1 Introduction

The homotopy theory we will discuss in these lectures has historical origins in very concrete geometric problems. We will begin, in this lecture, by describing some of those problems and their solutions. For the most part these solutions date from work by Frank Adams in the 1960's, but the problems themselves are much older.

The central problem we are concerned with involves the construction of interesting sets of tangent vector fields on the $(n-1)$-sphere

$$S^{n-1} = \{x \in \mathbb{R}^n : \|x\| = 1\} .$$

This embedding of the sphere into Euclidean space provides us with a concrete way to visualize tangent vectors: A tangent vector at x is a vector $v \in \mathbb{R}^n$ such that $x \cdot v = 0$. The tangent space to S^{n-1} at a point x is the orthogonal complement in $\mathbb{R}^n$ of the line through x. Write $T_x S^{n-1}$, or just T_x, for this vector space.

The identification of the tangent space at a point with a subspace of $\mathbb{R}^n$ provides us with an inner product on each tangent space. This inner product varies continuously with the point on the sphere; it is a "metric" on the tangent bundle. In particular each tangent vector v has a length $||v||$.

A *vector field* on S^{n-1} is a continuous section of the tangent bundle. In prosaic terms, this is a continuous function

$$v : S^{n-1} \to \mathbb{R}^n$$

such that

$$x \cdot v(x) = 0 \quad \text{for all} \quad x \in S^{n-1} .$$

A first question is whether there exists a *nowhere vanishing* vector field. If we have a nonwhere vanishing vector field, we may normalize it by dividing by its length (which is a continuous function on S^{n-1}) to obtain a *unit* tangent vector field: a map $v : S^{n-1} \to S^{n-1}$ such that $x \cdot v(x) = 0$ for all $x \in S^{n-1}$. There are advantages to replacing the open condition of being everywhere nonzero with the closed condition of being everywhere of length 1, just as there are advantages to thinking about the unit sphere S^{n-1} in place of the noncompact set of nonzero vectors in $\mathbb{R}^n$.

The answer to this question is very well known.

Proposition 1.1. *S^{n-1} admits a unit tangent vector field if and only if n is even.*

Proof. Suppose first that n is even; say $n = 2m$. Then S^{n-1} can be thought of as the unit vectors in $\mathbb{C}^m$, and $v(x) = ix$ satisfies $||v(x)|| = 1$ and $x \cdot v(x) = 0$ for all $x \in S^{n-1}$. This gives us a nonvanishing vector field on each odd sphere.

However, when n is odd there are none. Such a $v(x)$ would give a homotopy between the identity and the antipodal map $\alpha(x) = -x$:

$$h_t(x) = (\cos \pi t)x + (\sin \pi t)v(x)\,.$$

We can get a contradiction by using the *degree*, i.e. the effect of a self-map of S^{n-1} on the $(n-1)$-dimensional integral homology group $H_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}$. The degree of the identity is 1, so by homotopy invariance $\deg \alpha = 1$. But α is a composite of n reflections in $\mathbb{R}^n$, each of degree -1, so $\deg \alpha = (-1)^n = -1$, giving a contradiction. □

A more refined question now arises: for a given positive integer k, can we produce a sequence $(v_1, v_2, \ldots, v_{k-1})$ of $(k-1)$ vector fields on S^{n-1} that are everywhere linearly independent? How big can we make k and still get an affirmative answer to this question?

Once again, we can replace this by an equivalent compact question, this time using the Gram-Schmidt process (which continuous, so we can apply it fiber-wise): find a sequence $v_1, v_2, \ldots, v_{k-1}$ of everywhere orthonormal vector fields on S^{n-1}.

And, once again, in certain cases there are rather obvious constructions. If $n = 4m$, for example, we can regard S^{n-1} as the unit vectors in quaternionic m-space, and then

$$v_1(x) = ix \quad , \quad v_2(x) = jx \quad , \quad v_3(x) = ijx$$

provide three orthonormal vector fields. Similarly, the "octonion" multiplication on $\mathbb{R}^8$ provides seven everywhere independent vector fields on S^7. In these cases, the vector fields provide an isomorphism $S^{n-1} \times \mathbb{R}^{n-1} \to TS^{n-1}$ of vector bundles; that is, a *trivialization* of the tangent bundle or a *parallelization* of the sphere. In fact, Kervaire proved that these are the only spheres that are parallelizable.

While these questions are clearly interesting in their own right, we shall see, in this course, that they are equivalent to fundamental questions about the structure of unstable homotopy theory, especially the behavior of the homotopy groups of spheres.

For a start, notice that they can be phrased in terms of standard questions about the existence of sections of certain fiber bundles. To define these fiber bundles, recall that a *k-frame* (or an *orthogonal k-frame*) in $\mathbb{R}^n$ is a sequence $(v_1, \ldots, v_k)$ of mutually orthogonal unit vectors. Regarding them as column vectors, they can be assembled into an $n \times k$ matrix v with the property that the product $v^T v$ is the $k \times k$ identity matrix I_k.

Definition 1.2. *The* Stiefel manifold $V_k(E)$ *of an inner product space E is the set of all k-frames in E, topologized as a subset of the product E^k.*

If $E = \mathbb{R}^n$ with it standard inner product, we will write $V_{n,k}$ in place of $V_k(\mathbb{R}^n)$.

The Stiefel manifold is a compact manifold. It comes equipped with a smooth map

$$\pi : V_{n,k} \to S^{n-1}$$

sending the k-frame $(v_0, v_1, \ldots, v_{k-1})$ to $v_0 \in S^{n-1}$. The rest of the k-frame, $(v_1, \ldots, v_{k-1})$, can be viewed as giving a $(k-1)$-frame in the tangent space to the sphere at v_0. That is, $\pi^{-1}(v_0) = V_{k-1}(T_{v_0})$.

Lemma 1.3. *The map $\pi : V_{n,k} \to S^{n-1}$ sending v to v_0 is the projection map of a fiber bundle.*

Proof. This means that any point v_0 in the base has a neighborhood W over which the projection is isomorphic to a product projection. In this case we can take for W the upper open hemisphere centered at v_0, $W = \{v_0' \in S^{n-1} : v_0 \cdot v_0' > 0\}$. Then define a map $\kappa : \pi^{-1}W \to V_{k-1}(T_{v_0})$ as follows. Let $v' \in \pi^{-1}W$. The sequence $(v_0, v_1', \ldots, v_{k-1}')$ is still linearly independent, so we can apply the Gram-Schmidt process to it to obtain a sequence $(v_0, v_1, \ldots, v_{k-1})$. Define $\kappa(v') = (v_1, \ldots, v_{k-1})$. Then

$$(\pi, \kappa) : \pi^{-1}W \to W \times V_{k-1}(T_{v_0})$$

is a homeomorphism compatible with the projections to W. □

So the projection map always has local cross-sections; the sphere (or any $(n-1)$-manifold) admits local $(k-1)$-frames as long as $k \leq n$. A global cross-section is the same thing as a tangential $(k-1)$-frame field (or briefly a $(k-1)$-frame) on S^{n-1}. Finding obstructions to the existence of a cross-section of a fiber bundle is one of a small collection of standard problems in homotopy theory, and we will bring a number of homotopy-theoretic tools to bear on this particular sectioning problem.

We may ask how large k can be, given n. The answer is given in terms of a function $\rho(n)$, defined as follows. Let $\nu(n)$ be the largest integer such that 2^ν divides n. Then write $\nu = \nu(n)$ as $\nu = 4b + c$, $0 \leq c \leq 3$, and set

$$\rho(n) = 8b + 2^c\,.$$

Theorem 1.4 (Hurwitz, Radon, Eckmann; Adams). *There exist $\rho(n) - 1$ linearly independent vector fields on S^{n-1} (Hurwitz-Radon-Eckmann) and no more (Adams).*

Here's a table of a few values of this strange function $\rho(n)$. Since $\rho(n)$ depends only on $\nu(n)$, we tabulate its values in terms of ν.

ν	0	1	2	3	4	5	6	7	$\cdots$
2^ν	1	2	4	8	16	32	64	128	$\cdots$
$\rho(n)$	1	2	4	8	9	10	12	16	$\cdots$

Thus $\rho(n)$ is roughly twice the exponent of 2 in n; it quickly falls behind n. In fact the only cases in which $\rho(n) = n$ are $n = 1, 2, 4, 8$; the only parallelizable spheres are S^0, S^1, S^3, S^7. The sphere S^{127} admits fifteen everywhere independent vector fields and no more.

There are two steps to proving Theorem 1.4:

1. Construct enough vector fields. It turns out that linear algebra – in the form of representations of Clifford algebras – produces a $(\rho(n)-1)$-frame on S^{n-1}.
2. Show that no other method can produce more. This is much harder, and was the first major victory for K-theory in topology, the occasion for Frank Adam's introduction of the operations named after him.

Let's focus for the moment on the parallelizable case. There are many statements that are equivalent to saying that S^{n-1} is parallelizable.

If S^{n-1} is a Lie group, we can pick a basis for the tangent space at the identity element and translate it around the group; this gives a different way of thinking of the parallelization we got above by thinking of S^1 and S^3 as unit vectors in $\mathbb{C}$ and $\mathbb{H}$. The 7-sphere is not a Lie group, but it does posess a smooth (though non-associative) multiplication, which can be used to parallelize it.

Conversely, a parallelization of a sphere produces a multiplication on the sphere. More precisely, you get an H-space structure, where:

Definition 1.5. *An H-space is a pointed space X equipped with a map $\mu : X \times X \to X$ such that the diagram*

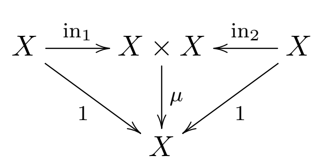

is homotopy commutative, where $i_1(x) = (x, *)$ *and* $i_2(x) = (*, x)$.

Lemma 1.6. *A parallelization of* S^{n-1} *determines an* H*-space structure on this space.*

Proof. The H-space structure arises via the "stereographic projection" $\sigma : TS^{n-1} \to S^{n-1}$, defined as follows. A tangent vector is a pair (x, v) where $x \in S^{n-1}$ and v is a vector in $\mathbb{R}^n$ perpendicular to x. Connect the sum $x+v$ to $-x$ by a line segment. It intersects S^{n-1} in two points: $-x$ is one, and the other we take as the definition of $\sigma(x, v)$. Thus $\sigma(x, 0) = x$, and for any nonzero v the limit of $\sigma(x, tv)$ as $t \to \infty$ is $-x$. This lets us extend the map σ slightly, by adjoining a point "at infinity" to each tangent space. This yields a fiber bundle $\overline{T}S^{n-1} \to S^{n-1}$ whose fiber over $x \in S^{n-1}$ is the one-point compactification of T_x; that is, an $(n-1)$-sphere. $\overline{T}S^{n-1}$ is a "twisted product" of S^{n-1} with itself. We obtain a map $\sigma : \overline{T}S^{n-1} \to S^{n-1}$ such that the left triangle commutes and the right one commutes up to homotopy.

$$\begin{array}{ccccc} S^{n-1} & \xrightarrow{s_0} & \overline{T}S^{n-1} & \xleftarrow{\iota_x} & S^{n-1} \\ & {\scriptstyle 1}\searrow & \downarrow{\scriptstyle\sigma} & \swarrow{\scriptstyle 1} & \\ & & S^{n-1} & & \end{array}$$

where s_0 is the inclusion of the zero-section and ι_x is the inclusion of the fiber over x.

Now suppose that S^{n-1} is parallelizable. A parallelization gives a trivialization of the tangent bundle, and hence of the compactified tangent bundle. This gives us a homeomorphism $S^{n-1} \times S^{n-1} \to \overline{T}S^{n-1}$ with the property that the zero section and fiber inclusion correspond to axis inclusions. So the parallelization determines an H-space structure on S^{n-1}. $\square$

Next, we can use an H-space structure to produce a corresponding "projective plane." This uses several general topological constructions. The *cone* CX on a space X is the quotient space

$$CX = (X \times I)/\sim$$

where I denotes the unit interval $[0, 1]$ and $(x, 0) \sim (x', 0)$ for any $x, x' \in X$. Note that CX has a distinguished basepoint, represented by $(x, 0)$ for any $x \in X$. (Categorical considerations demand that we say $C\varnothing = *$.) CX admits a canonical contracting homotopy. Write $i : X \to CX$ for the inclusion $x \mapsto (x, 1)$. Given a map $f : X \to Y$, the *mapping cone* $C(f)$

of f is the pushout in

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ {\scriptstyle i}\downarrow & & \downarrow \\ CX & \longrightarrow & C(f) \end{array}$$

We could also write $C(f) = Y \cup CX$. For example, the mapping cone of the unique map $X \to *$ is the *unreduced suspension* SX of X; this is also the quotient of CX by the subspace given by the image of i.

The *join* of two spaces, X and Y, is

$$X * Y = X \times I \times Y / \sim$$

where $(x, 0, y) \sim (x', 0, y)$ for any $x, x' \in X, y \in Y$, and $(x, 1, y) \sim (x, 1, y')$ for any $x \in X$ and $y, y' \in Y$.

The *Hopf construction* of a map $\mu : X \times Y \to Z$ is the map

$$h(\mu) : X * Y \to SZ$$

defined by the formula

$$h(\mu)(x, t, y) = (\mu(x, y), t)\,.$$

The *projective plane* of an H-space (X, μ) is the mapping cone of the Hopf construction of μ.

