:

n	1	2	3	4	5	6	7	8	9	10	11
a_n	24	240	504	480	264	65520	24	16320	28728	13200	552

In spite of appearances, there is great regularity in this sequence, but this becomes clear only when one looks at the prime factorization of a_n. Here are the rules for computing a_n:

1. The highest power of 2 dividing a_n is $2^{\ell+3}$ where 2^ℓ is the highest power of 2 dividing n.
2. An odd prime p divides a_n iff n is a multiple of $(p-1)/2$, and in this case the highest power of p dividing a_n is $p^{\ell+1}$ where p^ℓ is the highest power of p dividing n.

The first three cases $p = 2, 3, 5$ are shown in the following diagram, where a vertical chain of k connected dots above the number $4n-1$ means that the highest power of p dividing a_n is p^k.

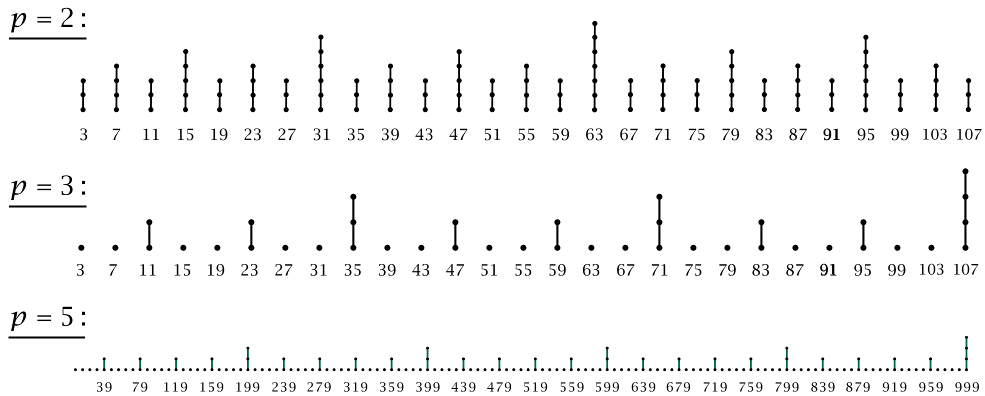

4.1 Lower Bounds for Im J

After starting this section with the definition of the J-homomorphism, we will use a homomorphism $K(X) \to H^*(X;\mathbb{Q})$ known as the Chern character to show that $a_n/2$ is a lower bound on the order of the image of J in dimension $4n-1$. Then using real K-theory this bound will be improved to a_n, and we will take care of the cases in which the domaim of the J-homomorphism is $\mathbb{Z}_2$.

The simplest definition of the J-homomorphism goes as follows. An element $[f] \in \pi_i(O(n))$ is represented by a family of isometries $f_x \in O(n)$, $x \in S^i$, with f_x the identity when x is the basepoint of S^i. Writing S^{n+i} as $\partial(D^{i+1}\times D^n) = S^i \times D^n \cup D^{i+1}\times S^{n-1}$ and S^n as $D^n/\partial D^n$, let $Jf(x,y) = f_x(y)$ for $(x,y) \in S^i \times D^n$ and let $Jf(D^{i+1}\times S^{n-1}) = \partial D^n$, the basepoint of $D^n/\partial D^n$. Clearly $f \simeq g$ implies $Jf \simeq Jg$, so we have a map $J:\pi_i(O(n)) \to \pi_{n+i}(S^n)$. We will tacitly exclude the trivial case $i = 0$.

Proposition 4.1. *J is a homomorphism.*

Proof: We can view Jf as a map $I^{n+i} \to S^n = D^n/\partial D^n$ which on $S^i \times D^n \subset I^{n+1}$ is given by $(x,v) \mapsto f_x(v)$ and which sends the complement of $S^i \times D^n$ to the basepoint ∂D^n. Taking a similar view of Jg, the sum $Jf + Jg$ is obtained by juxtaposing these two maps on either side of a hyperplane. We may assume f_x is the identity for x in the right half of S^i and g_x is the identity for x in the left half of S^i. Then we obtain a homotopy from $Jf + Jg$ to $J(f+g)$ by moving the two $S^i \times D^n$'s together until they coincide, as shown in the figure below. □

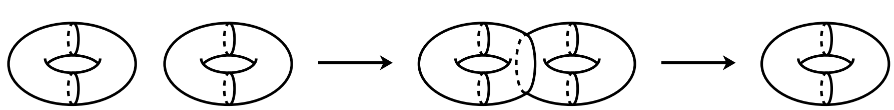

We know that $\pi_i(O(n))$ and $\pi_{n+i}(S^n)$ are independent of n for $n > i+1$, so we would expect the J-homomorphism defined above to induce a stable J–homomorphism $J:\pi_i(O)\to\pi_i^s$, via commutativity of the diagram at the right. We leave it as an exercise for the reader to verify that this is the case.

$$\begin{array}{ccc} \pi_i(O(n)) & \longrightarrow & \pi_i(O(n+1)) \\ \downarrow J & \gamma & \downarrow J \\ \pi_{n+i}(S^n) & \xrightarrow{S} & \pi_{n+i+1}(S^{n+1}) \end{array}$$

Composing the stable J–homomorphism with the map $\pi_i(U)\to\pi_i(O)$ induced by the natural inclusions $U(n)\subset O(2n)$ which give an inclusion $U\subset O$, we get the stable complex J–homomorphism $J_{\mathbb{C}}:\pi_i(U)\to\pi_i^s$. Our goal is to define via K–theory a homomorphism $e:\pi_i^s\to\mathbb{Q}/\mathbb{Z}$ for i odd and compute the composition $eJ_{\mathbb{C}}:\pi_i(U)\to\mathbb{Q}/\mathbb{Z}$. This will give a lower bound for the order of the image of the real J–homomorphism $\pi_i(O)\to\pi_i^s$ when $i=4n-1$.

