4. Cohomology

Reference for this section: [Hat02, Section 4.3], [Hat, Section 3.2].

Let us look again at the formula $[B, G_n]$. This is the set of homotopy classes of maps from B into G_n. One of the difficulties of indentifying this is the fact that this is a set: since it has no other structure it is quite difficult to work with. Thus one way we could make it easier to work with is by mapping it into a group of some sort.

What kinds of groups have elements that look like $[X, Y]$? One such group is $\pi_n(Y)$, the groups where $X = S^n$. However, it turns out that in general, computing homotopy groups is very difficult (in fact, even computing the homotopy groups of spheres is incredibly difficult!) so we would like a more computable approach.

We have the following theorem:

Theorem 4.1. *Suppose that we are given a sequence of pointed spaces $X_0, X_1, \ldots$ together with weak equivalences $X_n \longrightarrow \Omega X_{n+1}$. Then the sequence of functors $h^n \colon \mathbf{CW}_*^{\mathrm{op}} \longrightarrow \mathbf{AbGp}$ given by*

$$h^n(Y) \stackrel{\text{def}}{=} [Y, \Omega^{\max(0,-n)} X_{\max(0,n)}]$$

is a generalized cohomology theory.

In order to prove this theorem we first need to discuss the meanings of all of the terms.

Definition 4.2. A *pointed space* is a space X together with a distinguished basepoint $* \in X$. In the category of pointed spaces, the *suspension* functor $\Sigma \colon \mathbf{Top}_* \longrightarrow \mathbf{Top}_*$ is given by $S^1 \wedge \cdot$, the smash product with the circle.

Definition 4.3. A *generalized cohomology theory*† is a sequence of functors $h^i \colon \mathbf{CW}_*^{\mathrm{op}} \longrightarrow \mathbf{AbGp}$ for all $i \in \mathbb{Z}$, together with natural *suspension isomorphisms*

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

such that the following axioms hold:

- **homotopy invariance:** If $f_1, f_2 \colon X \longrightarrow Y$ are two maps which are homotopic (relative to the basepoint) then the induced maps $h^i(f_1)$ and $f^i(f_2)$.
- **exactness:** Let $\alpha \colon A \hookrightarrow X$ be an inclusion of pointed spaces, and let $C\alpha$ be the mapping cone of α with $\beta \colon X \longrightarrow C\alpha$ the inclusion. Then we have an induced exact sequence

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

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

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

  is an isomorphism.

Note that using the suspension isomorphism and the exactness axiom we can get a long exact sequence in cohomology. Note that $C\beta \simeq \Sigma A$. If we let $\gamma \colon C\alpha \longrightarrow C\beta$ then $C\gamma \simeq \Sigma X$. Let η be the inclusion $C\beta \longrightarrow C\gamma$. Applying the exactness axiom to both α, β and γ we get an exact sequence

$$h^i(\Sigma X) \longrightarrow h^i(\Sigma A) \longrightarrow h^i(C\alpha) \longrightarrow h^i(X) \longrightarrow h^i(A).$$

Applying the suspension isomorphism to the two left-hand entries and noting that the map η is exactly $\Sigma\alpha$, we can rewrite this sequence as

$$h^{i-1}(X) \xrightarrow{h^{i-1}\alpha} h^{i-1}(A) \longrightarrow h^i(C\alpha) \longrightarrow h^i(X) \longrightarrow h^i(A).$$

† We are working with reduced theories; to get the usual unreduced theory of a space we add a disjoint basepoint.

Gluing these together for all i gives the usual long exact sequence in cohomology.

Now we can prove the theorem.

Proof of Theorem 4.1. The formula in the theorem gives a sequence of functors. Note that for any spaces Z and Z' the set $[Z, \Omega^2 Z']$ has a natural abelian group structure given by concatenation of loops. To check that the functors land in abelian groups we note that when $n \geq 0$,

$$[Y, X_n] = [Y, \Omega X_{n+1}] = [Y, \Omega^2 X_{n+2}],$$

since $X_n \simeq \Omega X_{n+1} \simeq \Omega X_{n+2}$. It works analogously for $n \leq 0$.

To construct the suspension isomorphism, we note that

$$h^{n+1}(\Sigma Y) = [\Sigma Y, \Omega^{\max(0,-(n+1))} X_{\max(0,n+1)}] \cong [Y, \Omega\Omega^{\max(0,-(n+1))} X_{\max(0,n+1)}] \cong h^n(Y).$$

The middle isomorphism is the natural isomorphism given by the adjunction between Σ and Ω, so this gives a natural isomorphism.