Let's apply this to our H-space structures on spheres. When $n = 1$, we have S^0; there is only one multiplication μ on this two-point space; $S^0 * S^0$ and SS^0 are both homeomorphic to S^1, and $h(\mu)$ is a map $2\iota : S^1 \to S^1$ of degree 2 (with appropriate choice of orientations). In general, the join of two spheres is another sphere:

$$S^{p-1} * S^{q-1} \xrightarrow{\cong} S^{p+q-1}$$

by sending (x, t, y) to $(\cos(\pi t/2)x, \sin(\pi t/2)y)$; and of course the unreduced suspension of S^{n-1} is an n-sphere. The Hopf constructions on H-space structures on spheres give important maps, the *Hopf maps*

$$2\iota : S^1 \to S^1 \quad , \quad \eta : S^3 \to S^2 \quad , \quad \nu : S^7 \to S^4 \quad , \quad \sigma : S^{15} \to S^8\,.$$

The discovery of the essential (non null-homotopic) map η, by Heinz Hopf, was one of the kick-off events in the discipline of homotopy theory; and for many years these maps and their combinations were the only known elements in the homotopy groups of spheres:

$$2\iota \in \pi_1(S^1) \quad , \quad \eta \in \pi_3(S^2) \quad , \quad \nu \in \pi_7(S^4) \quad , \quad \sigma \in \pi_{15}(S^8) \,.$$

These maps really are essential. We can see that by exploiting the mapping cone construction, which uses a map to produce a space. This construction is "homotopical":

Lemma 1.7. *Let $f, g : X \to Y$. A homotopy from f to g determines a homotopy equivalence $Cf \to Cg$.*

Let $*$ denote both a point $* \in Y$ and a map from any space to Y whose image lies in $\{*\}$. Then

$$C(* : X \to Y) = Y \vee SX$$

where the $\vee$ denotes the "wedge" of the two pointed spaces, obtained from the disjoint union by identifying base points. So if the mapping cone of f does not split as a wedge then f is essential.

Suppose $\alpha \in \pi_{2n-1}(S^n)$. Its mapping cone has cells in dimensions $0, n$, and $2n$. If $n > 1$, the integral cohomology is rank one free abelian in these dimensions, and orientations for the spheres provide us with generators x_n and x_{2n}. The cup product structure is determined by an integer $H(\alpha)$:

$$x_n^2 = H(\alpha) x_{2n}$$

As we shall see, this defines a group homomorphism

$$H : \pi_{2n-1}(S^n) \to \mathbb{Z} \,,$$

the *Hopf invariant.*

In our case, the mapping cones of the various Hopf maps are the projective planes $\mathbb{R}P^2, \mathbb{C}P^2, \mathbb{H}P^2, \mathbb{O}P^2$. Homology groups suffice to distinguish $\mathbb{R}P^2$ from $S^1 \vee S^2$. For the rest, the cup product structure does; in each case, the cup product structure of projective space shows that the Hopf invariant of each of the Hopf maps (for $n > 1$) is ± 1.

We have seen a good portion of the following omnibus result. The hard part, (4) $\Longrightarrow$ (5), is due to Frank Adams.

Theorem 1.8. *The following are equivalent.*
(1) *There is a nonsingular bilinear pairing $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$.*

(2) S^{n-1} *is parallelizable.*
(3) S^{n-1} *admits the structure of an* H*-space.*
(4) *The Hopf invariant* $\pi_{2n-1}(S^n) \to \mathbb{Z}$ *is surjective.*
(5) $n \in \{1, 2, 4, 8\}$.

The Hopf maps arise in many contexts. A slightly different construction of them was generalized by George Whitehead. The orthogonal group $O(n)$ acts linearly on the sphere S^{n-1}. Any map $S^{k-1} \to SO(n) \subset O(n)$ thus gives us a composite

$$S^{k-1} \times S^{n-1} \to SO(n) \times S^{n-1} \to S^{n-1}$$

to which we can apply the Hopf constuction. The result is a homomorphism

$$J : \pi_{k-1}(SO(n)) \to \pi_{n+k-1}(S^n) .$$

The orthogonal group $O(n)$ is also known as $V_{n,n}$, so a parallelization of S^{n-1} compatible with its orientation provides an example of such a homotopy class α with $k = n$, and $J\alpha \in \pi_{2n-1}(S^n)$ is the class we constructed above.

The maps J are compatible as n increases. On the right, the groups stabilize to the "stable homotopy group" π^s_{k-1}. The increasing union of the topological groups SO is the "stable special orthogonal group" SO, and we get a map

$$J : \pi_{k-1}(SO) \to \pi^s_{k-1} .$$

Luckily, the homotopy groups of SO are known; this is part of Bott periodicity. The image of the J-homomorphism was studied, also by Adams, and (following the resolution of the "Adams conjecture") turns out to be a direct summand of known order. It is the "linear part" of the stable homotopy groups of spheres, and in this course we will get to know these groups very well.

2 Clifford algebras and their representations

In this lecture we will see how many vector fields on spheres we can construct using linear algebra. We'll begin by defining a family of associative algebras called Clifford algebras, then explain how their representations can be used to produce frame fields on spheres, identify them in simpler terms that will allow us to understand their representation theory, and then show that we have produced a $(\rho(n) - 1)$-frame on S^{n-1} for any n. Clifford algebras provide a means to study quadratic forms in some generality, but we will focus on the particular form (two forms, actually) that will be of use to us in constructing vector fields on spheres.

Definition 2.1. *Let $k \geq 0$. The* Clifford algebra C_k^+ *is the associative $\mathbb{R}$-algebra generated by elements $e_1, \ldots, e_k$, subject to relations*

$$e_i e_j + e_j e_i = 0 \text{ for } i \neq j\,,$$
$$e_i^2 = -1\,.$$

For example,

$$C_0^+ = \mathbb{R},$$
$$C_1^+ \cong \mathbb{C} \text{ by } e_1 \mapsto i$$
$$C_2^+ \cong \mathbb{H} \text{ by } e_1 \mapsto i\,,\, e_2 \mapsto j\,,\, e_1 e_2 \mapsto k$$

Remarks 2.2. 1. The relations imply that a basis for C_k^+ is given by the set of words

$$\{e_{i_1} \cdots e_{i_m} : m \geq 0\,,\, i_1 < \cdots < i_m\}$$

made up of ordered nonrepeating sequences of the generators. So

$$\dim C_k^+ = 2^k\,.$$

2. The set

$$G_k = \{\pm e_{i_1} \cdots e_{i_m} : m \geq 0\,,\, i_1 < \cdots < i_m\}$$

is a multiplicative subgroup of C_k^+, of order 2^{k+1}. For example G_2 is the quaternion group of order 8. A C_k^+-module is the same thing as a real representation of G_k with the property that $(-e_i) \cdot x = -(e_i \cdot x)$ for all i and x.

Before trying to identify more of these algebras, let's look at how they can be used to construct vector fields on spheres. Briefly, a representation of C_k^+ on $\mathbb{R}^n$ produces a k-frame on S^{n-1}.

Suppose V is an n-dimensional real vector space with a C_k^+-module structure. Choose an inner product on V. By averaging over the action of the finite group G_k we arrive at a G_k-invariant inner product $(-,-)$. Let $\mathbb{S}(V)$ denote the unit sphere of V, and note that multiplication by e_i sends $\mathbb{S}(V)$ to itself. We claim:

Lemma 2.3. *The function*

$$x \mapsto (e_1 x, \ldots, e_k x)$$

defines an orthonormal k-frame on $\mathbb{S}(V)$.

Proof. The proof consists in two simple calculations. First,

$$(x, e_i x) = (e_i x, e_i e_i x) = -(e_i x, x) = -(x, e_i x)\,,$$

so $(x, e_i x) = 0$ and $e_i x$ is tangent to $\mathbb{S}(V)$ at x. Second, if $i \neq j$, we have

$$(e_i x, e_j x) = (e_i e_j e_i x, e_i e_j e_j x) = (-e_i^2 e_j x, e_i e_j^2 x) = (e_j x, -e_i x) = -(e_i x, e_j x)\,,$$

so $(e_i x, e_j x) = 0$, and $e_i x$ and $e_j x$ are orthogonal. □

This recovers the 1-frames on S^{n-1} with n even and the 3-frames on S^{n-1} with n divisible by 4 that we saw in in the last lecture.

It turns out that the algebras C_k^+ can be expressed in terms of $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, and matrix algebras over them. If A is an associative $\mathbb{R}$-algebra, write A^n for the n-fold product of A with itself, and write $A(n)$ for the associative $\mathbb{R}$-algebra of $n \times n$ matrices over A. We will identify the Clifford algebras inductively in terms of these basic examples and operations, building up the table below.

k	0	1	2	3	4	5	6	7	8
C_k^+	$\mathbb{R}$	$\mathbb{C}$	$\mathbb{H}$	$\mathbb{H}^2$	$\mathbb{H}(2)$	$\mathbb{C}(4)$	$\mathbb{R}(8)$	$\mathbb{R}(8)^2$	$\mathbb{R}(16)$

$$C_k^+ = C_{k-8}^+(16)\,, \quad k \geq 8$$

In carrying out this induction it is useful to introduce a Clifford algebra associated to a different quadratic form. Let C_k^- be the associative $\mathbb{R}$-algebra generated by $e_1, \ldots, e_k$, subject to relations $e_i e_j + e_j e_i = 0$ (for $i \neq j$) and $e_i^2 = 1$.

We need to compute the first few of the C_k^-'s as well. C_1^- has basis $\{1, e_1\}$, and $e_1^2 = 1$, so $C_1^- \cong \mathbb{R}^2$, with $e_1 \mapsto (1, -1)$ (or $(-1, 1)$). C_2^- admits a representation on $\mathbb{R}^2$: Let e_1 act by reflecting across the x-axis and e_2 by reflecting across the line $x = y$. Then $e_1 e_2$ gives a clockwise rotation by 90° and $e_2 e_1$ gives the inverse rotation. This representation gives us an algebra isomorphism $C_2^- \cong \mathbb{R}(2)$.

This simple geometry starts our induction. For the inductive step, we have:

Lemma 2.4. *For $k \geq 2$, $C_k^\pm \cong C_2^\pm \otimes C_{k-2}^\mp$.*

Proof. Take the upper signs first. Define a linear map $\mathbb{R}^k \to C_2^+ \otimes C_{k-2}^-$ by

$$e_1 \mapsto e_1 \otimes 1 \quad , \quad e_2 \mapsto e_2 \otimes 1$$

and for $i > 2$
$$e_i \mapsto e_1 e_2 \otimes e_{i-2} \,.$$
Note that in C_2^+
$$(e_1 e_2)^2 = e_1(-e_1 e_2)e_2 = -1 \,,$$
so that in $C_2^+ \otimes C_{k-2}^-$ we have
$$(e_1 \otimes 1)^2 = -1 \otimes 1 \quad , \quad (e_2 \otimes 1)^2 = -1 \otimes 1 \,,$$
and, for $i > 1$
$$(e_1 e_2 \otimes e_{i-2})^2 = -1 \otimes 1 \,.$$
The anti-commutativity is easily checked, so this linear map extends to an algebra map
$$C_k^+ \to C_2^+ \otimes C_{k-2}^- \,.$$
This map sends basis elements to basis elements and so is an isomorphism.

The other case is identical up to sign. □

Using the known values of $C_2^{\pm}$, these isomorphisms may be rewritten
$$C_k^+ \cong \mathbb{H} \otimes C_{k-2}^- \quad , \quad C_k^- \cong C_{k-2}^+(2) \tag{2.1}$$
where for the second isomorphism we use the fact that $A \otimes \mathbb{R}(n) \cong A(n)$.

Finally we need to observe:

Lemma 2.5. *We have the following* $\mathbb{R}$*-algebra isomorphisms*
$$\begin{aligned} \mathbb{C} \otimes \mathbb{C} &\cong \mathbb{C}^2 \\ \mathbb{C} \otimes \mathbb{H} &\cong \mathbb{C}(2) \\ \mathbb{H} \otimes \mathbb{H} &\cong \mathbb{R}(4) \,. \end{aligned}$$

Proof. The complex numbers admit two complex-valued real bilinear forms, xy and $x\bar{y}$. Together they provide the first isomorphism.

For the second isomorphism, observe that there is a homomorphism of $\mathbb{R}$-algebras $\mathbb{H} \to \mathbb{C}(2)$ sending
$$i \mapsto \begin{bmatrix} 0 & 1 \\ i & 0 \end{bmatrix} \quad , \quad j \mapsto \begin{bmatrix} 0 & i \\ 1 & 0 \end{bmatrix} .$$
This map extends to an isomorphism $\mathbb{C} \otimes \mathbb{H} \to \mathbb{C}(2)$.