The Chern Character

The total Chern class $c=1+c_1+c_2+\cdots$ takes direct sums to cup products, and the idea of the Chern character is to form an algebraic combination of Chern classes which takes direct sums to sums and tensor products to cup products, thus giving a natural ring homomorphism from K–theory to cohomology. In order to make this work one must use cohomology with rational coefficients, however. The situation might have been simpler if it had been possible to use integer coefficients instead, but on the other hand, the fact that one has rational coefficients instead of integers makes it possible to define a homomorphism $e:\pi_{2m-1}(S^{2n})\to\mathbb{Q}/\mathbb{Z}$ which gives some very interesting information about the difficult subject of homotopy groups of spheres.

In order to define the Chern character it suffices, via the splitting principle, to do the case of line bundles. The idea is to define the Chern character $ch(L)$ for a line bundle $L\to X$ to be $ch(L)=e^{c_1(L)}=1+c_1(L)+c_1(L)^2/2!+\cdots\in H^*(X;\mathbb{Q})$, so that $ch(L_1\otimes L_2)=e^{c_1(L_1\otimes L_2)}=e^{c_1(L_1)+c_1(L_2)}=e^{c_1(L_1)}e^{c_1(L_2)}=ch(L_1)ch(L_2)$. If the sum $1+c_1(L)+c_1(L)^2/2!+\cdots$ has infinitely many nonzero terms, it will lie not in the direct sum $H^*(X;\mathbb{Q})$ of the groups $H^n(X;\mathbb{Q})$ but rather in the direct product. However, in the examples we shall be considering, $H^n(X;\mathbb{Q})$ will be zero for sufficiently large n, so this distinction will not matter.

For a direct sum of line bundles $E\approx L_1\oplus\cdots\oplus L_n$ we would then want to have

$$ch(E)=\sum_i ch(L_i)=\sum_i e^{t_i}=n+(t_1+\cdots+t_n)+\cdots+(t_1^k+\cdots+t_n^k)/k!+\cdots$$

where $t_i=c_1(L_i)$. The total Chern class $c(E)$ is then $(1+t_1)\cdots(1+t_n)=1+\sigma_1+\cdots+\sigma_n$, where $\sigma_j=c_j(E)$ is the j^{th} elementary symmetric polynomial in the t_i's, the sum of all products of j distinct t_i's. As we saw in §2.3, the Newton polynomials s_k satisfy $t_1^k+\cdots+t_n^k=s_k(\sigma_1,\cdots,\sigma_k)$. Since $\sigma_j=c_j(E)$, this means that the preceding displayed formula can be rewritten

$$ch(E)=\dim E+\sum_{k>0}s_k(c_1(E),\cdots,c_k(E))/k!$$

The right side of this equation is defined for arbitrary vector bundles E, so we take this as our general definition of $ch(E)$.

Proposition 4.2. *$ch(E_1 \oplus E_2) = ch(E_1) + ch(E_2)$ and $ch(E_1 \otimes E_2) = ch(E_1)ch(E_2)$.*

Proof: The proof of the splitting principle for ordinary cohomology in Proposition 2.3 works with any coefficients in the case of complex vector bundles, in particular for $\mathbb{Q}$ coefficients. By this splitting principle we can pull E_1 back to a sum of line bundles over a space $F(E_1)$. By another application of the splitting principle to the pullback of E_2 over $F(E_1)$, we have a map $F(E_1, E_2) \to X$ pulling both E_1 and E_2 back to sums of line bundles, with the induced map $H^*(X;\mathbb{Q}) \to H^*(F(E_1,E_2);\mathbb{Q})$ injective. So to prove the proposition it suffices to verify the two formulas when E_1 and E_2 are sums of line bundles, say $E_i = \oplus_j L_{ij}$ for $i = 1, 2$. The sum formula holds since $ch(E_1 \oplus E_2) = ch(\oplus_{i,j} L_{ij}) = \sum_{i,j} e^{c_1(L_{ij})} = ch(E_1) + ch(E_2)$, by the discussion preceding the definition of ch. For the product formula, $ch(E_1 \otimes E_2) = ch\big(\oplus_{j,k}(L_{1j} \otimes L_{2k})\big) = \sum_{j,k} ch(L_{1j} \otimes L_{2k}) = \sum_{j,k} ch(L_{1j})ch(L_{2k}) = ch(E_1)ch(E_2)$. □

In view of this proposition, the Chern character automatically extends to a ring homomorphism $ch : K(X) \to H^*(X;\mathbb{Q})$. By naturality there is also a reduced form $ch : \widetilde{K}(X) \to \widetilde{H}^*(X;\mathbb{Q})$ since these reduced rings are the kernels of restriction to a point.

As a first calculation of the Chern character, we have:

Proposition 4.3. *$ch : \widetilde{K}(S^{2n}) \to H^{2n}(S^{2n};\mathbb{Q})$ is injective with image equal to the subgroup $H^{2n}(S^{2n};\mathbb{Z}) \subset H^{2n}(S^{2n};\mathbb{Q})$.*

Proof: Since $ch(x \otimes (H-1)) = ch(x) \smile ch(H-1)$ we have the commutative diagram shown at the right, where the upper map is external tensor product with $H - 1$, which is an isomorphism by Bott periodicity, and the lower map is cross product with $ch(H-1) = ch(H) - ch(1) = 1 + c_1(H) - 1 = c_1(H)$, a generator of $H^2(S^2;\mathbb{Z})$. From Theorem 3.16 of [AT] the lower map is an isomorphism and restricts to an isomorphism of the $\mathbb{Z}$-coefficient subgroups. Taking $X = S^{2n}$, the result now follows by induction on n, starting with the trivial case $n = 0$. □

$$\begin{array}{ccc} \widetilde{K}(X) & \xrightarrow{\approx} & \widetilde{K}(S^2X) \\ \downarrow{\scriptstyle ch} & & \downarrow{\scriptstyle ch} \\ \widetilde{H}^*(X;\mathbb{Q}) & \xrightarrow{\approx} & \widetilde{H}^*(S^2X;\mathbb{Q}) \end{array}$$