Homotopy invariance is clear from the definition. Additivity follows from the universal property of coproducts. Thus it remains to check exactness. Consider any inclusion $\alpha\colon A \longrightarrow Y$ and let $C\alpha$ be the mapping cone. Consider the sequence

$$[C\alpha, X_n] \longrightarrow [Y, X_n] \longrightarrow [A, X_n]$$

with the maps given by precomposition. Since both α and β are inclusions, the maps are actually given by restriction. Consider any map $f\colon C\alpha \longrightarrow X_n$. Since $C\alpha = CA \cup Y$ this gives a map $CA \longrightarrow X_n$, which is a map $A \times I \longrightarrow X_n$ which is constant on $A \times \{0\}$. This is exactly a homotopy from the map $A \times \{0\} \longrightarrow X_n$ to a constant map; thus given any map $C\alpha \longrightarrow X_n$, when the map is restricted to a map $A \longrightarrow X_n$ it will be null-homotopic. Thus the image of the left map is contained in the kernel of the right.

Now suppose that we have a map $f\colon Y \longrightarrow X_n$ which is null-homotopic when restricted to A. Let $h\colon A \times I \longrightarrow X_n$ be a null-homotopy of this map. Then we can define a map $C\alpha \longrightarrow X_n$ by defining it to be f on Y and h on CA; by definition, these give a continuous map $C\alpha \longrightarrow X_n$, as desired. Thus the sequence is exact. □

In fact, by the *Brown representability theorem* the converse of this theorem also holds: any cohomology theory arises from such a sequence of spaces. Such a sequence is called an Ω-*spectrum*; the theory of spectra is an integral part of algebraic topology. We will see other spectra later in the course when we discuss K-theory.

Example 4.4. Let $X_n = K(G, n)$ for $n \geq 0$, and let $X_n = *$ for $n < 0$. Recall that for an abelian group G, $K(G, n)$ is the space which has $\pi_n K(G, n) = G$ and all others trivial. It is a theorem that such spaces are uniquely defined up to weak equivalence. Since $\pi_i \Omega X \cong \pi_{i+1} X$, we see that $\Omega K(G, n) = K(G, n-1)$; thus by Theorem 4.1 this produces a generalized cohomology theory.

Consider S^n. We have

$$[S^n, K(G, m)] = \begin{cases} G & \text{if } m = n \\ 0 & \text{otherwise.} \end{cases}$$

In particular, when $n = 0$ this is G at $m = 0$ and 0 otherwise. Therefore this cohomology theory satisfies all of the Eilenberg–Steenrod axioms and it agrees with cellular cohomology with coefficients in G.

How does this help us? Because of exactness, it is possible to compute cohomology in a vast variety of cases by buliding up from small CW complexes to larger ones. Cohomology also has a lot of structure that makes it an exceptionally effective computational tool. For example, unlike homology, cohomology has a ring structure. Suppose that $h^n(Y) = [Y, X_n]$. If the spaces X_i are a *ring spectrum*, in the sense that there are maps $\mu_{mn}: X_m \wedge X_n \longrightarrow X_{m+n}$ (which are appropriately associative) then there is a product operation

$$[Y, X_m] \times [Y, X_n] \longrightarrow [Y, X_m \times X_n] \xrightarrow{\mu_{mn}\circ} [Y, X_{m+n}]$$

which gives a ring structure on the cohomology $h^*(Y) \stackrel{\text{def}}{=} \bigoplus_{n\in\mathbb{Z}} h^n(Y)$.

Example 4.5. Let $G = \mathbb{Z}/2$. Take the CW structure on S^∞ which has 2 cells in each dimension. $\mathbb{Z}/2$ acts on this by swapping opposite cells. Then $\mathbf{R}P^\infty$, which is $S^\infty/(\mathbb{Z}/2)$ has exactly once cell in each dimension. The boundary of an even-dimensional cell is twice the cell in one dimension lower; the boundary of an odd-dimensional cell is 0. Thus $H^n(\mathbf{R}P^\infty; \mathbb{Z}/2) \cong \mathbb{Z}/2$ for all $n \geq 0$.

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

In fact, similar proofs to the above show that $H^*(\mathbf{C}P^\infty; \mathbb{Z}) \cong \mathbb{Z}[x]$ with $|x| = 2$ and $H^*(\mathbf{H}P^\infty; \mathbb{Z}) \cong \mathbb{Z}[x]$ with $|x| = 4$.