The third isomorphism arises from the $\mathbb{H} \otimes \mathbb{H}$-module structure on $\mathbb{H}$ given by the action

$$(x \otimes y) \cdot z = xz\overline{y}$$

where $y \mapsto \overline{y}$ is the anti-involution defined by $\overline{i} = -i$, $\overline{j} = -j$. □

The remaining $C_k^{\pm}$ in the table can be computed in succession using these lemmas, rather like lacing a skate. The arrows pointing southeast tensor with $\mathbb{H}$ and those pointing southwest tensor with $\mathbb{R}(2)$:

k	C_k^+	C_k^-
0	$\mathbb{R}$	$\mathbb{R}$
1	$\mathbb{C}$	$\mathbb{R}^2$
2	$\mathbb{H}$	$\mathbb{R}(2)$
3	$\mathbb{H}^2$	$\mathbb{C}(2)$
4	$\mathbb{H}(2)$	$\mathbb{H}(2)$
5	$\mathbb{C}(4)$	$\mathbb{H}(2)^2$
6	$\mathbb{R}(8)$	$\mathbb{H}(4)$
7	$\mathbb{R}(8)^2$	$\mathbb{C}(8)$
8	$\mathbb{R}(16)$	$\mathbb{R}(16)$

If we continue to lace, we will repeat the same pattern but now tensored with $\mathbb{R}(16)$:

$$C_k^{\pm} = C_{k-8}^{\pm}(16) \ , \quad k \geq 8 \, .$$

These identifications of the algebras C_k^+ provide us with representations, using the following constructions.

(1) $\mathbb{R}$, $\mathbb{C}$, and $\mathbb{H}$ have representations on themselves by left multiplication.

(2) If $\mathbb{R}^n$ has the structure of an A-module, then $\mathbb{R}^{dn}$ has the structure of an $A(d)$-module, adapting the usual way in which a matrix acts on a column vector.

(3) If the associative $\mathbb{R}$-algebra A acts on $\mathbb{R}^n$, then A^2 acts on $\mathbb{R}^n$ as well, in two ways — through the projections to the two factors. We obtain representations of real dimension a_k, as given in the following table.

k	0	1	2	3	4	5	6	7	8
C_k^+	$\mathbb{R}$	$\mathbb{C}$	$\mathbb{H}$	$\mathbb{H}^2$	$\mathbb{H}(2)$	$\mathbb{C}(4)$	$\mathbb{R}(8)$	$\mathbb{R}(8)^2$	$\mathbb{R}(16)$
a_k	1	2	4	4	8	8	8	8	16

$$a_k = 16a_{k-8}\,, \quad k \geq 8$$

By taking direct sums of these representations we also get representations of C_k^+ on $\mathbb{R}^{ca_k}$ for any $c \geq 1$. It is not hard to see that these representations are of minimal dimension; in fact every representation of C_k^+ is a direct sum if these (remembering that there are two of dimension a_k if $k \equiv 3, 7 \bmod 8$).

We can now harvest some families of vector fields on spheres. A representation of C_k^+ on $\mathbb{R}^n$ provides us with a k-frame on S^{n-1}. We have found such representations provided that n is a multiple of a_k. To maximize k for given n, then, we want to maximize k such that a_k divides n. The dimension a_k is always a power of 2, so this maximal k will depend only on $\nu(n)$, the power of 2 dividing n. The first few cases are:

$\nu(n)$	0	1	2	3	4	5	6
k_{max}	0	1	3	7	8	9	11

and

$$k_{max}(\nu) = k_{max}(\nu - 4) + 8 \quad \text{for} \quad \nu \geq 4\,.$$

This is exactly $\rho(n) - 1$ from Lecture 1. So we have succeeded in constructing $\rho(n) - 1$ linearly independent vector fields on the $(n-1)$-sphere, and explained, from this linear algebra perspective, the occurence of the strange number $\rho(n)$. It turns out that this is the best we can do using linear algebra — in fact this really is the best we can do period. To prove this we will have to explain how the number $\rho(n)$ occurs from a purely homotopy-theoretic perspective.

3 Projective space

We now set about finding obstructions to sectioning the projection map $V_{n,k} \downarrow S^{n-1}$. The key to this is problem is hidden in the homotopy theory of real projective space, $\mathbb{R}P^{k-1}$. This space will be a central actor throughout these lectures.

By definition, $\mathbb{R}P^{k-1}$ is the quotient of S^{k-1} obtained by identifying x and $-x$. Thus it is the orbit space of the action of the cyclic group C_2

of order 2 on S^{k-1} in which the nontrivial element acts antipodally. The projection $S^{k-1} \downarrow \mathbb{R}P^{k-1}$ is a principal C_2-bundle.

Every line through the origin in $\mathbb{R}^k$ meets S^{k-1} precisely in a pair of antipodal points, so we can also regard $\mathbb{R}P^{k-1}$ as the space of lines through the origin in $\mathbb{R}^k$. From this perspective, the natural structure over $\mathbb{R}P^{k-1}$ is a real line bundle λ, the *canonical line bundle*. The total space of λ is

$$\mathbb{E}(\lambda) = \{(x, l) \in \mathbb{R}P^{k-1} \times \mathbb{R}^k : x \in l\} .$$

These two structures are equivalent to each other. The line bundle λ has a natural metric, given by restricting the Euclidean inner product in $\mathbb{R}^n$, and the double covering by the sphere is the unit sphere-bundle in λ (with fiber sphere S^0):

$$\begin{array}{ccc} S^{k-1} & \xrightarrow{\cong} & \mathbb{S}(\lambda) \\ & \searrow \quad \swarrow & \\ & \mathbb{R}P^{k-1} & \end{array}$$

In fact the sphere with its C_2 action may be thought of as the "principalization" of λ: an element of S^n can be thought of as an isometric linear isomorphism from $\mathbb{R}$ to a fiber of λ. Conversely, λ is obtained from the principal C_2-bundle $S^{k-1} \downarrow \mathbb{R}P^{k-1}$ via the "Borel construction": C_2 acts on $\mathbb{R}$ linearly by sending x to $-x$; this is the "sign representation" of C_2. The Borel construction is given by

$$S^{k-1} \times \mathbb{R}/\sim \quad , \, (x, t) \sim (-x, -t) .$$

The map sending (x, t) to $\pm x \in \mathbb{R}^{k-1}$ factors through the quotient and displays this space as the total space of a line bundle over $\mathbb{R}P^{k-1}$, one which is none other than the canonical line bundle λ.

Another (less interesting but just as important) line bundle over $\mathbb{R}P^{k-1}$ is the trivial line bundle $\mathbb{R}P^{k-1} \times \mathbb{R} \downarrow \mathbb{R}P^{k-1}$, which we denote by ϵ. The canonical line bundle comes to us embedded as a sub-bundle of the trivial k-plane bundle $k\epsilon$ over $\mathbb{R}P^{k-1}$, and as such has a complementary $(k-1)$-plane bundle $\lambda^\perp$.

The covering projection $S^{k-1} \downarrow \mathbb{R}P^{k-1}$ shows that $\mathbb{R}P^{k-1}$ is locally isomorphic to the manifold S^{k-1}, and so is itself a compact smooth manifold. We will identify its tangent bundle.

A tangent vector $v \in (\mathbb{R}x)^\perp$ at $x \in S^{k-1}$ descends to a tangent vector at $\mathbb{R}x \in \mathbb{R}P^{k-1}$. The opposite vector $-v$ at $-x$ descends to the same tangent vector. This behavior can be expressed by saying that a tangent vector at $l \in \mathbb{R}P^{k-1}$ is the same thing as a linear transformation from l to $l^{\perp}$. The tangent space is the vector bundle whose fibers are these spaces of homomorphisms,

$$T\mathbb{R}P^{k-1} = \mathrm{Hom}(\lambda, \lambda^{\perp}) \,.$$

The connection of all this to our sectioning problem is the following simple lemma.

Lemma 3.1. *A section s of the projection $V_{n,k} \downarrow S^{n-1}$ determines a bundle map*

$$\begin{array}{ccc} \mathbb{R}P^{k-1} \times S^{n-1} = \mathbb{S}(n\epsilon) & \xrightarrow{\hat{s}} & \mathbb{S}(n\lambda) = S^{k-1} \times_{C_2} S^{n-1} \\ & \searrow \qquad \swarrow & \\ & \mathbb{R}P^{k-1} & \end{array}$$

that is of degree 1 on each fiber.

Proof. To begin with, recall that the image of $x \in S^{n-1}$ under the section s defines a linear isometric embedding $s(x) : \mathbb{R}^k \to \mathbb{R}^n$. So we can define a map $\overline{s} : \mathbb{R}^k \times S^{n-1} \to \mathbb{R}^k \times \mathbb{R}^n$ by

$$\overline{s}(x, v) = (x, s(v)x)$$

that restricts to a map $S^{k-1} \times S^{n-1} \to S^{k-1} \times S^{n-1}$. Since $\overline{s}(-x, v) = (-x, -s(v)x)$, this map descends to a map $\hat{s}$ between the indicated quotients. Both projection maps are induced by projection to the first factor, so $\hat{s}$ is a map of fiber bundles over $\mathbb{R}P^{k-1}$. The map s is a section of the map that sends a linear isometric embedding to its value on $e_1 \in \mathbb{R}^k$: $s(v)e_1 = v$. So $\overline{s}(e_1, v) = (e_1, v)$, inducing a map of degree 1 on the fiber over $\pm e_1 \in \mathbb{R}P^{k-1}$. The degree of the restriction to the fiber over $\pm x$ is an integer depending continuously on $\pm x \in \mathbb{R}P^{k-1}$ and hence is constant since $\mathbb{R}P^{k-1}$ is connected. $\square$

When the section s is one of those constructed in Lecture 2, the map $\overline{s}$ defines a map of vector bundles $n\epsilon \to n\lambda$ that is an isomorphism on fibers, and so provides a trivialization of the bundle $n\lambda$ over $\mathbb{R}P^{k-1}$. In general, the map $\hat{s}$ does not extend to a linear isomorphism, so we cannot conclude that $n\lambda$ is trivial as a vector bundle.

Nevertheless, the map $\hat{s}$ does provide us with a trivialization of the spherical fibration $\mathbb{S}(n\lambda) \downarrow \mathbb{R}P^{k-1}$ is trivial. We had better be clear about what this means.

Definition 3.2. *Let E and E' be fiber bundles over a common base B. Two bundle maps $f, g : E \to E'$ are* fiber homotopic *if they are homotopic through bundle maps. The bundle map f is a* fiber homotopy equivalence *if there is a bundle map $g : E' \to E$ such that fg and gf are fiber homotopic to the respective identity maps. A* fiber homotopy trivialization *is a fiber homotopy equivalence from a trivial bundle.*

We have a bundle map of sphere bundles that is a homotopy equivalence on each fiber. Fortunately this forces it to be a fiber homotopy equivalence by the following theorem of Albrecht Dold.

Theorem 3.3 (Dold)**.** *Suppose E and E' are fibrations over B with a bundle map f inducing a homotopy equivalence on each fiber. If E and E' have the homotopy type of CW complexes then f is a fiber homotopy equivalence.*

Proof. See [3]. □

If B is connected, then it is enough to check that the map is a homotopy equivalence on a single fiber. This is clearest if the fiber is a sphere; then the degree of the map is a continuous function with discrete values. The homotopy inverse of a fiber homotopy trivialization of a fiber bundle with fiber F amounts to a map from the total space to F which restricts to a homotopy equivalence on each fiber.

In our context, the lemma gives

Corollary 3.4. *A section of $V_{n,k} \downarrow S^{n-1}$ determines a fiber homotopy trivialization of $\mathbb{S}(n\lambda) \downarrow \mathbb{R}P^{k-1}$.*

So the topology of certain vector bundles over real projective space become important in our study of the vector field problem. Let's get a little more familiar with these bundles.

First of all, a linear embedding of $\mathbb{R}^k$ into $\mathbb{R}^{n+k}$ induces an embedding of manifolds

$$\mathbb{R}P^{k-1} \hookrightarrow \mathbb{R}P^{n+k-1}.$$

For definiteness, let's take $\mathbb{R}^k$ to be cut out by setting the last n coordinates equal to zero. Its orthogonal complement is an appropriately embedded copy of $\mathbb{R}^n$, which defines a complementary embedding of $\mathbb{R}P^{n-1}$ into $\mathbb{R}P^{n+k-1}$.

An element $l \in \mathbb{R}P^{k-1}$ together with a linear map $f : l \to V$ determines an element of the larger projective space $\mathbb{R}P^{n+k-1}$, namely, the graph of f, as a subspace of $l \times \mathbb{R}^n \subseteq \mathbb{R}^{n+k}$. If $f = 0$, the graph is just l again. On the other hand, if l' is almost any line through the origin in $\mathbb{R}^{n+k}$, we may project it to a line in $\mathbb{R}^k$, and express it as the graph of a linear function on the image line. The exception is if $l' \subseteq \mathbb{R}^n$, in which case it projects to the origin of $\mathbb{R}^k$ rather than a line.