An interesting by-product of this is:

Corollary 4.4. *A class in $H^{2n}(S^{2n};\mathbb{Z})$ occurs as a Chern class $c_n(E)$ iff it is divisible by $(n-1)!$.*

Proof: For vector bundles $E \to S^{2n}$ we have $c_1(E) = \cdots = c_{n-1}(E) = 0$, so $ch(E) = \dim E + s_n(c_1, \cdots, c_n)/n! = \dim E \pm n c_n(E)/n!$ by the recursion relation for s_n derived in §2.3, namely, $s_n = \sigma_1 s_{n-1} - \sigma_2 s_{n-2} + \cdots + (-1)^{n-2}\sigma_{n-1}s_1 + (-1)^{n-1} n\sigma_n$. □

Even when $H^*(X;\mathbb{Z})$ is torsionfree, so that $H^*(X;\mathbb{Z})$ is a subring of $H^*(X;\mathbb{Q})$, it is not always true that the image of ch is contained in $H^*(X;\mathbb{Z})$. For example, if

$L \in K(\mathbb{C}P^n)$ is the canonical line bundle, then $ch(L) = 1 + c + c^2/2 + \cdots + c^n/n!$ where $c = c_1(L)$ generates $H^2(\mathbb{C}P^n;\mathbb{Z})$, hence c^k generates $H^{2k}(\mathbb{C}P^n;\mathbb{Z})$ for $k \le n$.

The Chern character can be used to show that for finite cell complexes X, the only possible differences between the groups $K^*(X)$ and $H^*(X;\mathbb{Z})$ lie in their torsion subgroups. Since these are finitely generated abelian groups, this will follow if we can show that $K^*(X)\otimes\mathbb{Q}$ and $H^*(X;\mathbb{Q})$ are isomorphic. Thus far we have defined the Chern character $K^0(X)\to H^{even}(X;\mathbb{Q})$, and it is easy to extend this formally to odd dimensions by the commutative diagram at the right.

$$\begin{array}{ccc} \widetilde{K}^0(SX) & \xrightarrow{ch} & \widetilde{H}^{even}(SX;\mathbb{Q}) \\ \| & & \wr\wr \\ K^1(X) & \xrightarrow{ch} & H^{odd}(X;\mathbb{Q}) \end{array}$$

Proposition 4.5. *The map $K^*(X)\otimes\mathbb{Q}\to H^*(X;\mathbb{Q})$ induced by the Chern character is an isomorphism for all finite cell complexes X.*

Proof: We proceed by induction on the number of cells of X. The result is trivially true when there is a single cell, a 0-cell, and it is also true when there are two cells, so that X is a sphere, by the preceding proposition. For the induction step, let X be obtained from a subcomplex A by attaching a cell. Consider the five-term sequence $X/A\to SA\to SX\to SX/SA\to S^2A$. Applying the rationalized Chern character $K^*(-)\otimes\mathbb{Q}\to H^*(-;\mathbb{Q})$ then gives a commutative diagram of five-term exact sequences since tensoring with $\mathbb{Q}$ preserves exactness. The spaces X/A and SX/SA are spheres. Both SA and S^2A are homotopy equivalent to cell complexes with the same number of cells as A, by collapsing the suspension or double suspension of a 0-cell. Thus by induction four of the five maps between the two exact sequences are isomorphisms, all except the map $K^*(SX)\otimes\mathbb{Q}\to H^*(SX;\mathbb{Q})$, so by the five-lemma this map is an isomorphism as well. Finally, to obtain the result for X itself we may replace X by S^2X since the Chern character commutes with double suspension, as we have seen, and a double suspension is in particular a single suspension, with the same number of cells, up to homotopy equivalence. □

The e Invariant

Now let us define the main object we will be studying in this section, the homomorphism $e:\pi_{2m-1}(S^{2n})\to\mathbb{Q}/\mathbb{Z}$. For a map $f:S^{2m-1}\to S^{2n}$ we have the mapping cone C_f obtained by attaching a cell e^{2m} to S^{2n} by f. The quotient C_f/S^{2n} is S^{2m} so we have a commutative diagram of short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \widetilde{K}(S^{2m}) & \longrightarrow & \widetilde{K}(C_f) & \longrightarrow & \widetilde{K}(S^{2n}) & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle ch} & & \downarrow{\scriptstyle ch} & & \downarrow{\scriptstyle ch} & & \\ 0 & \longrightarrow & \widetilde{H}^*(S^{2m};\mathbb{Q}) & \longrightarrow & \widetilde{H}^*(C_f;\mathbb{Q}) & \longrightarrow & \widetilde{H}^*(S^{2n};\mathbb{Q}) & \longrightarrow & 0 \end{array}$$

There are elements $\alpha,\beta \in \widetilde{K}(C_f)$ mapping from and to the standard generators $(H-1)*\cdots*(H-1)$ of $\widetilde{K}(S^{2m})$ and $\widetilde{K}(S^{2n})$, respectively. In a similar way there are elements $a, b \in \widetilde{H}^*(C_f;\mathbb{Q})$ mapping from and to generators of $H^{2m}(S^{2m};\mathbb{Z})$ and $H^{2n}(S^{2n};\mathbb{Z})$. After perhaps replacing a and b by their negatives we may assume that $ch(\alpha) = a$ and $ch(\beta) = b + ra$ for some $r \in \mathbb{Q}$, using Proposition 4.3. The elements β and b are not uniquely determined but can be varied by adding any integer multiples of α and a. The effect of such a variation on the formula $ch(\beta) = b + ra$ is to change r by an integer, so r is well-defined in the additive group $\mathbb{Q}/\mathbb{Z}$, and we define $e(f)$ to be this element $r \in \mathbb{Q}/\mathbb{Z}$. Since $f \simeq g$ implies $C_f \simeq C_g$, we have a well-defined map $e : \pi_{2n-1}(S^{2m}) \to \mathbb{Q}/\mathbb{Z}$.