The spaces $\mathbf{R}P^\infty$, $\mathbf{C}P^\infty$ and $\mathbf{H}P^\infty$ are exactly the Grassmannians of lines in real, complex and quaternionic space, respectively. Thus we may hope that it might be possible to determine the cohomology of G_n for general n and use it to produce invariants of vector bundles.

Theorem 4.6.

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

There exists a relatively straightforward proof of this theorem, which unfortunately relies on existence of Steifel–Whitney classes (which we will discuss soon). See [MS74, Section 7]. However, I do not like this proof because it has always felt morally circular to me, as the existence of these characteristic classes should morally follow from this theorem. Therefore we will do the calculation in a more difficult but "cleaner" way. The reader is strongly encouraged to go back and read the slicker proof later, as the ideas in it are quite beautiful.

Our proof relies on two theorems. We begin with the notion of the Thom space associated to a vector bundle; this will be an important concept later when we discuss cobordism.

Definition 4.7. Given a vector bundle $p: E \longrightarrow B$ on a paracompact base B, the *Thom space* $\mathrm{Th}(E)$ is the space $D(E)/S(E)$. Here, $D(E)$—the *disk bundle*—is the subspace of all vectors of length at most 1 and $S(E)$—the *sphere bundle*—is the subspace of all vectors of length exactly 1. (When B is not paracompact we can define the Thom space to be the 1-point compactification of E.)

Theorem 4.8 (Thom Isomorphism Theorem, [Hat02, Theorem 4D.10])**.** *There exists a class $c \in \widetilde{H}^n(\mathrm{Th}(E); \mathbb{Z}/2)$ such that the restriction of c to any fiber is a generator of $\widetilde{H}^n(S^n; \mathbb{Z}/2)$. The map $\Phi: \widetilde{H}^i(B_+; \mathbb{Z}/2) \longrightarrow \widetilde{H}^{i+n}(\mathrm{Th}(E); \mathbb{Z}/2)$ given by $b \longmapsto p^*(b) \smile c$ is an isomorphism for all i. (Here, B_+ is B with a disjoint basepoint added.)*

When E is orientable this theorem also holds with $\mathbb{Z}$ coefficients. (See [Hat02, Theorem 4D.10].) We will work purely with $\mathbb{Z}/2$ coefficients from here on, and thus we supress them from the notation.

We have a long exact sequence in cohomology

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

By the Thom Isomorphism Theorem and the fact that $D(E)$ is homotopy equivalent to B, this is isomorphic to the exact sequence

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

where here $e = (p^*)^{-1}(j^*c)$. (This e is called the $\mathbb{Z}/2$-*Euler class*.) This sequence is called the *Gysin sequence*, and it exists for any sphere bundle.

There is a nice geometric description of the $\mathbb{Z}/2$-Euler class when B is a smooth manifold. In this case, $D(E)$ and $S(E)$ are also smooth manifolds, so we can apply Poincare duality to its definition. It turns out that the Poincare dual of the $\mathbb{Z}/2$-Euler class is the homology class of the intersection of the zero section with a generic section of $D(E)$.[‡]

Even though the Thom class c is necessarily nonzero, it is entirely possible that e is zero. However, in nice cases this is not the case.

Lemma 4.9. *For the universal bundle γ_n over G_n, the $\mathbb{Z}/2$-Euler class e is nonzero.*

The analog of this statement is also true with $\mathbb{Z}$ coefficients if we work with oriented bundles.

Proof. Note that by the definition of the Thom class, the $\mathbb{Z}/2$-Euler class is natural in the following sense: given any bundle $p\colon E \longrightarrow B$ with $\mathbb{Z}/2$-Euler class $e \in H^n(B)$ and any map $f\colon B' \longrightarrow B$, the $\mathbb{Z}/2$-Euler class of f^*E is $f^*(e)$. (This is true because in a pullback bundle the restriction of a cohomology class to the fiber is the same as the restriction of the cohomology class to a fiber in the original bundle.) Thus to show that the universal bundle has a nonzero $\mathbb{Z}/2$-Euler class it suffices to check that there exists an n-bundle with a nonzero $\mathbb{Z}/2$-Euler class. Since this will (by the classification theorem) be a pullback of the universal bundle, this will mean that the Euler class of the universal bundle is nonzero.