The complement $\mathbb{R}P^{n+k-1} \backslash \mathbb{R}P^{n-1}$ is an open set containing $\mathbb{R}P^{k-1}$, which we have just identified with the total space of the vector bundle

$$\mathrm{Hom}(\lambda, n\epsilon) = n\lambda^*$$

This is a particular case of the following general theorem from differential topology:

Theorem 3.5. (Tubular neighborhood theorem) *Let $M \subset N$ be an embedded submanifold with normal bundle ν. There is a neighborhood U of M in N such that (U, M) is homeomorphic rel M to $(\mathbb{E}(\nu), M)$.*

The metric on λ further identifies λ^* with λ, so we have identified the normal bundle of $\mathbb{R}P^{k-1} \subseteq \mathbb{R}P^{n+k-1}$ with $\mathbb{E}(n\lambda)$. Since this neighborhood is dense in $\mathbb{R}P^{n+k-1}$, the quotient is precisely the one-point compactification of $\mathbb{E}(n\lambda)$.

$$\mathbb{R}P^{n+k-1}/\mathbb{R}P^{n-1} \cong \mathbb{E}(n\lambda)^+ .$$

This quotient is called a *stunted projective space*, and written

$$\mathbb{R}P_n^{n+k-1} = \mathbb{R}P^{n+k-1}/\mathbb{R}P^{n-1} .$$

Exercise 3.6. *Identify the stunted projective space $\mathbb{R}P_n^{n+1}$.*

Real projective space $\mathbb{R}P^n$ has a well-known cell structure, with k skeleton given by $\mathbb{R}P^k$. The attaching map $S^{k-1} \to \mathbb{R}P^{k-1}$ for the k-cell is simply the double cover $S^{k-1} \to \mathbb{R}P^{k-1}$. To see this, think of the k-disk embedded in $\mathbb{R}^{k+1}$ as the upper hemisphere. Every line through the origin passes through a unique point of this disk, unless it is actually a line in $\mathbb{R}^k$, when it passes through a pair of antipodal points of the boundary of this disk.

The stunted projective space $\mathbb{R}P_n^{n+k-1}$ has a cell structure obtained from the given one on $\mathbb{R}P^{n+k-1}$ by collapsing the $(n-1)$-skeleton. The attaching maps are given by composing the double cover with the projection map. $\mathbb{R}P_n^{n+k-1}$ has one cell of dimension i for $i = 0$ and $n \leq i < n + k$.

4 Thom spaces

The projection map of a vector bundle is a homotopy equivalence, with homotopy inverse given by the zero section. This would seem to make it hard to detect features of the vector bundle in homotopy theoretic terms. René Thom explained how to capture deep information about a vector bundle in a space, the *Thom space* of the bundle.

So suppose that ξ is a vector bundle with projection map $\pi : E \downarrow B$. We might as well give it a metric, so the total space $\mathbb{E}(\xi)$ contains the unit disk bundle $\mathbb{D}(\xi)$ and the unit sphere bundle $\mathbb{S}(\xi)$.

Definition 4.1. *The* Thom space *of the vector bundle ξ is*

$$Th(\xi) = \mathbb{D}(\xi)/\mathbb{S}(\xi)\,.$$

This construction may be described in terms of the fiberwise one-point compactification $\overline{\mathbb{E}}(\xi)$ defined in Lecture 1: The points at infinity define a section s_∞ of $\overline{\mathbb{E}}(\xi) \downarrow B$, and

$$Th(\xi) = \overline{\mathbb{E}}(\xi)/s_\infty(B)\,.$$

If B is a compact Hausdorff space, this is simply the one-point compactification of $\mathbb{E}(\xi)$.

For trivial bundle $n\epsilon$ over B,

$$Th(n\epsilon) = \frac{B \times D^n}{B \times S^{n-1}} = \frac{B_+ \times S^n}{+ \times S^{n-1} \cup B_+ \times *} = B_+ \wedge S^n\,.$$

In particular, the case $n = 0$ expresses B^0 as B with a disjoint basepoint adjoined.

For another example, we have shown in Lecture 3 that

$$Th(n\lambda \downarrow \mathbb{R}P^{k-1}) = \mathbb{R}P_n^{n+k-1}\,. \tag{4.1}$$

As we have defined it, the Thom space is a pointed space functorially attached to a vector bundle (perhaps with metric). But the construction actually only depends upon the underlying sphere bundle. This is because the total space of the disk bundle is none other than the mapping cylinder of the projection $\pi : \mathbb{S}(\xi) \downarrow B$. So the Thom space is simply the mapping cone of this projection map.

Consequently, fiber homotopy equivalent sphere bundles have homotopy equivalent Thom spaces. In particular, the Thom space of a fiber homotopically trival sphere bundle (with fiber dimension $n-1$) is simply an n-fold suspension:

$$Th(\xi) \simeq \Sigma^n B_+ \,.$$

For any point $* \in B$ we have a "fiber inclusion" $\pi^{-1}* \hookrightarrow \mathbb{S}(\xi)$, and an induced pointed map on Thom spaces. The Thom space of a bundle over a point is a sphere, and if we are given an orientation for the vector bundle we get an element of $\pi_n(B^\xi)$ that is independent of the chosen point if B is path-connected. A "coreduction" of a vector bundle ξ is a map $Th(\xi) \to S^n$ splitting this element up to homotopy. If the bundle is fiber homotopically trivial, we certainly get a coreduction, by collapsing B to a point:

$$Th(\xi) \simeq \Sigma^n B_+ \to S^n \,.$$

Putting this together, we have shown:

Proposition 4.2. *If S^{n-1} admits $k-1$ everywhere independent vector fields, then $n\lambda \downarrow \mathbb{R}P^{k-1}$ admits a coreduction.*

So we will spend a lot of time finding obstructions to coreducing these bundles. In terms of the cell structure given in Lecture 3, you can say that a coreduction shows that the bottom cell splits off as a wedge summand, so that

$$\mathbb{R}P_n^{n+k-1} \simeq \mathbb{R}P_{n+1}^{n+k-1} \vee S^n \,.$$

In Lecture 6 we will see what Steenrod operations have to say about this question. They do not suffice, and later we will follow Adams in applying K-theory to the problem.

Exercise 4.3. *What does your determination of the homotopy type of $\mathbb{R}P_n^{n+1}$ tell us about the non-existence of vector fields on spheres?*

A coreduction is weaker than a fiber homotopy trivialization, but not by much.

Lemma 4.4. *A coreduction of a vector bundle ξ determines a fiber homotopy trivialization of $\mathbb{S}(\xi \oplus \epsilon)$.*

Proof. We begin by giving yet another description of the Thom space of a vector bundle ξ over B. Consider the Whitney sum $\xi \oplus \epsilon$, and its unit sphere bundle $\mathbb{S}(\xi \oplus \epsilon)$. It is equipped with a nowhere vanishing section s_1 taking the value $(0,1)$ at each point. Stereographic projection from $(0,1)$ provides a map $\mathbb{E}(\xi) \to \mathbb{S}(\xi \oplus \epsilon)$, and as a vector in $E(\xi)$ goes to infinity its image goes to $(0,1)$. We thus obtain a homeomorphism

$$\overline{\mathbb{E}}(\xi) \to \mathbb{S}(\xi \oplus \epsilon)$$

under which the section at infinity s_∞ corresponds to the section s_1. So the Thom space, obtained by collapsing the image of s_∞, can be identified with the space obtained from $\mathbb{S}(\xi \oplus \epsilon)$ by collapsing the image of s_1 to a point. Under this identification, the fiber inclusion into $\mathbb{S}(\xi \oplus \epsilon)$ over a point in the base corresponds to the fiber inclusion into the Thom space.

Suppose that the bundle ξ admits a coreduction. The composite

$$\mathbb{S}(\xi \oplus \epsilon) \to Th(\xi) \to S^n$$

is a homotopy equivalence when restricted to each fiber. It therefore defines a fiber homotopy trivialization of $\mathbb{S}(\xi \oplus \epsilon)$. □

The Thom space construction interacts well with products. If $\xi \downarrow X$ and $\eta \downarrow Y$ are vector bundles, then the product of total spaces maps to the product of the base spaces and defines an "product" bundle $\xi \times \eta \downarrow X \times Y$. If $X = Y$, we can pull this back under the diagonal map $\Delta : X \to X \times Y$, and obtain the Whitney sum $\xi \oplus \eta \downarrow X$.

If we equip ξ and η with metrics, the we can take for the disk bundle of $\xi \times \eta$ the product of the disk bundles of the factors. Since the boundary sphere of a product of disks decomposes as

$$\partial(D^p \times D^q) = D^p \times S^{q-1} \cup S^{p-1} \times D^q \,,$$

we can form the Thom space of the product bundle as

$$Th(\xi \times \eta) = \frac{\mathbb{D}(\xi) \times \mathbb{D}(\eta)}{\mathbb{D}(\xi) \times \mathbb{S}(\eta) \cup \mathbb{S}(\xi) \times \mathbb{D}(\eta)} = \frac{Th(\xi) \times Th(\eta)}{Th(\xi) \vee Th(\eta)} \,,$$

so there is a natural identification

$$Th(\xi \times \eta) = Th(\xi) \wedge Th(\eta) \,. \tag{4.2}$$

By naturality, this is compatible with the inclusions of the fibers.

For example, suppose that $\eta = q\epsilon$, a trivial bundle over $*$. Since $Th(q\epsilon \downarrow *) = S^q$, (4.2) gives us

$$Th(\xi \oplus q\epsilon) = Th(\xi) \wedge S^q = \Sigma^q Th(\xi) \,. \tag{4.3}$$

The product of the 0-dimensional bundle with xi gives us a vector bundle over $X \times X$. It can be described as the pullback of ξ under $pr_2 : X \times X \to X$. The diagonal map $\Delta : X \to X \times X$ is covered by a bundle map $0 \times \xi \to \xi$, and, using the fact that $Th(0 \downarrow X) = X^+$, we receive a map

$$Th(\xi) \to X^+ \wedge Th(\xi)\,.$$

This map enters into an important theorem, due to René Thom, about the cohomology of a Thom space. In cohomology (with coefficients in a commutative ring R) the map provides us with a pairing

$$H^*(X) \otimes \overline{H}^*(Th(\xi)) \to \overline{H}^*(Th(\xi))$$

which, as can be easily checked, renders $\overline{H}^*(Th(\xi))$ a module over the graded ring $H^*(X)$.

To study this structure in more detail, suppose that the fiber dimension of ξ is n. Pick a metric. Write ξ_x for the fiber over x, and $\mathbb{D}(\xi_x)$ and $\mathbb{S}(\xi_x)$ for the fibers of the disk and sphere bundles. The relative cohomology groups

$$H^*(\mathbb{D}(\xi_x), \mathbb{S}(\xi_x); R)$$

with coefficients in a commutative ring R are nonzero only in dimension n, where they are free of rank 1 over R. As x varies, these abelian groups form a local coefficient system over X, which we denote by $\mathcal{O}_\xi$. It is a local coefficent system of R-modules, and this provides the cohomology $H^*(X; \mathcal{O}_\xi)$ with the structure of a module over the graded ring $H^*(X; R)$.

Theorem 4.5 (General Thom isomorphism theorem)**.** *Let ξ be an n-plane bundle over X. There is a natural isomorphism*

$$\overline{H}^*(Th(\xi)) \cong \sigma^n H^*(X; \mathcal{O}_\xi)$$

of modules over $H^*(X)$, where σ^n shifts the module up by n dimensions.

Proof. Once you have the Serre spectral sequence for a pair of fibrations, this is easy:

$$E_2^{s,t} = \begin{cases} H^s(X; \mathcal{O}_\xi) & \text{for} \quad t = n \\ 0 & \text{otherwise} \end{cases}$$

There is only one nonzero row, and the result follows by convergence and the multiplicative structure of the spectral sequence. □

The vector bundle is "R-oriented" by a section of $\mathcal{O}_\xi$ that provides a generator for each fiber. If $R = \mathbb{F}_2$, this is no data; the nonzero elements in the fiber give the unique R-orientation in that case. If $R = \mathbb{Z}$, this is the same as the standard notion of an orientation for a vector bundle. An orientation provides an isomorphism $\mathcal{O}_\xi$ with the trivial coefficient system with fiber R. The conclusion:

Theorem 4.6 (Oriented Thom isomorphism theorem)**.** *An R-orientation of an n-plane bundle ξ determines an element*

$$U \in \overline{H}^n(Th(\xi)) ,$$

the Thom class*, that restricts to the chosen generator of $H^n(\mathbb{D}(\xi_x), \mathbb{S}(\xi_x))$ for every $x \in X$. It serves as a generator of $\overline{H}^*(Th(\xi))$ as a free $H^*(X)$-module of rank 1.*

For example, with $R = \mathbb{F}_2$, $X = \mathbb{R}P^{k-1}$, $\xi = n\lambda$, we find that the Thom class is the generator of $\overline{H}^n(\mathbb{R}P_n^{n+k-1})$, and $\overline{H}^*(\mathbb{R}P_n^{n+k-1})$ becomes a free module of rank 1 over $H^*(\mathbb{R}P^{k-1})$.

The behavior of Thom spaces under products, (4.2), implies that the product of two R-oriented vector bundles is naturally R-oriented, and the Thom class is the product of the Thom classes':

$$U_{\xi\oplus\eta} = U_\xi \wedge U_\eta . \tag{4.4}$$

If we restrict the Thom class to $D(\xi) \simeq X$, we receive an element

$$e(\xi) \in H^n(X) .$$

the *Euler class*. This is the first and fundamental example of a "characteristic class." It is "characteristic" in the sense that it is naturally associated to the R-oriented vector bundle ξ: If $f : Y \to X$ is a map, the pulled back vector bundle $f^*\xi$ receives an R orientation in the form of $f^*U \in \overline{H}^n(\mathbb{D}(f^*\xi), \mathbb{S}(f^*\xi); R)$, and so

$$e(f^*\xi) = f^*e(\xi) .$$

We will have more to say about characteristic classes in a later lecture, but for now note just two things:

Lemma 4.7. *If a vector bundle ξ admits a nowhere vanishing section, then its Euler class is zero.*

Proof. The section can be normalized to a section of the sphere bundle; so the inclusion $\mathbb{S}(\xi) \hookrightarrow \mathbb{D}(\xi)$ induces a monomorophism in cohomology, and the previous map in the long exact sequence, $H^*(\mathbb{D}(\xi), \mathbb{S}(\xi)) \to H^*(\mathbb{D}(\xi)) = H^*(B)$ is trivial. Thus the Thom class pulls back to the zero class. □

Lemma 4.8. *$e(\xi \oplus \eta) = e(\xi)e(\eta)$.*

Proof. This follows from (4.4). □

5 Steenrod operations

Our next objective is to construct Steenrod operations and see what they tell us about our coreduction problem. Luckily, this is a 2-primary problem, so we can get away with constructing these operations in mod 2 cohomology.