Proposition 4.6. *e is a homomorphism.*

Proof: Let $C_{f,g}$ be obtained from S^{2n} by attaching two $2m$-cells by f and g, so $C_{f,g}$ contains both C_f and C_g. There is a quotient map $q : C_{f+g} \to C_{f,g}$ collapsing a sphere S^{2m-1} that separates the $2m$-cell of $C_{f,g}$ into a pair of $2m$-cells. In the upper row of the commutative diagram at the right we have generators α_f and α_g mapping to α_{f+g} and $\beta_{f,g}$ mapping to β_{f+g}, and similarly in the second row with generators a_f, a_g, a_{f+g}, $b_{f,g}$, and b_{f+g}. By restriction to the subspaces C_f and C_g of $C_{f,g}$ we obtain $ch(\beta_{f,g}) = b_{f,g} + r_f a_f + r_g b_g$, so $ch(\beta_{f+g}) = b_{f+g} + (r_f + r_g)a_{f+g}$. □

$$\begin{array}{ccc} \widetilde{K}(C_{f,g}) & \xrightarrow{q^*} & \widetilde{K}(C_{f+g}) \\ \downarrow{\scriptstyle ch} & & \downarrow{\scriptstyle ch} \\ \widetilde{H}^*(C_{f,g};\mathbb{Q}) & \xrightarrow{q^*} & \widetilde{H}^*(C_{f+g};\mathbb{Q}) \end{array}$$

There is a commutative diagram involving the double suspension:

$$\begin{array}{ccc} \pi_{2n-1}(S^{2m}) & \xrightarrow{S^2} & \pi_{2n+1}(S^{2m+2}) \\ & {\scriptstyle e}\searrow \quad \nearrow{\scriptstyle e} & \\ & \mathbb{Q}/\mathbb{Z} & \end{array}$$

Commutativity follows from the fact that $C_{S^2 f} = S^2 C_f$ and ch commutes with the double suspension, as we saw in the proof of Proposition 4.3. From the commutativity of the diagram there is induced a stable e-invariant $e : \pi^s_{2k-1} \to \mathbb{Q}/\mathbb{Z}$ for each k.

Theorem 4.7. *If the map $f : S^{2k-1} \to U(n)$ represents a generator of $\pi_{2k-1}(U)$, then $e(J_{\mathbb{C}} f) = \pm\beta_k/k$ where β_k is defined via the power series*

$$x/(e^x - 1) = \textstyle\sum_i \beta_i x^i / i!$$

Hence the image of J in π^s_{2k-1} has order divisible by the denominator of β_k/k.

The numbers β_k are known in number theory as Bernoulli numbers. After proving the theorem we will show how to compute the denominator of β_k/k.

Thom Spaces

For a vector bundle $E \to B$ with unit disk bundle $D(E)$ and unit sphere bundle $S(E)$ the quotient space $D(E)/S(E)$ is called the *Thom space* $T(E)$. When B is compact this can also be regarded as the one-point compactification of E. Thom spaces arise in the present context through the following:

Lemma 4.8. *C_{Jf} is the Thom space of the bundle $E_f \to S^{2k}$ determined by the clutching function $f: S^{2k-1} \to U(n)$.*

Proof: By definition, E_f is the union of two copies of $D^{2k} \times \mathbb{C}^n$ with the subspaces $\partial D^{2k} \times \mathbb{C}^n$ identified via $(x, v) \sim (x, f_x(v))$. Collapsing the second copy of $D^{2k} \times \mathbb{C}^n$ to $\mathbb{C}^n$ via projection produces the same vector bundle E_f, so E_f can also be obtained from $D^{2k} \times \mathbb{C}^n \amalg \mathbb{C}^n$ by the identification $(x, v) \sim f_x(v)$ for $x \in \partial D^{2k}$. Restricting to the unit disk bundle $D(E_f)$, we have $D(E_f)$ expressed as a quotient of $D^{2k} \times D^{2n} \amalg D_0^{2n}$ by the same identification relation, where the subscript 0 labels this particular disk fiber of $D(E_f)$. In the quotient $T(E_f) = D(E_f)/S(E_f)$ we then have the sphere $S^{2n} = D_0^{2n}/\partial D_0^{2n}$, and $T(E_f)$ is obtained from this S^{2n} by attaching a cell e^{2k+2n} with characteristic map the quotient map $D^{2k} \times D^{2n} \to D(E_f) \to T(E_f)$. The attaching map of this cell is precisely Jf, since on $\partial D^{2k} \times D^{2n}$ it is given by $(x, v) \mapsto f_x(v) \in D^{2n}/\partial D^{2n}$ and all of $D^{2k} \times \partial D^{2n}$ maps to the point $\partial D^{2n}/\partial D^{2n}$. □

In order to compute $eJ_{\mathbb{C}}(f)$ we need to compute $ch(\beta)$ where $\beta \in \widetilde{K}(C_{Jf}) = \widetilde{K}(T(E_f))$ restricts to a generator of $\widetilde{K}(S^{2n})$. The S^{2n} here is $D_0^{2n}/\partial D_0^{2n}$ for a fiber D_0^{2n} of $D(E_f)$. For a general complex vector bundle $E \to X$ a class in $\widetilde{K}(T(E))$ that restricts to a generator of $\widetilde{K}(S^n)$ for each sphere S^n coming from a fiber of E is called a *Thom class* for the bundle.