Consider $G_n(\mathbf{R}^{n+1})$. This is a smooth manifold which is homeomorphic to $\mathbf{R}P^n$ (by taking any n-plane to its orthogonal complement). It comes with a natural inclusion $G_n(\mathbf{R}^{n+1}) \longrightarrow G_n$, and the pullback of the universal bundle is the universal bundle over $G_n(\mathbf{R}^{n+1})$. We can visualize this universal bundle in the following manner. For a line in $\mathbf{R}^{n+1}$ we take the n-space orthogonal to that line; if we think of $\mathbf{R}P^n$ as being a quotient of the sphere, this plane is the tangent plane to the sphere. Thus this bundle is the quotient of TS^n by the map which glues opposite points on the sphere together, and glues the associated tangent planes by the identity map. Thus our goal is to construct a section of this bundle which intersects the zero section transversely.

We do this by constructing a section of the tangent bundle to the sphere which is compatible with this quotient. Let $v = (0, \dots, 0, 1)$. We define a section $S^n \longrightarrow TS^n$ by taking a point x to the orthogonal projection of v onto the tangent plane at x. For points diametrically opposite on the sphere tangent planes are parallel and thus the two projections are equal. The only points at which this projection is 0 are v and $-v$. Thus this section intersects the zero section at one point. This is a nonzero term in $H_0(G_n(\mathbf{R}^{n+1}))$, and is thus Poincare dual to a nonzero term in $H^n(G_n(\mathbf{R}^{n+1}))$, and thus has nonzero $\mathbb{Z}/2$-Euler class, as desired. □

Remark 4.10. It is interesting to compare this result with Example 2.7. In that case, we showed that the universal line bundle has no nonvanishing sections. Here, we showed something more: all universal bundles have no nonvanishing sections. This makes sense, since a single nonvanishing section allows us to "slice off" a line bundle. Thus if a universal n-plane bundle had a nonvanishing section this would mean that *all* n-bundles had a line bundle that could be sliced off. If that were the case it would be strange and unnatural, so it is good that that is not the case.

We are now ready to prove the theorem.

Proof of Theorem 4.6. We prove this by induction on n using the Gysin sequence for the universal bundle $\gamma_n \longrightarrow G_n$. We can start the induction at $n = 0$, where $G_0 = *$ and the result holds trivially.[§] Thus we can assume that $H^*(G_{n-1}) \cong \mathbb{Z}/2[w_1, \dots, w_{n-1}]$.

The sphere bundle $S(\gamma_n)$ is the space of pairs (ω, v) with ω an n-plane in $\mathbf{R}^\infty$ and $v \in \omega$ a unit vector. Thus we have a natural projection $p'\colon S(\gamma_n) \longrightarrow G_{n-1}$ given by mapping (ω, v) to the $n-1$-plane $(v^\perp \cap \omega)$. Then p' is a fiber bundle with fiber S^∞ (all unit vectors orthogonal to

[‡]We may prove this later if we have time.

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

$(v^\perp \cap \omega))$. Since S^∞ is contractible, p' induces isomorphisms on all homotopy groups, and thus also on cohomology rings. This gives a map (of rings) $\eta\colon H^*(G_n) \longrightarrow H^*(G_{n-1})$.

The Gysin sequence for p is therefore

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

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

When $i = 0$, $\smile e\colon H^0(G_n) \longrightarrow H^n(G_n)$ is injective by Lemma 4.9, and thus $\widetilde{H}^{n-1}(G_n) \cong \widetilde{H}^{n-1}(G_{n-1})$; thus w'_{n-1} exists in $H^*(G_n)$ as well. In addition, since $\widetilde{H}^*(G_{n-1})$ is generated by the w_j and η is a ring homomorphism, it is surjective in each degree. Thus the Gysin sequence above splits for all i into short exact sequences

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

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

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

Remark 4.11. Note that we did not use any properties of the real numbers other than that Thom classes exist with $\mathbb{Z}/2$ coefficients. It turns out that Thom classes exist for all complex vector bundles, so the same proof shows that $H^*G_n(\mathbf{C}^\infty) \cong \mathbb{Z}[c_1, \ldots, c_n]$ with $|c_i| = 2i$.

Remark 4.12. Note that we have uniquely characerized the w_i in terms of the Thom classes: each w_i appears as the image of a Thom class in $H^*(G_i)$, which is then uniquely translated into the w_i for G_n via the sphere bundle construction.

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

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

Corollary 4.13. *The map $\eta\colon H^*(G_n) \longrightarrow H^*(G_{n-1})$ is the map taking w_n to 0.*

The important thing to note is that pullback along the projection $G_{n-1} \simeq S(\gamma_n) \longrightarrow G_n$ is exactly η, so this corollary states that the pullback of a class $q(w_1, \ldots, w_n) \in H^*(G_n)$ along the classifying map of $\gamma_{n-1} \oplus \epsilon^1$ is $q(w_1, \ldots, w_n) \in H^*(G_{n-1})$.