Steenrod operations arise from the permutation action of the symmetric group on a product of a space with itself. The Künneth map is "commutative," in the sense that the diagram

$$\begin{array}{ccc} H^*(X) \otimes H^*(Y) & \longrightarrow & H^*(X \times Y) \\ \downarrow T & & \downarrow T^* \\ H^*(Y) \otimes H^*(X) & \longrightarrow & H^*(Y \times X) \end{array}$$

is commutative. Here we take coefficients in a commutative ring; and T denotes the signed switch map on the left and the switch map on spaces on the right. But the product arises from the Alexander-Whitney map, which is not commutative; you evaluate the first cocycle on the front face and the second on the back face. It is this failure of commutativity on the "chain level" that gives rise to Steenrod operations.

We will detect this failure of commutatity homotopy theoretically, by analysing the "extended power" construction using the Serre spectral sequence.

So to begin with, we will fix a subgroup π of the symmetric group Σ_n, which we regard as the group of permutations of the set $\{1, 2, \ldots, n\}$. There is a natural action of Σ_n on X^n, in which the action of σ given by sending $(x_1, \ldots, x_n)$ to $(x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(n)})$. This restricts to an action of π.

We will follow a standard topological gambit by studying this group action by means of its Borel construction. So let $E\pi$ be a contractible space with a free right action of π. For example, if $\pi = \Sigma_2$ we can take $E\pi = S^\infty$, the union of the unit spheres in Euclidean space, with the antipodal action. The Borel construction of an action of π on Y is given by first freeing up the action without changing the homotopy type, by looking at the diagonal action of π on $E\pi \times Y$, and then forming the orbit space with respect to this action: $E\pi \times_\pi Y$. This is a fiber bundle over $E\pi/\pi = B\pi$ with fiber Y. The "extended power" is the case $Y = X^n$ with π acting by permuting the factors.

We begin with a general observation about the relationship between the cohomology of the total space of a fibration $E \downarrow B$ and the cohomology of the fiber F over a point $* \in B$. Let $\sigma : I \to B$ be a loop at the basepoint $*$. Using the fibration condition, we can complete the following path lifting diagram.

$$\begin{array}{ccc} F & \xrightarrow{\;i\;} & E \\ {\scriptstyle in_0}\downarrow & \overset{h}{\nearrow} & \downarrow \\ F \times I & \xrightarrow{\sigma\circ pr_1} & B\,. \end{array}$$

Then the other end of the homotopy h gives us a self-map of F:

$$\begin{array}{ccc} F & \xrightarrow{\;\sigma_\#\;} & F \\ \downarrow{\scriptstyle in_1} & & \downarrow{\scriptstyle i} \\ F \times I & \xrightarrow{\;h\;} & E\,. \end{array}$$

This construction gives us an action of $\pi_1(B,*)$ on the homotopy type of the fiber F, and hence on the cohomology of F. But the map h gives us something more: it provides us with a homotopy from $i : F \to E$ to $i \circ \sigma_\# : F \to E$. It follows from this that the map induced in cohomology by i lands in the invariants under the action of $\pi_1(B,*)$:

$$i^* : H^*(E) \to H^*(F)^{\pi_1(B,*)}\,. \tag{5.1}$$

The extended power $E\pi \times_\pi X^n$ is our primary object of interest. But in order to expose its most interesting features, it is useful to work modulo a certain subspace that is defined by means of a choice of basepoint $* \in X$. Using the usual construction of the product of pairs,

$$(X, A) \times (Y, B) = (X \times Y, X \times B \cup A \times Y)\,,$$

we have the pair

$$(X, *)^n = (X^n, W_n X)$$

where W_nX is the "fat wedge"

$$W_nX = (X,*)^n = \{(x_1,\ldots,x_n) \in X^n : \text{at least of the } x_i\text{'s is } *\}\,.$$

Collapsing this subspace to a point gives the n-fold smash product $X^{\wedge n}$.

The Serre spectral sequence can be thought of as a device for determining how close the map (5.1) is to being an isomorphism. In the case of an extended power, we find the following result. Let's take coefficients in a field K and suppose that X is a CW complex such that $H_*(X;K)$ is of finite type. Then the relative Künneth map gives an equivariant isomorphism

$$\overline{H}^*(X) \to H^*(X^n, W_nX)\,. \tag{5.2}$$

Proposition 5.1. *Suppose in addition that $\overline{H}_i(X) = 0$ for $i < q$. Then*

$$H^i(E\pi \times_\pi (X^n, W_nX)) = 0 \quad \textit{for} \quad i < nq$$

and the natural map

$$H^{nq}(E\pi \times_\pi (X^n, W_nX)) \to (\overline{H}^q(X)^{\otimes n})^\pi$$

is an isomorphism.

Proof. We will use the Serre spectral sequence for the pair of fibrations $E\pi \times_\pi (X^n, W_nX^n) \downarrow B\pi$. In this spectral sequence,

$$E_2^{s,t} = H^s(B\pi; H^t(X^n, W_nX))\,,$$

where the local coefficient system is the one associated to the action of $\pi = \pi_1(B\pi)$ on the cohomology of the fiber pair. The Künneth map (5.2) shows that

$$E_2^{s,t} = 0 \quad \text{for} \quad t < nq\,.$$

It follows that

$$E_\infty^{0,nq} = E_2^{0,nq} = H^0(B\pi; H^q(X^n, W_nX)) = (\overline{H}^q(X)^{\otimes n})^\pi\,.$$

This is the only possibly nonzero group in total degreen nq, so convergence of the spectral sequence, along with the standard identification of the edge homomorphism, then gives the result. □

If we collapse the fat wedge to a point in the product we get the n-fold smash product,

$$X^n/W_nX = X^{\wedge n}\,.$$

So we define the *reduced extended π-power* to be

$$D_\pi X = \frac{E\pi \times_\pi X^n}{E\pi \times_\pi W_nX} = E\pi_+ \wedge_\pi X^{\wedge n}\,. \tag{5.3}$$

It comes equipped with a "fiber inclusion"

$$i : X^{\wedge n} \to D_\pi X \tag{5.4}$$

defined by sending x to $[e, x]$ where e is a chosen point of $E\pi$. Up to homotopy this map is independent of the choice of e.

In this lecture we will focus on the case in which $n = 2$, π is the group of order 2, and coefficients are in the field $\mathbb{F}_2$. Then $W_2X = X \vee X$, embedded in X^2 as the axes. We will write D_2X for $D_\pi X$ in this case.

Theorem 5.2. *Work in the category of pointed CW complexes of such that $\overline{H}_*(X)$ is of finite type. There is a unique natural transformation*

$$P : \overline{H}^q(X) \to \overline{H}^{2q}(D_2X)$$

such that

$$i^*Px = x \wedge x \in \overline{H}^{2q}(X^{\wedge 2})\,.$$

Proof. On pointed CW complexes, $\overline{H}^q(-)$ is representable. The universal q-dimensional cohomology class is the fundamental class

$$\iota_q \in \overline{H}^q(K_q)\,,$$

where K_q denotes the Eilenberg Mac Lane space $K(\mathbb{F}_2, q)$. By the Hurewicz theorem, $\overline{H}^i(K_q) = 0$ for $i < q$, and work of Serre implies that $\overline{H}^*(K_q)$ is of finite type. So by Proposition 5.1, the map

$$\overline{H}^{2q}(D_2K_q) \to (\overline{H}^q(K_q)^{\otimes 2})^\pi$$

is an isomorphism. The class $\iota_q \wedge \iota_q$ is invariant, so there is a unique class $P\iota_q \in \overline{H}^{2q}(D_2K_q)$ restricting to it. The theorem follows from this universal case. □

This operation is the *external square*. It yields internal operations in mod 2 cohomology by pulling back along a map induced by the diagonal map

$$\Delta : X \to X \times X .$$

This map is equivariant when π is made to act trivially on X, and hence induces a map

$$j : B\pi_+ \wedge X = E\pi_+ \wedge_\pi X \to E\pi_+ \wedge_\pi X^{\wedge 2} = D_2 X . \tag{5.5}$$

Recall that

$$H^*(B\pi) = \mathbb{F}_2[t] , \quad |t| = 1 .$$

If we pull Px back under the map j, we get a class which can be expressed using the Künneth formula as a polynomial with coefficients in $\overline{H}^*(X)$:

$$j^* Px = 1 \wedge x_{2q} + t \wedge x_{2q-1} + \cdots + t^{2q} \wedge x_0 , \quad x_i \in \overline{H}^i(X) .$$

The classes x_i depend naturally on x; $x \mapsto x_i$ is a natural transformation from $\overline{H}^q$ to $\overline{H}^i$.

We can say some things about these operations right away. Since K_q is $(q-1)$-connected, there are no cohomology operations that lower degree. This implies that $x_i = 0$ for $i < q$. For the others, let us write

$$\mathrm{Sq}^i x = x_{q+i} , \quad i \geq 0 ,$$

so that

$$j^* Px = 1 \wedge \mathrm{Sq}^q x + t \wedge \mathrm{Sq}^{q-1} x + \cdots + t^q \wedge \mathrm{Sq}^0 x$$

We have defined operations, the *Steenrod squares*

$$\mathrm{Sq}^i : \overline{H}^q \to \overline{H}^{q+i} .$$

(We will not indicate the source dimension q in the notation.) It will do no harm to agree that

$$\mathrm{Sq}^i = 0 \quad \text{on} \quad \overline{H}^q \quad \text{for} \quad i > q .$$

We have defined these operations on the reduced cohomology of pointed spaces. By applying them to the pointed space X^+, X with a disjoint basepoint adjoined, we get natural transformations

$$\mathrm{Sq}^i : H^q(X) \to H^{q+i}(X)$$

natural on unpointed spaces.

One of these operations is easy to identify. A point in $E\pi$ determines vertical maps fitting into a commutative diagram

$$\begin{CD} B\pi_+ \wedge X @>j>> D_2X \\ @AAA @AAiA \\ X @>\Delta>> X^{\wedge 2} \end{CD}\ .$$

Up to homotopy, these vertical maps are independent of choice of point in $E\pi$. Chasing the element Px, for $x \in \overline{H}^q(X)$, shows that

$$\mathrm{Sq}^q x = x^2 \qquad |x| = q\,.$$

This operation is linear, since $(x+y)^2 = x^2 + y^2$ mod 2, and multiplicative, since $(xy)^2 = x^2y^2$. How about the others? In the next lecture, we will focus first on the effect of the squares on cup-products. Additivity will turn out to be a consequence.

6 The Cartan formula

We will study the effect of the total power operation on products using the diagonal map

$$D_2(X\wedge Y) = E\pi_+\wedge_\pi (X\wedge Y)^{\wedge 2} \xrightarrow{\delta} (E\pi_+\wedge_\pi X^{\wedge 2})\wedge(E\pi_+\wedge_\pi Y^{\wedge 2}) = D_2X\wedge D_2Y$$

$$(e,(x_1,y_1),(x_2,y_2)) \mapsto (e,(x_1,x_2)),(e,(y_1,y_2))$$

Note that the following diagram commutes.

$$\begin{CD} (X\wedge Y)^{\wedge 2} @>i>> D_2(X\wedge Y) @<j<< B\pi_+\wedge(X\wedge Y) \\ @VVTV @VV\delta V @VVT\Delta_+V \\ X^{\wedge 2}\wedge Y^{\wedge 2} @>i\wedge i>> D_2X\wedge D_2Y @<j\wedge j<< B\pi_+\wedge X\wedge B\pi_+\wedge Y \end{CD} \tag{6.1}$$

Lemma 6.1. *$\delta^*(Pu\wedge Pv) = P(u\wedge v)$.*

Proof. It suffices to take the universal case,

$$X = K_p\,, \quad Y = K_q\,, \quad u = \iota_p \in \overline{H}^p(K_p)\,, \quad v = \iota_q \in \overline{H}^q(K_q)\,.$$