Let us show how to find a Thom class $U \in \widetilde{K}(T(E))$ for a complex vector bundle $E \to X$ with X compact Hausdorff. We may view $T(E)$ as the quotient $P(E \oplus 1)/P(E)$ since in each fiber $\mathbb{C}^n$ of E we obtain $P(\mathbb{C}^n \oplus \mathbb{C}) = \mathbb{C}\mathrm{P}^n$ from $P(\mathbb{C}^n) = \mathbb{C}\mathrm{P}^{n-1}$ by attaching the $2n$-cell $\mathbb{C}^n \times \{1\}$, so the quotient $P(\mathbb{C}^n \oplus \mathbb{C})/P(\mathbb{C}^n)$ is S^{2n}, which is the part of $T(E)$ coming from this fiber $\mathbb{C}^n$. From Example 2.24 we know that $K^*(P(E \oplus 1))$ is the free $K^*(X)$-module with basis $1, L, \cdots, L^n$, where L is the canonical line bundle over $P(E \oplus 1)$. Restricting to $P(E) \subset P(E \oplus 1), K^*(P(E))$ is the free $K^*(X)$-module with basis the restrictions of $1, L, \cdots, L^{n-1}$ to $P(E)$. So we have a short exact sequence

$$0 \longrightarrow \widetilde{K}^*(T(E)) \longrightarrow K^*(P(E \oplus 1)) \xrightarrow{\rho} K^*(P(E)) \longrightarrow 0$$

and $\mathrm{Ker}\,\rho$ must be generated as a $K^*(X)$-module by some polynomial of the form $L^n + a_{n-1}L^{n-1} + \cdots + a_0 1$ with coefficients $a_i \in K^*(X)$, namely the polynomial $\sum_i (-1)^i \lambda^i(E) L^{n-i}$ in Proposition 2.25, regarded now as an element of $K(P(E \oplus 1))$. The class $U \in \widetilde{K}(T(E))$ mapping to $\sum_i (-1)^i \lambda^i(E) L^{n-i}$ is the desired Thom class since when we restrict over a point of X the preceding considerations still apply, so the kernel of $K(\mathbb{C}\mathrm{P}^n) \to K(\mathbb{C}\mathrm{P}^{n-1})$ is generated by the restriction of $\sum_i (-1)^i \lambda^i(E) L^{n-i}$ to a fiber.

Observe that $\widetilde{K}^*(T(E))$ is a free $K^*(X)$-module with the single basis element U since $K^*(P(E \oplus 1))$ is a free $K^*(X)$-module with basis $1, L, \cdots, L^{n-1}, U$. In particular we have an isomorphism $\widetilde{K}^*(T(E)) \approx K^*(X)$, known as the Thom isomorphism.

Similar constructions can be made for ordinary cohomology. The defining relation for $H^*(P(E))$ as $H^*(X)$-module has the form $\sum_i (-1)^i c_i(E) x^{n-i} = 0$ where $x =$

$x(E) \in H^2(P(E))$ restricts to a generator of $H^2(\mathbb{C}P^{n-1})$ in each fiber. Viewed as an element of $H^*(P(E\oplus 1))$, the element $\sum_i(-1)^i c_i(E)x^{n-i}$, with $x = x(E\oplus 1)$ now, generates the kernel of the map to $H^*(P(E))$ since the coefficient of x^n is 1. So $\sum_i(-1)^i c_i(E)x^{n-i} \in H^*(P(E\oplus 1))$ is the image of a Thom class $u \in H^{2n}(T(E))$. For future reference we note two facts:

(1) $x = c_1(L) \in H^*(P(E\oplus 1))$, since the defining relation for $c_1(L)$ is $x(L) - c_1(L) = 0$ and $P(L) = P(E\oplus 1)$, the bundle $L \to E\oplus 1$ being a line bundle, so $x(E\oplus 1) = x(L)$.

(2) If we identify u with $\sum_i(-1)^i c_i(E)x^{n-i} \in H^*(P(E\oplus 1))$, then $xu = 0$ since the defining relation for $H^*(P(E\oplus 1))$ is $\sum_i(-1)^i c_i(E\oplus 1)x^{n+1-i} = 0$ and $c_i(E\oplus 1) = c_i(E)$.

For convenience we shall also identify U with $\sum_i(-1)^i\lambda^i(E)L^{n-i} \in K(P(E\oplus 1))$. We are omitting notation for pullbacks, so in particular we are viewing E as already pulled back over $P(E\oplus 1)$. By the splitting principle we can pull this bundle E back further to a sum $\bigoplus_i L_i$ of line bundles over a space $F(E)$ and work in the cohomology and K–theory of $F(E)$. The Thom class $u = \sum_i(-1)^i c_i(E)x^{n-i}$ then factors as a product $\prod_i(x - x_i)$ where $x_i = c_1(L_i)$, since $c_i(E)$ is the i^{th} elementary symmetric function σ_i of $x_1, \cdots, x_n$. Similarly, for the the K–theory Thom class U we have $U = \sum_i(-1)^i\lambda^i(E)L^{n-i} = L^n\lambda_t(E) = L^n\prod_i\lambda_t(L_i) = L^n\prod_i(1 + L_i t)$ for $t = -L^{-1}$, so $U = \prod_i(L - L_i)$. Therefore we have

$$ch(U) = \textstyle\prod_i ch(L - L_i) = \prod_i(e^x - e^{x_i}) = u\prod_i[(e^{x_i} - e^x)/(x_i - x)]$$

This last expression can be simplified to $u\prod_i[(e^{x_i} - 1)/x_i]$ since after writing it as $u\prod_i e^{x_i}\prod_i[(1 - e^{x - x_i})/(x_i - x)]$ and expanding the last product out as a multivariable power series in x and the x_i's we see that because of the factor u in front and the relation $xu = 0$ noted earlier in (2) all the terms containing x can be deleted, or what amounts to the same thing, we can set $x = 0$.