Apply cohomology to the left square in (6.1). The element $P\iota_p \wedge P\iota_q \in \overline{H}^{q(p+q)}(D_2K_p \wedge D_2K_q)$ gets sent by $T^*(i \wedge i)^*$ to $(\iota_p \wedge \iota_q)^{\wedge 2} \in \overline{H}^{2(p+q)}((X \wedge Y)^{\wedge 2})$. So $i^*\delta^*(P\iota_p \wedge P\iota_q = (\iota_p \wedge \iota_q)^{\wedge 2}$ as well. Also, $i^*(P(\iota_p \wedge \iota_q) = (\iota_p \wedge \iota_q)^{\wedge 2}$. But $\overline{H}_i(K_p \wedge K_q) = 0$ for $i < p+q$, so by Proposition 5.1 the map i^* is injective in this dimension. Thus

$$\delta^*(P\iota_p \wedge P\iota_q) = P(\iota_p \wedge \iota_q)\,,$$

which is the universal case of our formula. □

Corollary 6.2. (External Cartan formula) *Let* $x \in \overline{H}^p(X)$ *and* $y \in \overline{H}^q(Y)$. *In* $\overline{H}^*(X \wedge Y)$,

$$\mathrm{Sq}^k(x \wedge y) = \sum_{i+j=k} \mathrm{Sq}^i x \wedge \mathrm{Sq}^j y\,.$$

Proof. Chase $Px \wedge Py \in \overline{H}^{2(p+q)}(D_2X \wedge D_2Y)$ in cohomology applied to the right square of (6.1) and use Lemma 6.1:

$$\begin{array}{c} (T\Delta_+)^*(j^*Px \wedge j^*Py) = \sum_{i,j} t^{(p+q)-(i+j)} \wedge \mathrm{Sq}^i x \wedge \mathrm{Sq}^j y \\ || \\ j^*\delta^*(Px \wedge Py) = j^*P(x \wedge y) = \sum_k t^{(p+q)-k} \wedge \mathrm{Sq}^k(x \wedge y)\,. \end{array}$$

Now equate coefficients of $t^{(p+q)-k}$. (Please forgive the re-use of the letters i and j.) □

Pulling this back under a diagonal map gives us the *Internal Cartan formula*

$$\mathrm{Sq}^k(xy) = \sum_{i+j=k} (\mathrm{Sq}^i x)(\mathrm{Sq}^j y)\,.$$

The Cartan formula can be conveniently packaged in the following way. Introduce a formal variable t of dimension -1 and define the "total square"

$$\mathrm{Sq}_t : H^*(X) \to H^*(X)[t]\,, \quad \mathrm{Sq}_t x = \sum (\mathrm{Sq}^k x) t^k$$

Then

$$\mathrm{Sq}_t(xy) = (\mathrm{Sq}_t x)(\mathrm{Sq}_t y)\,.$$

We don't (yet) know that $\mathrm{Sq}^0 = 1$ in all dimensions, but we do at least know that $\mathrm{Sq}^0 : H^0 \to H^0$ is the identity, since the top square squares, and in dimension zero squaring is the identity. Thus

$$\mathrm{Sq}_t 1 = 1\,.$$

So the Cartan formula says that the total square is a ring homomorphism.

At this point the only squares we know to be non-trivial are the top squares, $\mathrm{Sq}^n : \overline{H}^n \to \overline{H}^{2n}$. At the other extreme, we would like to know how

$$\mathrm{Sq}^0 : \overline{H}^q \to \overline{H}^q$$

acts. Since $\overline{H}^q(K_n) = \mathbb{F}_2$ (by the Hurewicz and universal coefficient theorems), there are only two natural endomorphisms of $\overline{H}^q$: the zero map and the identity map. Which is Sq^0? To see that it is not the zero map, we need just one example. Let's start with $q = 1$.

Lemma 6.3. $\mathrm{Sq}^0 : \overline{H}^1(S^1) \to \overline{H}^1(S^1)$ *is nonzero.*

To prove this we need to analyze $D_2(S^1) = E\pi_+ \wedge_\pi (S^1)^{\wedge 2}$. Reduced extended powers (5.3) of spheres have a description in terms of Thom spaces.

Proposition 6.4. *Let* $\pi \subseteq \Sigma_n$. *The reduced extended power*

$$D_\pi(S^q) = Th(q\xi \downarrow B\pi)\,,$$

where $\xi = E\pi \times_\pi \mathbb{R}\pi$ *is the vector bundle associated to the regular representation of* π *on* $\mathbb{R}\pi$.

Proof. Since $(S^q)^n / W_n(S^q) = (S^q)^{\wedge n} = S^{qn}$,

$$D_\pi(S^q) = \frac{E\pi \times_\pi (S^q)^n}{E\pi_\pi W_n(S^q)} = \frac{E\pi \times_\pi S^{qn}}{E\pi \times_\pi *}\,.$$

But as a π-space, S^{qn} is just the one-point compactification of $q\mathbb{R}\pi$, so the right hand side is $Th(q\xi)$. □

Thus (with π the group of order 2)

$$D_2 S^1 = Th(\xi \downarrow B\pi)$$

where ξ is the vector bundle associated to the regular representation of π on $\mathbb{R}\pi$. This representation is isomorphic to the direct sum of the sign representation and the trivial representation. By (4.3), this implies that

$$D_2 S^1 = \Sigma Th(\lambda \downarrow \mathbb{R}P^\infty) = \Sigma \mathbb{R}P^\infty\,.$$

In order to compute Steenrod operations on $\sigma \in H^1(S^1)$, we have to go on to identify the fiber inclusion (5.4) and diagonal (5.5):

$$(S^1)^\wedge 2 \xrightarrow{i} D_2 S^1 \xleftarrow{j} \mathbb{R}P^\infty_+ \wedge S^1 \,.$$

under the identification of D_2S^1 with a Thom space. The map i is the inclusion of the Thom space of ξ over a point. The map j is induced by the inclusion of the trivial representation into $\mathbb{R}\pi$. The class $P\sigma$ is the Thom class:

$$P\sigma = t \wedge \sigma \in \overline{H}^2(\Sigma \mathbb{R}P^\infty) \,.$$

The map j restricts to the identity on $\mathbb{R}P^\infty$, so $j^*(t \wedge \sigma) = t \wedge \sigma$ and in

$$j^*(t \wedge \sigma) = 1 \wedge \mathrm{Sq}^1\sigma + t \wedge \mathrm{Sq}^0\sigma \,,$$

we must have $\mathrm{Sq}^1\sigma = 0$ (of course) and $\mathrm{Sq}^0\sigma = \sigma$. This completes the proof of 6.3.

Combined with the Cartan formula, this example implies that the Steenrod squares are "stable" operations: smashing with the generator $\sigma \in \overline{H}^1(S^1)$ gives the suspension isomorphism

$$\overline{H}^q(X) \xrightarrow{\cong} \overline{H}^{q+1}(\Sigma X) \,.$$

Since Sq^0 is the only Steenrod operation that is nontrivial on σ, the Cartan formula collapses to

$$\mathrm{Sq}^k(x \wedge \sigma) = (\mathrm{Sq}^k x) \wedge \sigma \,. \tag{6.2}$$

It is now easy to deduce:

Theorem 6.5. *For any n, $\mathrm{Sq}^0 : \overline{H}^n \to \overline{H}^n$ is the identity map.*

Proof. Again, a single example suffices. We can use $S^n = (S^1)^{\wedge n}$, by stability. □

Proposition 6.6. *For all q and k, $\mathrm{Sq}^k : \overline{H}^q \to \overline{H}^{q+k}$ is a homomorphism.*

Proof. This is implied by the fact that Sq^k is stable. Here is one way to see that. Since cohomology is represented by Eilenberg Mac Lane spaces, we can think of $\mathrm{Sq}^k : \overline{H}^q \to \overline{H}^{q+k}$ as a homotopy class

$$\mathrm{Sq}^k : K_q \to K_{q+k} \,,$$

namely the one pulling ι_{q+k} back to $\mathrm{Sq}^k \iota_q$. The functor $X \mapsto \overline{H}^{q+1}(\Sigma X)$ is representable too: The functor Σ has a right adjoint given by the looping functor

$$\Omega X = \mathrm{map}_*(S^1, X)\,,$$

so

$$\overline{H}^{q+1}(\Sigma X) = [\Sigma X, K_{q+1}]_* = [X, \Omega K_{q+1}]_*\,.$$

The suspension isomorphism is represented by a homotopy equivalence

$$K_q \xrightarrow{\sigma} \Omega K_{q+1}\,.$$

and the stability of Sq^k is equivalent to the commutativity of the diagram

$$\begin{array}{ccc} K_q & \xrightarrow{\mathrm{Sq}^k} & K_{q+k} \\ \downarrow{\scriptstyle\sigma} & & \downarrow{\scriptstyle\sigma} \\ \Omega K_{q+1} & \xrightarrow{\Omega \mathrm{Sq}^k} & \Omega K_{q+k+1}\,. \end{array} \tag{6.3}$$

The identification of K_q with ΩK_{q+1} imposes a natural addition on $[X, K_q]_*$. This addition is none other than the additive structure of $\overline{H}^q(X)$, and the diagram (6.3) implies that Sq^k respects that addition. □

Let's see what Steenrod operations have to say about our co-reducibility question. We have to compute their effect in the cohomology of projective space. This is easy, since $H^*(\mathbb{R}P^\infty)$ is generated by the class t in dimension 1. Only Sq^0 and Sq^1 can be nonzero on it, and both are. In terms of the total square, we can say

$$\mathrm{Sq}\, t = t(1+t)$$

so, by the Cartan formula,

$$\mathrm{Sq}\, t^n = t^n(1+t)^n = \sum_{i=0}^{n} \binom{n}{i} t^{n+i}\,.$$

In $H^*(\mathbb{R}P^{k-1})$ the formulas are the same; but $t^k = 0$.

We are interested in the stunted projective space $\mathbb{R}P_n^{n+k-1}$. Its reduced cohomology is the subspace of $H^*(\mathbb{R}P^{n+k-1})$ generated by t^i for $n \le i < n+k$, so we know what Steenrod operations do in it as well. We are interested in whether the bottom cell splits off. If it does, the map $\mathbb{R}P_n^{n+k-1} \to S^n$

induces an injection in cohomology, so all positive dimensional operations on t^n must vanish. This is equivalent to requiring that

$$\binom{n}{i} = 0 \mod 2 \quad \text{for all } i \text{ with } 0 < i < k\,.$$

Writing n in its binary expansion gives an expression for $(1+t)^n$ as a product of terms of the for $(1+t^{2^i})$. The lowest positive dimensional term corresponds to the lowest power of 2 in the expression for n, which is to say the largest power of 2 dividing n:

$$\max\left\{k : \binom{n}{i} = 0 \mod 2 \quad \text{for all } i \text{ with } 0 < i < k\right\} = 2^{\nu(n)}\,.$$

Therefore the bottom cell cannot split off of $\mathbb{R}P_n^{n+k-1}$ unless $k \leq 2^{\nu(n)}$. From this, and Proposition 4.2, we conclude:

Theorem 6.7. *S^{n-1} admits at most $2^{\nu(n)}-1$ everywhere linearly indpendent vector fields.*

For example, if n is odd then there are no nowhere vanishing vector fields on S^{n-1}. For another example:

Corollary 6.8. *If S^{n-1} is parallelizable then n is a power of 2.*

This is a celebrated theorem of José Adem. It was a major vindication of the "Adem relations" among Steenrod operations. But using the coreduction approach avoids the need to invoke these relations.

Nevertheless, this is not best possible. Theorem 1.4 gives the optimal statement. It turns out for example that S^{15} is not parallelizable, and in fact does not admit even nine everywhere independent vector fields. For the general case, we will need to probe stunted projective spaces using a different tool, namely, Adams operations in K-theory. In general it does turn out that the maximal number of independent vector fields depends only on $\nu(n)$, but for $\nu(n) > 3$ it is significantly less than $2^{\nu(n)}$.

7 The Adem relations

placeholder

8 Chern classes

It may come as a surprise that we have come this far in analyzing the structure of various vector bundles without ever mentioning characteristic classes. This state of affairs cannot continue, however, and in this lecture we will set up the theory of Chern classes.

The starting point is the functor from spaces to sets given by sending a space X to the set of isomorphism classes of complex vector bundles of fiber dimension n. To reassure you that this is a set (and not something bigger), recall that a vector bundle is by definition locally trivial, so any n-plane bundle is isomorphic to one obtained by gluing trivial bundles over intersections of open subsets. Write $\mathrm{Vect}_n(X)$ for this set. The pull-back of vector bundles renders it a contravariant functor

$$\mathrm{Vect}_n : \mathbf{Top} \to \mathbf{Set} \,.$$

This functor is quite difficult to understand in general. One approach to studying it is to define natural transformations from it to more understandable functors. One example, that is in some sense universal, is given by K-theory, and we will study this in detail in later lectures. Another example is ordinary cohomology, and we will study that in this lecture. The relevant definition is:

Definition 8.1. *A* characteristic class *for complex n-plane bundles is a natural transformation*

$$\mathrm{Vect}_n(X) \to H^q(X; R)$$

for some q and some coefficient group R.

A similar definition applies to real vector bundles, oriented real vector bundles, and a variety of other structured bundles. We have already seen one example: the Euler class is characteristic for R-oriented n-plane bundles.

Theorem 8.2. *There is a family of characteristic classes (the* Chern classes*)*

$$c_i^{(n)} : \mathrm{Vect}_n(X) \to H^{2i}(X; \mathbb{Z}) \quad , 0 \leq i \leq n \,,$$

such that
(1) $c_0^{(n)}(\xi) = 1 \in H^0(X)$,
(2) $c_1^{(1)}(\lambda) = -e(\lambda) \in H^1(X)$ *for any line bundle λ, regarding it as a real*

vector bundle oriented by saying that for any nonzero vector v the ordered basis (v, iv) is compatible with the orientation, and
(3) the "Whitney sum formula" holds: if ξ is of fiber dimension m and η is of fiber dimension n, then

$$c_k^{(m+n)}(\xi \oplus \eta) = \sum_{i+j=k} c_i^{(m)}(\xi) \cdot c_j^{(n)}(\eta) \, .$$

The axioms identify the Chern classes in the case of complex line bundles; $c_0 = 1$, and $c_1^{(1)}$ is just the negative of the Euler class of the underlying oriented real 2-plane bundle. The Euler class of a trivial bundle (of positive dimension) is zero, so the trivial line bundle $\epsilon^{\mathbb{C}}$ has trivial first Chern class.

So by the Whitney sum fomula,

$$c_k^{(n+1)}(\xi \oplus \epsilon^{\mathbb{C}}) = c_k^{(n)}(\xi) \, .$$

This "stability" suggests that we just drop the superscript, and we do so.

Under certain very general conditions the functor Vect_n is representable. One may restrict the type of vector bundle: A vector bundle is *numerable* if it admits a trivializing open cover that supports a subordinate partition of unity. Or one may restrict the type of base space: over a paracompact space (e.g. a CW complex) any vector bundle is numerable.

Under these conditions there is a universal n-plane bundle ξ_n over a space $BU(n)$. The Yoneda lemma assures us that there is a bijection

$$H^q(BU(n); R) \to \mathrm{nt}_X(\mathrm{Vect}_n(X), H^q(X; R)) \, .$$

The cohomology of $BU(n)$ carries complete information about characteristic classes for complex n-plane bundles.

The "classifying space" $BU(n)$ is only well defined up to homotopy equivalence, but it may be taken to be a CW complex. In fact, it admits the following very concrete description. Form the Stiefel manifold $V_n(\mathbb{C}^{n+k})$ of Hermitian orthogonal n-frames in $\mathbb{C}^{n+k}$. This is the space of inner-product preserving linear maps $\mathbb{C}^n \to \mathbb{C}^{n+k}$, and as such it admits a free right action of the unitary group $U(n)$ by precomposition. The orbit space for this action is the Grassmannian $\mathrm{Gr}_n(\mathbb{C}^{n+k})$ of n-dimensional subspaces in $\mathbb{C}^n$. For example, $V_1(\mathbb{C}^{1+k})$ is the unit sphere in $\mathbb{C}^{1+k}$, and $\mathrm{Gr}_1(\mathbb{C}^{1+k}) = \mathbb{C}P^k$.

The Grassmannian supports a canonical vector bundle, built as the Borel construction

$$\xi_{n,k} : V_n(\mathbb{C}^{n+k}) \times_{U(n)} \mathbb{C}^n \downarrow \mathrm{Gr}_n(\mathbb{C}^{n+k}) \, .$$

Theorem 8.2 has an interpretation in terms of the cohomology of $BU(n)$. Write

$$c_i \in H^{2i}(BU(n)) \quad , 0 \leq i \leq n \quad , c_0 = 1$$

for the Chern classes of the universal n-plane bundle; these are the universal Chern classes. To capture the Whitney sum formula, consider the $(m+n)$-plane bundle $\xi_m \times \xi_n$ over $BU(m) \times BU(n)$. It is represented by a well defined map

$$\mu : BU(m) \times BU(n) \to BU(m+n)\,,$$

and these maps are associative and unital since the Whitney sum operation is.

Theorem 8.3. *There exists a unique family of classes*

$$c_i \in H^{2i}(BU(n)) \quad , 0 \leq i \leq n\,,$$

such that

$$c_0 = 1 \in H^0(BU(n))\,,$$

$$c_1(\lambda) = -e(\lambda) \in H^2(BU(1))\,,$$

and

$$\mu^* c_k = \sum_{i+j=k} c_i \times c_j \in H^*(BU(m) \times BU(n))\,.$$

Furthermore,

$$H^*(BU(n)) = \mathbb{Z}[c_1, \ldots, c_n]\,.$$

This theorem goes beyond the previous one in two ways. First, it tells us that polynomials in the Chern classes exhaust all characteristic classes for complex vector bundles – with any coefficients, not just with $\mathbb{Z}$ coefficients, since the universal coefficient theorem tells us that

$$H^*(BU(n); R) = R[c_1, \ldots, c_n]$$

for any coefficient ring R. Second, it tells us that there are no algebraic relations among the Chern classes that are valid for all vector bundles.

The proof of this theorem will go by induction on n, using the simplest relationship between successive classifying spaces: the map

$$BU(n-1) \to BU(n)$$

representing $\xi_{n-1} \oplus \epsilon^{\mathbb{C}}$.

Lemma 8.4. *There is a homotopy equivalence*

$$\begin{array}{ccc} BU(n-1) & \xrightarrow{\simeq} & \mathbb{S}(\xi_n) \\ & \searrow \quad \swarrow & \\ & BU(n) & \end{array}$$

compatible with the indicated maps to $BU(n)$*.*

Proof. There are many ways to see this. One is to use the Grassmannian models. These are homogeneous spaces;

$$\mathrm{Gr}^{(}_n\mathbb{C}^{n+k}) = U(n+k)/U(n) \times U(k)\,,$$

where the subgroup sits in $U(n+k)$ as block diagonal matrices. We have the following fiber bundle:

$$\frac{U(n+k-1)}{U(n-1) \times U(k)} \to \frac{U(n+k)}{U(n-1) \times 1 \times U(k)} \to \frac{U(n+k)}{U(n+k-1)}\,.$$

Using the identification of $U(n)/U(n-1)$ with the unit sphere in $\mathbb{C}^n$, we can rewrite this fiber bundle as

$$\mathrm{Gr}_{n-1}(\mathbb{C}^{n+k-1}) \to \mathbb{S}(\xi_{n,k} \downarrow \mathrm{Gr}_n(\mathbb{C}^{n+k})) \to S^{2(n+k)-1}\,.$$

These bundles are compatible as k increases, and in the limit the base becomes contractible and

$$\mathrm{Gr}_{n-1}(\mathbb{C}^{\infty}) \to \mathbb{S}(\xi_n \downarrow \mathrm{Gr}_n(\mathbb{C}^{\infty}))$$

becomes the homotopy equivalence we asserted. □

Another proof uses the fact that the classifying space of a topological group G is characterized as the orbit space of a free (principal, to be strictly correct) action of G on a contractible space (written EG). If H is a closed subgroup of G, then H acts freely on EG and so its orbit space is a model for BH. But

$$EG/H = EG \times_G (G/H)$$

fibers over BG with fiber G/H: we have a fiber bundle

$$G/H \to BH \to BG\,.$$

Now apply this principle to the inclusion $U(n-1) \to U(n)$ and remember again that $U(n)$ acts transitively on the unit sphere in $\mathbb{C}^n$ with isotropy group $U(n-1)$.

In any case, we have a sphere bundle $BU(n-1) \to BU(n)$. Let's analyze the relationship between the cohomology of the sphere bundle $\pi : S(\xi) \downarrow B$ of a vector bundle over B and the cohomology of B. This is the "Gysin sequence," and actually the analysis works if you have a fibration whose fibers have the homology of a sphere. We will use the Thom isomorphism theorem 4.5. The cofiber sequence

$$\mathbb{S}(\xi) \to \mathbb{D}(\xi) \to Th(\xi)$$

gives rise to a long exact sequence in cohomology, shown in the top row below.

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & \overline{H}^q(Th(\xi)) & \longrightarrow & H^q(\mathbb{D}(\xi)) & \longrightarrow & H^q(\mathbb{S}(\xi)) & \xrightarrow{\delta} & \overline{H}^{q+1}(Th(\xi)) & \longrightarrow & \cdots \\
 & & \uparrow{\scriptstyle \cup U} & & \uparrow{\scriptstyle \cong} & & \uparrow{\scriptstyle =} & & \uparrow{\scriptstyle \cup U} & & \\
\cdots & \longrightarrow & H^{q-n}(B) & \xrightarrow{\cdot e} & H^q(B) & \xrightarrow{\pi^*} & H^q(\mathbb{S}(\xi)) & \xrightarrow{\pi_*} & H^{q-n+1}(B) & \longrightarrow & \cdots
\end{array}$$

The vertical arrows are given by the Thom isomorphism and the projection map $\mathbb{D}(\xi) \downarrow B$ (a homotopy equivalence). The Euler class $e \in H^n(B)$ is the restriction of the Thom class to $H^n(B)$ (via $H^n(\mathbb{D}(\xi))$), so the map $H^0(B) \to H^n(B)$ in the bottom row sends 1 to e.

The cohomology long exact sequence is a sequence of modules over $H^*(\mathbb{D}(\xi))$, so the maps in the bottom row are $H^*(B)$-module maps, and $H^{q-n}(B) \to H^q(B)$ sends a to $a \cdot e$.

The map π_* is variously known as the "Gysin map," an "umkher homomorphism," or a "pushforward." We will have more to say about such maps later, but for now let's make note of the fact that it too is a homomorphism of $H^*(B)$-modules:

$$\pi_*(\pi^* a \cdot b) = a \cdot \pi_* b\,.$$

We now make a couple of deductions from the Gysin sequence, that will serve us well in the induction.

Assume that

$$n \text{ is even and } H^q(\mathbb{S}(\xi)) = 0 \text{ for } q \text{ odd.}$$

Then the Gysin sequence implies that

$$e|H^q(B) \text{ is monic for } q \text{ even}$$

$$e|H^q(B) \text{ is epic for } q \text{ odd.}$$

But $H^q(B) = 0$ for $q < 0$, for sure, so we conclude that

$$H^q(B) = 0 \text{ for } q \text{ odd.}$$

It follows that

$$\pi_* = 0\,,$$

since it relates even and odd dimensions. So

$$H^*(\mathbb{S}(\xi)) = H^*(B)/H^*(B)e$$

and

$$H^*(B) = H^*(\mathbb{S}(\xi))[e]\,.$$

We can now run our induction, starting with the trivial case $H^*(BU(0)) = \mathbb{Z}$. The Gysin sequence assocated to the sphere bundle of the canonical line bundle $\lambda = \xi_1$ over $BU(1) = \mathbb{C}P^\infty$ shows that

$$H^*(BU(1)) = \mathbb{Z}[e_1]$$

where e_1 is the Euler class of λ. Define $c_1 = -e_1$. For the inductive step, suppose we have classes $c_i \in H^{2i}(BU(n-1)), 1 < i < n$, such that

$$H^*(BU(n-1)) = \mathbb{Z}[c_1, \cdots, c_{n-1}]\,.$$

Then

$$H^*(BU(n)) = H^*(BU(n-1))[e]$$

where $e_n \in H^{2n}(BU(n))$ is the Euler class of ξ_n. Note that $\pi^* : H^q(BU(n)) \to H^*(BU(n-1))$ is an isomorphism for $q < 2n$, so for $0 < i < n$ there are unique classes $c_i \in H^{2i}(BU(n))$ pulling back to classes of the same name in $H^*(BU(n-1))$.

To proceed, we could take for c_n any class of the form $\pm e_n + c$ where c is a polynomial in $c_1, \ldots, c_{n-1}$. The preferred choice is

$$c_n = (-1)^n e_n\,.$$

This completes a computation of $H^*(BU(n))$ and construction of classes c_i satisfying two of the three conditions in the theorem. We now need to check the Whitney sum formula. (Whitney once called this the hardest theorem he ever proved.) For this we will use the "splitting principle."

9 The Splitting Principle

Theorem 9.1 (The Splitting Princple). *Let ξ be a complex n-plane bundle over a space B. There is a map $f : X \to B$ such that*
(1) *$f^*\xi$ splits as a sum of line bundles, and*
(2) *$f^* : H^*(B) \to H^*(X)$ is a monomorphism.*

As a result, to check an identity among Chern classes it is enough to check it on sums of line bundles.

Proof. By induction, it suffices to split off one line bundle. There is a canonical way to do this, using the *projectivization* of ξ. This is the fiber bundle over B with total space

$$\mathbb{P}(\xi) = \{(x, l) : x \in B,\, l \text{ is a one-dimensional subspace of } \xi_x\}\,.$$

The point (x, l) projects to $x \in B$. The projectivization of an n-plane bundle over a point is simply complex projective space of complex dimension $(n-1)$. In general the fiber over x is the projective space of ξ_x.

The space $\mathbb{P}(\xi)$ supports a "tautologous" line bundle λ with

$$\mathbb{E}(\lambda) = \{(l, v) : v \in l \in \mathbb{P}(\xi)\}\,.$$

Now

$$\mathbb{E}(f^*\xi) = \{(x, l, v) : x \in V_x, l \subseteq V_x\}\,,$$

so $\lambda \subseteq f^*\xi$ as the subbundle where $x \in l$. A choice of metric expresses $f_*\xi$ as $\lambda \oplus \lambda^\perp$.

It remains to show that $f^* : \mathbb{P}(\xi) \to B$ induces a monomorphism in cohomology. Let's analyze $H^*(\mathbb{P}(\xi))$ using the Serre spectral sequence for the fibration:

$$E_2^{s,t} = H^s(B; H^*(\mathbb{C}P^{n-1})) \implies H^*(\mathbb{P}(\xi))\,.$$

Can the coefficient system be non-trivial? Certainly B need not be simply connected. But ξ is a complex vector bundle, so it is a pullback of the universal n-plane bundle $\xi_n \downarrow BU(n)$. $BU(n)$ *is* simply connected (because $U(n)$ is connected), so the coefficient system given by $\mathbb{P}(\xi_n)$ is trivial. It pulls back to the coefficient system for $\mathbb{P}(\xi)$, which is therefore also trivial.