Since the Thom isomorphism $\Phi : H^*(X) \to H^*(D(E), S(E)) \approx \widetilde{H}^*(T(E))$ is given by cup product with the Thom class u, the result of the preceding calculation can be written as $\Phi^{-1}ch(U) = \prod_i[(e^{x_i} - 1)/x_i]$. When dealing with products such as this it is often convenient to take logarithms. There is a power series for $\log[(e^y - 1)/y]$ of the form $\sum_j \alpha_j y^j/j!$ since the function $(e^y - 1)/y$ has a nonzero value at $y = 0$. Then we have

$$\begin{aligned}\log\Phi^{-1}ch(U) &= \log\textstyle\prod_i[(e^{x_i} - 1)/x_i] = \sum_i \log[(e^{x_i} - 1)/x_i] = \sum_{i,j}\alpha_j x_i^j/j! \\ &= \textstyle\sum_j \alpha_j ch^j(E)\end{aligned}$$

where $ch^j(E)$ is the component of $ch(E)$ in dimension $2j$. Thus we have the general formula $\log\Phi^{-1}ch(U) = \sum_j \alpha_j ch^j(E)$ which no longer involves the splitting of the bundle $E \to X$ into the line bundles L_i, so by the splitting principle this formula is valid back in the cohomology of X.

Proof of 4.7: Let us specialize the preceding to a bundle $E_f \to S^{2k}$ with clutching function $f : S^{2k-1} \to U(n)$ where the earlier dimension m is replaced now by k. As described earlier, the class $\beta \in \widetilde{K}(C_{Jf}) = \widetilde{K}(T(E_f))$ is the Thom class U, up to a sign which we can make $+1$ by rechoosing β if necessary. Since $ch(U) = ch(\beta) = b + ra$, we have $\Phi^{-1}ch(U) = 1 + rh$ where h is a generator of $H^{2k}(S^{2k})$. It follows that $\log \Phi^{-1}ch(U) = rh$ since $\log(1+z) = z - z^2/2 + \cdots$ and $h^2 = 0$. On the other hand, the general formula $\log \Phi^{-1}ch(U) = \sum_j \alpha_j ch^j(E)$ specializes to $\log \Phi^{-1}ch(U) = \alpha_k ch^k(E_f)$ in the present case since $\widetilde{H}^{2j}(S^{2k};\mathbb{Q}) = 0$ for $j \neq k$. If f represents a suitable choice of generator of $\pi_{2k-1}(U(n))$ then $ch^k(E_f) = h$ by Proposition 4.3. Comparing the two calculations of $\log \Phi^{-1}ch(U)$, we obtain $r = \alpha_k$. Since $e(J_{\mathbb{C}}f)$ was defined to be r, we conclude that $e(J_{\mathbb{C}}f) = \alpha_k$ for f representing a generator of $\pi_{2k-1}(U(n))$.

To relate α_k to Bernoulli numbers β_k we differentiate both sides of the equation $\sum_k \alpha_k x^k/k! = \log[(e^x - 1)/x] = \log(e^x - 1) - \log x$, obtaining

$$\begin{aligned}\sum_{k\geq 1} \alpha_k x^{k-1}/(k-1)! &= e^x/(e^x - 1) - x^{-1} = 1 + (e^x - 1)^{-1} - x^{-1} \\ &= 1 - x^{-1} + \sum_{k\geq 0} \beta_k x^{k-1}/k! \\ &= 1 + \sum_{k\geq 1} \beta_k x^{k-1}/k!\end{aligned}$$

where the last equality uses the fact that $\beta_0 = 1$, which comes from the formula $x/(e^x - 1) = \sum_i \beta_i x^i/i!$. Thus we obtain $\alpha_k = \beta_k/k$ for $k > 1$ and $1 + \beta_1 = \alpha_1$. It is not hard to compute that $\beta_1 = -1/2$, so $\alpha_1 = 1/2$ and $\alpha_k = -\beta_k/k$ when $k = 1$. □

Bernoulli Denominators

The numbers β_k are zero for odd $k > 1$ since the function $x/(e^x - 1) - 1 + x/2 = \sum_{i\geq 2} \beta_i x^i/i!$ is even, as a routine calculation shows. Determining the denominator of β_k/k for even k is our next goal since this tells us the order of the image of $eJ_{\mathbb{C}}$ in these cases.

Theorem 4.9. *For even $k > 0$ the denominator of β_k/k is the product of the prime powers $p^{\ell+1}$ such that $p - 1$ divides k and p^{ℓ} is the highest power of p dividing k. More precisely*:

(1) *The denominator of β_k is the product of all the distinct primes p such that $p - 1$ divides k.*

(2) *A prime divides the denominator of β_k/k iff it divides the denominator of β_k.*

The first step in proving the theorem is to relate Bernoulli numbers to the numbers $S_k(n) = 1^k + 2^k + \cdots + (n-1)^k$.

Proposition 4.11. $S_k(n) = \sum_{i=0}^{k} \binom{k}{i} \beta_{k-i} n^{i+1}/(i+1)$.