So (by the Künneth theorem)

$$E_2^{s,t} = H^s(B) \otimes \mathbb{Z}[e]/(e^n)\,,$$

where e is the generator of $H^*(\mathbb{C}P^{n-1})$.

Does the spectral sequence collapse? If so, we are done, because the kernel of $f^* : H^*(B) \to H^*(\mathbb{P}(\xi))$ is the sum of the images of the differentials hitting the bottom row in the spectral sequence. We should try to compute differentials on the algebra generators of $E_2^{*,*}$. There are the generators of $H^*(B)$ along the base, and these certainly survive because all differntials on them land below the horizontal axis. There is only one more generator, namely the element $e \in E_2^{0,2}$. (Actually, there's one such element for every component of the base space. But we can handle the components separately.)

Under the identification

$$E_2^{0,t} \cong H^t(B)$$

(assuming B is connected and the coefficient system is trivial!) an element on the vertical axis survives if and only if it extends to cohomology class in the total space.

Well, what is the generator $e \in H^2(\mathbb{C}P^{n-1})$? It's the Euler class of the tautologous line bundle λ_x over the fiber. Since e is a characteristic class, it will extend to a class in $H^2(\mathbb{P}(\xi))$ if λ_x extends to a line bundle over $\mathbb{P}(\xi)$ ...and it does!

This completes the proof of the splitting principle. □

The spectral sequence comes close to computing the cohomology of the projectivization. The fact that it collapses gives us a natural multiplicative isomorphism

$$\mathrm{gr}^s H^*(\mathbb{P}(\xi)) \cong H^s(B; H^*(\mathbb{C}P^{n-1}))$$

where we are forming the associated graded algebra with respect to the filtration coming from the spectral sequence. This decreasing filtration does not exhibit the module structure over $H^*(B)$, but it can be modified to an increasing filtration that does (following Quillen [4]): Define

$$F_t H^n(\mathbb{P}(\xi)) = F^{n-t} H^n(\mathbb{P}(\xi))\,.$$

Then, for example,

$$F_0 H^n(\mathbb{P}(\xi)) = F^n H^n(\mathbb{P}(\xi)) = \mathrm{gr}^n H^n(\mathbb{P}(\xi)) \cong H^n(B)$$

since $F^{n+1}H^n(\mathbb{P}(\xi)) = 0$. This is a filtration of $H^*(\mathbb{P}(\xi))$ by $H^*(B)$-submodules, with

$$\mathrm{gr}_* H^*(\mathbb{P}(\xi)) = H^*(B)[e]/(e^n)$$

Since the associated quotients are free modules over $H^*(B)$, we find by induction that $H^*(\mathbb{P}(\xi))$ is a free module; and we know a set of generators: $\{1, e, \ldots, e^{n-1}\}$. Then e^n can be expressed as an $H^*(B)$-linear combination of these generators. We choose to express this relation as a monic polynomial satisfied by $e \in H^*(\mathbb{P}(\xi))$:

$$e^n + c_1 e^{n-1} + \cdots + c_{n-1} e + c_n = 0\,.$$

We have written the coefficients this way because, in fact, this gives a different definition of the Chern classes of ξ. This approach is due to Alexander Grothendieck. On the other hand, this can be viewed as a functorial description of the cohomology of $\mathbb{P}(\xi)$ in terms of $H^*(B)$ and the Chern classes of ξ.

By iterating the projectivization process, we can produce a map $f : X \to B$ canonically associated to the complex vector bundle ξ such that $f^*\xi$ splits into a sum of line bundles. This is the *flag bundle* associated to ξ. The total space is given by

$$Fl(\xi) = \left\{ (x, l_1, \ldots, l_n) : \begin{array}{l} x \in B, l_1, \ldots, l_n \text{ one-dimensional} \\ \text{subspaces of } \xi_x \text{ whose sum is } \xi_x \end{array} \right\}.$$

Exercise 9.2. *There are many variants of the flag bundle construction. Given a partition*

$$n = n_1 + \cdots + n_k$$

we may consider the fiber bundle $Fl_{n_1,\ldots,n_k}(\xi)$ over X whose fiber over x is the space of direct sum decompositions of ξ_x as sum of k subspaces where the ith one has dimension n_i. Find a functorial expression for $H^*(Fl_{n_1,\ldots,n_k}(\xi))$ in terms of $H^*(B)$ and the Chern classes of ξ. In particular, give a description of the generalized flag manifold, which is the case $B = *$.

We are now aiming at a proof of the Whitney sum formula. Along the way, we will need to better understand the notion of a classifying space. In Lecture 5 we used classifying spaces for finite groups, and we are now studying the classifying space for n-plane bundles. The relationship between these concepts occurs through the notion of a principal bundle.

Definition 9.3. *Let G be a topological group. A* principal G-bundle *is a fiber bundle $P \downarrow B$ which is expressed as the orbit projection of a free continuous right action of G on P.*

A complex n-plane bundle ξ with a Hermitian metric determines a principal $U(n)$-bundle, namely the *frame bundle*, with total space

$$Fr(\xi) = \{(x,t) : x \in B, t : \mathbb{C}^n \to \xi_x \text{an isometry}\}\,.$$

The structure group $U(n)$ acts from the right on this space by precomposition, and

$$E(\xi) \cong Fr(\xi) \times_{U(n)} \mathbb{C}^n\,.$$

In fact all the fiber bundles associated to ξ can be built using the Borel construction from $Fr(\xi)$:

$$\mathbb{S}(\xi) = Fr(\xi) \times_{U(n)} S^{2n-1}\,, \quad \mathbb{P}(\xi) = Fr(\xi) \times_{U(n)} \mathbb{C}P^{n-1}\,,$$

and

$$Fl(\xi) = Fr(\xi) \times_{U(n)} (U(n)/T^n)\,.$$

Two principal G-bundles P, P' over a base B are "isomorphic" if there is an equivariant homeomorphism $P \to P'$ covering the identity on B. Principal bundles pull back under maps, as do isomorphisms between them. The set of isomorphism classes of principal G bundles over B forms a set, and we have contravariant functor

$$X \mapsto \mathrm{Bun}_G(X)\,.$$

Under numerability or paracompactness assumptions, this functor is representable, and the representing object is the "universal" principal G-bundle

$$EG \downarrow BG\,.$$

It is a theorem that the universal principal G-bundle is characterized by the fact that EG is contractible.

Suppose that $f : H \to G$ is a continuous homomorphism and $P \downarrow B$ is a principal H-bundle. We can construct a principal G-bundle over B with total space

$$P \times_H G$$

where $g \in G$ acts by $(x, g')g = (x, gg')$. This operation respects isomorphisms, and defines a natural transformation $f_* : \mathrm{Bun}_H \to \mathrm{Bun}_G$.

By the Yoneda lemma, this natural transformation is represented by a unique homotopy class of maps $Bf : BH \to BG$. Uniqueness shows that we obtain a functor B from topological groups and continuous homomorphisms into the homotopy category.

Given $g \in G$, write $c_g : G \to G$ for the continuous homomorphism $\gamma \mapsto g\gamma g^{-1}$. These homomorphisms have a simple effect on BG:

Lemma 9.4. *Bc_γ is homotopic to the identity map on BG.*

Proof. To prove that $Bc_g \simeq 1$, we need to show that $(c_g)_* = 1 : \mathrm{Bun}_G \to \mathrm{Bun}_G$. So let $P \downarrow B$ be a principal G bundle. $(c_g)_*P$ is again P as a space but with G acting through c_g. Using $\cdot$ to denote this twisted action,

$$x \cdot h = xghg^{-1}$$

The isomorphisms are given by

$$(c_g)_*P \to P \quad \text{by} \quad x \mapsto xg\,,$$

$$P \to (c_g)_*P \quad \text{by} \quad x \mapsto xg^{-1}\,.$$

We leave it to the listener to check that these maps are equivariant. □

We can now use this machinery, and the splitting theorem, to give a different interpretion of the Chern classes, one which will immediately imply the Whitney sum formula. We are studying the inclusion

$$H^*(BU(n)) \hookrightarrow H^*(T^n)\,.$$

Conjugation by any element of $U(n)$ acts as the identity on $BU(n)$. The normalizer of T^n, $N_{U(n)}T^n$, is the set of elements that conjugate T^n into itself. The only elements in $U(n)$ that do that are the unitary matrices with n nonzero entries. There must be one in each row and one in each column, and the nonzero entries must be complex numbers of absolute value 1. This subgroup contains T^n as a normal subgroup, and the quotient group is the symmetric group Σ_n, acting by permuting the entries, and we have a split extension sequence

$$1 \to T^n \to N_{U(n)}T^n \to \Sigma_n \to 1\,.$$

In the parlance of Lie groups, T^n is a maximal torus in $U(n)$, and Σ_n is the Weyl group.

$N_{U(n)}T^n$ acts on the homotopy type of BT^n, and T^n acts trivially (since T^n is abelian, but actually conjugation by elements of T^n act as the identity on BT^n anyway), so we receive an action of Σ_n on the homotopy type of BT^n and hence on its cohomology. Σ_n also acts on $BU(n)$ and on its cohomology, but trivially, by Lemma 9.4. So the inclusion in cohomology factors as

$$H^*(BU(n)) \hookrightarrow H^*(BT^n)^{\Sigma_n}\,. \tag{9.1}$$

The action of Σ_n simply permutes the generators of $H^*(BT^n)$, so we are looking at a well-studied situation. Over any coefficient ring R, the symmetric invariants of the Σ_n action on $H^*(BT^n; R) = R[t_1, \ldots, t_n]$ form a polynomial subalgebra on generators given by the "elementary symmetric polynomials," $\sigma_j \in H^{2j}(BT^n; R)$, $1 \leq j \leq n$.

The inclusion (9.1) is valid for any coefficient ring; for example, for every field. Then a dimension count implies that the map is an isomorphism for every field, and that is enough to see that it is an isomorphism with coefficients in $\mathbb{Z}$ (and hence for any coefficient ring):

Proposition 9.5. *The map $BT^n \to BU(n)$ induces an isomorphism*

$$H^*(BU(n)) \xrightarrow{\cong} H^*(BT^n)^{\Sigma_n} .$$

We should be more explicit about the structure of the ring of symmetric invariants. The elementary symmetric polynomials

$$\sigma_j = \sigma_j(t_1, \ldots, t_j)$$

are defined using the equation

$$\prod_{i=1}^{n}(t - t_i) = \sum_{j=0}^{n} \sigma_j t^{n-j} . \tag{9.2}$$

One should think of the ring of invariants as generated by the coefficients of the general degree n monic polynomial. The larger ring $H^*(BT^n)$ is then obtained by formally adjoining the n roots of the polynomial.

We have taken the liberty of omitting the rank n from the notation for the elementary symmetric polynomials. This is justified by the observation that the map $R[t_1, \ldots, t_n] \to R[t_1, \ldots, t_{n-1}]$ sending t_n to zero is equivariant with respect to Σ_{n-1} (embedded into Σ_n as the isotropy group of $n \in \{1, \ldots, n\}$), and sends σ_i to σ_i for $i < n$ and to zero for $i = n$. This is the effect in cohomology of the inclusion of the first $(n-1)$ factors into BT^n and on invariants it reflects the map $BU(n-1) \to BU(n)$.

Where do the Chern classes map to under (9.1)? The top Chern class, c_n, is $(-1)^n$ times the Euler class of the canonical n-plane bundle. The map $BT^n \to BU(n)$ pulls that bundle back to the product of the canonical line bundles over the factors, so by Lemma 4.8 the Euler class pulls back to the product of the Euler classes t_i. By (9.2),

$$\sigma_n = (-1)^n t_1 \cdots t_n$$

as well, so c_n maps to σ_n. Since lower Chern classes correspond under $BU(n-1) \to BU(n)$, as do lower elementary symmetric polynomials, we find:

Proposition 9.6. *Under the map $BT^n \to BU(n)$ induced by the inclusion of the diagonal matrices,*

$$c_i \mapsto \sigma_i \,.$$

This gives us a different and very convenient way of thinking of the Chern classes, and, indeed, of computing with them. For example, since we know the effect of the Steenrod squares in $H^*(BT^n)$, we can read off the effect of the squares on the Chern classes.

For another example, we can now prove the Whitney sum formula. Let's first re-express it using the total Chern class

$$c_t(\xi) = \sum_i c_i(\xi) t^i$$

using a formal variable of dimension -2. The formula is then

$$c_t(\xi) c_t(\eta) = c_t(\xi \oplus \eta) \,.$$

We might as well do this in the universal case, with $\xi = \xi_m$ and $\eta = \xi_n$. We will take both sides of this equation and pull them back to the tori.

$$c_t(\xi) \times i_n^* c_t(\eta) = \prod_{i=1}^{m} (t - t_i) \prod_{i=n+1}^{m+n} (t - t_i) = \prod_{i=1}^{m+n} (t - t_i) = c_t(\xi_{m+n}) \,.$$

References

[1] J. Frank Adams,

[2] J. Frank Adams,

[3] Ioan James,

[4] Daniel Quillen,