Proof: The function $f(t) = 1 + e^t + e^{2t} + \cdots + e^{(n-1)t}$ has the power series expansion

$$\sum_{\ell=0}^{n-1} \sum_{k=0}^{\infty} \ell^k t^k/k! = \sum_{k=0}^{\infty} S_k(n) t^k/k!$$

On the other hand, $f(t)$ can be expressed as the product of $(e^{nt}-1)/t$ and $t/(e^t-1)$, with power series

$$\sum_{i=1}^{\infty} n^i t^{i-1}/i! \sum_{j=0}^{\infty} \beta_j t^j/j! = \sum_{i=0}^{\infty} n^{i+1} t^i/(i+1)! \sum_{j=0}^{\infty} \beta_j t^j/j!$$

Equating the coefficients of t^k we get

$$S_k(n)/k! = \sum_{i=0}^{k} n^{i+1}\beta_{k-i}/(i+1)!(k-i)!$$

Multiplying both sides of this equation by $k!$ gives the result. □

Proof of 4.9: We will be interested in the formula for $S_k(n)$ when n is a prime p:

$$(*) \qquad S_k(p) = \beta_k p + \binom{k}{1}\beta_{k-1}p^2/2 + \cdots + \beta_0 p^{k+1}/(k+1)$$

Let $\mathbb{Z}_{(p)} \subset \mathbb{Q}$ be the ring of *p-integers*, that is, rational numbers whose denominators are relatively prime to p. We will first apply $(*)$ to prove that $p\beta_k$ is a p-integer for all primes p. This is equivalent to saying that the denominator of β_k contains no square factors. By induction on k, we may assume $p\beta_{k-i}$ is a p-integer for $i > 0$. Also, $p^i/(i+1)$ is a p-integer since $p^i \geq i+1$ by induction on i. So the product $\beta_{k-i}p^{i+1}/(i+1)$ is a p-integer for $i > 0$. Thus every term except $\beta_k p$ in $(*)$ is a p-integer, and hence $\beta_k p$ is a p-integer as well.

Next we show that for even k, $p\beta_k \equiv S_k(p) \bmod p$ in $\mathbb{Z}_{(p)}$, that is, the difference $p\beta_k - S_k(p)$ is p times a p-integer. This will also follow from $(*)$ once we see that each term after $\beta_k p$ is p times a p-integer. For $i > 1$, $p^{i-1}/(i+1)$ is a p-integer by induction on i as in the preceding paragraph. Since we know $\beta_{k-i}p$ is a p-integer, it follows that each term in $(*)$ containing a β_{k-i} with $i > 1$ is p times a p-integer. As for the term containing β_{k-1}, this is zero if k is even and greater than 2. For $k = 2$, this term is $2(-1/2)p^2/2 = -p^2/2$, which is p times a p-integer.

Now we assert that $S_k(p) \equiv -1 \bmod p$ if $p-1$ divides k, while $S_k(p) \equiv 0 \bmod p$ in the opposite case. In the first case we have

$$S_k(p) = 1^k + \cdots + (p-1)^k \equiv 1 + \cdots + 1 = p - 1 \equiv -1 \bmod p$$

since the multiplicative group $\mathbb{Z}_p^* = \mathbb{Z}_p - \{0\}$ has order $p-1$ and $p-1$ divides k. For the second case we use the elementary fact that $\mathbb{Z}_p^*$ is a cyclic group. (If it were not cyclic, there would exist an exponent $n < p-1$ such that the equation $x^n - 1$ would have $p-1$ roots in $\mathbb{Z}_p$, but a polynomial with coefficients in a field cannot have more roots than its degree.) Let g be a generator of $\mathbb{Z}_p^*$, so $\{1, g^1, g^2, \cdots, g^{p-2}\} = \mathbb{Z}_p^*$. Then

$$S_k(p) = 1^k + \cdots + (p-1)^k = 1^k + g^k + g^{2k} + \cdots + g^{(p-2)k}$$

and hence $(g^k - 1)S_k(p) = g^{(p-1)k} - 1 = 0$ since $g^{p-1} = 1$. If $p-1$ does not divide k then $g^k \neq 1$, so we must have $S_k(p) \equiv 0 \bmod p$.

Statement (1) of the theorem now follows since if $p-1$ does not divide k then $p\beta_k \equiv S_k(p) \equiv 0 \bmod p$ so β_k is p-integral, while if $p-1$ does divide k then $p\beta_k \equiv S_k(p) \equiv -1 \bmod p$ so β_k is not p-integral and p divides the denominator of β_k.

To prove statement (2) of the theorem we will use the following fact:

Lemma 4.12. *For all* $n \in \mathbb{Z}$, $n^k(n^k-1)\beta_k/k$ *is an integer.*

Proof: Recall the function $f(t) = (e^{nt}-1)/(e^t-1)$ considered earlier. This has logarithmic derivative

$$f'(t)/f(t) = (\log f(t))' = [\log(e^{nt}-1) - \log(e^t-1)]' = ne^{nt}/(e^{nt}-1) - e^t/(e^t-1)$$

We have

$$e^x/(e^x-1) = 1/(1-e^{-x}) = x^{-1}[-x/(e^{-x}-1)] = \sum_{k=0}^{\infty}(-1)^k\beta_k x^{k-1}/k!$$

So

$$f'(t)/f(t) = \sum_{k=1}^{\infty}(-1)^k(n^k-1)\beta_k t^{k-1}/k!$$

where the summation starts with $k=1$ since the $k=0$ term is zero. The $(k-1)^{st}$ derivative of this power series at 0 is $\pm(n^k-1)\beta_k/k$. On the other hand, the $(k-1)^{st}$ derivative of $f'(t)(f(t))^{-1}$ is $(f(t))^{-k}$ times a polynomial in $f(t)$ and its derivatives, with integer coefficients, as one can readily see by induction on k. From the formula $f(t) = \sum_{k\geq 0} S_k(n)t^k/k!$ derived earlier, we have $f^{(i)}(0) = S_i(n)$, an integer. So the $(k-1)^{st}$ derivative of $f'(t)/f(t)$ at 0 has the form $m/f(0)^k = m/n^k$ for some $m \in \mathbb{Z}$. Thus $(n^k-1)\beta_k/k = \pm m/n^k$ and so $n^k(n^k-1)\beta_k/k$ is an integer. □

Statement (2) of the theorem can now be proved. If p divides the denominator of β_k then obviously p divides the denominator of β_k/k. Conversely, if p does not divide the denominator of β_k, then by statement (1), $p-1$ does not divide k. Let g be a generator of $\mathbb{Z}_p^*$ as before, so g^k is not congruent to 1 mod p. Then p does not divide $g^k(g^k-1)$, hence β_k/k is the integer $g^k(g^k-1)\beta_k/k$ divided by the number $g^k(g^k-1)$ which is relatively prime to p, so p does not divide the denominator of β_k/k.

The first statement of the theorem follows immediately from (1) and (2). □

There is an alternative definition of e purely in terms of K–theory and the Adams operations ψ^k. By the argument in the proof of Theorem 2.17 there are formulas $\psi^k(\alpha) = k^m\alpha$ and $\psi^k(\beta) = k^n\beta + \mu_k\alpha$ for some $\mu_k \in \mathbb{Z}$ satisfying $\mu_k/(k^m-k^n) = \mu_\ell/(\ell^m-\ell^n)$. The rational number $\mu_k/(k^m-k^n)$ is therefore independent of k. It is easy to check that replacing β by $\beta + p\alpha$ for $p \in \mathbb{Z}$ adds p to $\mu_k/(k^m-k^n)$, so $\mu_k/(k^m-k^n)$ is well-defined in $\mathbb{Q}/\mathbb{Z}$.

Proposition 4.13. $e(f) = \mu_k/(k^m-k^n)$ *in* $\mathbb{Q}/\mathbb{Z}$.

Proof: This follows by computing $ch\,\psi^k(\beta)$ in two ways. First, from the formula for $\psi^k(\beta)$ we have $ch\,\psi^k(\beta) = k^n ch(\beta) + \mu_k ch(\alpha) = k^n b + (k^n r + \mu_k)a$. On the other hand, there is a general formula $ch^q\,\psi^k(\xi) = k^q ch^q(\xi)$ where ch^q denotes the component of ch in H^{2q}. To prove this formula it suffices by the splitting principle and additivity to take ξ to be a line bundle, so $\psi^k(\xi) = \xi^k$, hence

$$ch^q\psi^k(\xi) = ch^q(\xi^k) = [c_1(\xi^k)]^q/q! = [kc_1(\xi)]^q/q! = k^q c_1(\xi)^q/q! = k^q ch^q(\xi)$$

In the case at hand this says $ch^m\psi^k(\beta) = k^m ch^m(\beta) = k^m ra$. Comparing this with the coefficient of a in the first formula for $ch\,\psi^k(\beta)$ gives $\mu_k = r(k^m - k^n)$. □

Bibliography

Books

M. F. ATIYAH, *K–Theory*, W. A. Benjamin, 1967.

R. BOTT, *Lectures on K(X)*, W. A. Benjamin, 1969.

J. DIEUDONNÉ, *A History of Algebraic and Differential Topology 1900-1960*, Birkhäuser, 1989.

W. FULTON, *Young Tableaux*, Cambridge Univ. Press, 1997.

H. HILLER, *Geometry of Coxeter Groups*, Pitman, 1982.

D. HUSEMOLLER, *Fibre Bundles*, McGraw-Hill, 1966 (later editions by Springer-Verlag).

M. KAROUBI, *K–Theory: An Introduction*, Springer-Verlag, 1978.

S. LANG, *Algebra*, Addison-Wesley, 1965.

J. W. MILNOR and J. D. STASHEFF, *Characteristic Classes*, Princeton Univ. Press, 1974.

N. STEENROD, *Topology of Fiber Bundles*, Princeton Univ. Press, 1951.

R. STONG, *Notes on Cobordism Theory*, Princeton Univ. Press, 1968.

B. L. VAN DER WAERDEN, *Modern Algebra*, Ungar, 1949.

Papers

J. F. ADAMS, On the non-existence of elements of Hopf invariant one, *Ann. of Math.* 72 (1960), 20–104.

J. F. ADAMS, Vector fields on spheres, *Ann. of Math.* 75 (1962), 603–632.

J. F. ADAMS, On the groups $J(X)$ IV, *Topology* 5 (1966), 21–71.

J. F. ADAMS and M. F. ATIYAH, K–theory and the Hopf invariant, *Quart. J. Math.* 17 (1966), 31–38.

M. ATIYAH and F. HIRZEBRUCH, Vector Bundles and Homogeneous Spaces, *Proc. Sym. Pure Math.*. III (1961), 7–38.

M. ATIYAH, Vector bundles and the Künneth formula, *Topology* 1 (1962), 245–248.

M. ATIYAH, Power Operations in K–Theory, *Quart. J. Math. Oxford.* 17 (1966), 165–193.

M. ATIYAH, K–Theory and reality, *Quart. J. Math. Oxford.* 17 (1966), 367–386.

M. ATIYAH and R. BOTT, On the periodicity theorem for complex vector bundles, *Acta Math.* 112 (1964), 229–247.

M. ATIYAH, R. BOTT and A. SHAPIRO, Clifford Modules, *Topology* 3 (1964), 3–38.

R. BOTT and J. MILNOR, On the parallelizability of spheres, *Bull. A.M.S.* 64 (1958), 87–89.

R. BOTT, The stable homotopy of the classical groups, *Ann. of Math.* 70 (1959), 313–337.

L. HODGKIN, On the K–Theory of Lie groups, *Topology.* 6 (1967), 1–36.

M. MAHOWALD, The order of the image of the J–homomorphism, *Bull. A.M.S.* 76 (1970), 1310–1313.

J. MILNOR, Some consequences of a theorem of Bott, *Ann. of Math.* 68 (1958), 444–449.

D. QUILLEN, The Adams conjecture, *Topology* 10 (1971), 67–80.

D. SULLIVAN, Genetics of homotopy theory and the Adams conjecture, *Ann. of Math.* 100 (1974), 1–79.

L. M. WOODWARD, Vector fields on spheres and a generalization, *Quart. J. Math. Oxford.* 24 (1973), 357–366.