Additional Topics

4.A Basepoints and Homotopy

In the first part of this section we will use the action of π_1 on π_n to describe the difference between $\pi_n(X,x_0)$ and the set of homotopy classes of maps $S^n \to X$ without conditions on basepoints. More generally, we will compare the set $\langle Z,X\rangle$ of basepoint-preserving homotopy classes of maps $(Z,z_0)\to(X,x_0)$ with the set $[Z,X]$ of unrestricted homotopy classes of maps $Z\to X$, for Z any CW complex with basepoint z_0 a 0-cell. Then the section concludes with an extended example exhibiting some rather subtle nonfinite generation phenomena in homotopy and homology groups.

We begin by constructing an action of $\pi_1(X,x_0)$ on $\langle Z,X\rangle$ when Z is a CW complex with basepoint 0-cell z_0. Given a loop γ in X based at x_0 and a map $f_0:(Z,z_0)\to(X,x_0)$, then by the homotopy extension property there is a homotopy $f_s:Z\to X$ of f_0 such that $f_s(z_0)$ is the loop γ. We might try to define an action of $\pi_1(X,x_0)$ on $\langle Z,X\rangle$ by $[\gamma][f_0]=[f_1]$, but this definition encounters a small problem when we compose loops. For if η is another loop at x_0, then by applying the homotopy extension property a second time we get a homotopy of f_1 restricting to η on x_0, and the two homotopies together give the relation $([\gamma][\eta])[f_0]=[\eta]([\gamma][f_0])$, in view of our convention that the product $\gamma\eta$ means first γ, then η. This is not quite the relation we want, but the problem is easily corrected by letting the action be an action on the right rather than on the left. Thus we set $[f_0][\gamma]=[f_1]$, and then $[f_0]([\gamma][\eta])=([f_0][\gamma])[\eta]$.

Let us check that this right action is well-defined. Suppose we start with maps $f_0,g_0:(Z,z_0)\to(X,x_0)$ representing the same class in $\langle Z,X\rangle$, together with homotopies f_s and g_s of f_0 and g_0 such that $f_s(z_0)$ and $g_s(z_0)$ are homotopic loops. These various homotopies define a map $H:Z\times I\times\partial I\cup Z\times\{0\}\times I\cup\{z_0\}\times I\times I\longrightarrow X$ which is f_s on $Z\times I\times\{0\}$, g_s on $Z\times I\times\{1\}$, the basepoint-preserving homotopy from f_0 to g_0 on $Z\times\{0\}\times I$, and the homotopy from $f_s(z_0)$ to $g_s(z_0)$ on $\{z_0\}\times I\times I$. We would like to extend H over $Z\times I\times I$. The pair $(I\times I, I\times\partial I\cup\{0\}\times I)$ is homeomorphic to $(I\times I, I\times\{0\})$, and via this homeomorphism we can view H as a map $Z\times I\times\{0\}\cup\{z_0\}\times I\times I\longrightarrow X$, that is, a map $Z\times I\to X$ with a homotopy on the subcomplex $\{z_0\}\times I$. This means the homotopy extension property can be applied to produce an extension of the original H to $Z\times I\times I$. Restricting this extended H to $Z\times\{1\}\times I$ gives a basepoint-preserving homotopy $f_1\simeq g_1$, which shows that $[f_0][\gamma]$ is well-defined.

Note that in this argument we did not have to assume the homotopies f_s and g_s were constructed by applying the homotopy extension property. Thus we have proved the following result:

Proposition 4A.1. *There is a right action of* $\pi_1(X,x_0)$ *on* $\langle Z,X\rangle$ *defined by setting* $[f_0][\gamma]=[f_1]$ *whenever there exists a homotopy* $f_s:Z\to X$ *from* f_0 *to* f_1 *such that* $f_s(z_0)$ *is the loop* γ*, or any loop homotopic to* γ*.* □

It is easy to convert this right action into a left action, by defining $[\gamma][f_0]=[f_0][\gamma]^{-1}$. This just amounts to choosing the homotopy f_s so that $f_s(z_0)$ is the inverse path of γ.

When $Z=S^n$ this action reduces to the usual action of $\pi_1(X,x_0)$ on $\pi_n(X,x_0)$ since in the original definition of γf in terms of maps $(I^n,\partial I^n)\to(X,x_0)$, a homotopy from γf to f is obtained by restricting γf to smaller and smaller concentric cubes, and on the 'basepoint' ∂I^n this homotopy traces out the loop γ.

Proposition 4A.2. *If* (Z,z_0) *is a CW pair and* X *is a path-connected space, then the natural map* $\langle Z,X\rangle\to[Z,X]$ *induces a bijection of the orbit set* $\langle Z,X\rangle/\pi_1(X,x_0)$ *onto* $[Z,X]$*.*

In particular, this implies that $[Z,X]=\langle Z,X\rangle$ if X is simply-connected.

Proof: Since X is path-connected, every $f:Z\to X$ can be homotoped to take z_0 to the basepoint x_0, via homotopy extension, so the map $\langle Z,X\rangle\to[Z,X]$ is onto. If f_0 and f_1 are basepoint-preserving maps that are homotopic via the homotopy $f_s:Z\to X$, then by definition $[f_1]=[f_0][\gamma]$ for the loop $\gamma(s)=f_s(z_0)$, so $[f_0]$ and $[f_1]$ are in the same orbit under the action of $\pi_1(X,x_0)$. Conversely, two basepoint-preserving maps in the same orbit are obviously homotopic. □

Example 4A.3. If X is an H–space with identity element x_0, then the action of $\pi_1(X,x_0)$ on $\langle Z,X\rangle$ is trivial since for a map $f:(Z,z_0)\to(X,x_0)$ and a loop γ in X based at x_0, the multiplication in X defines a homotopy $f_s(z)=f(z)\gamma(s)$. This starts and ends with a map homotopic to f, and the loop $f_s(z_0)$ is homotopic to γ, both these homotopies being basepoint-preserving by the definition of an H–space.

The set of orbits of the π_1 action on π_n does not generally inherit a group structure from π_n. For example, when $n=1$ the orbits are just the conjugacy classes in π_1, and these form a group only when π_1 is abelian. Basepoints are thus a necessary technical device for producing the group structure in homotopy groups, though as we have shown, they can be ignored in simply-connected spaces.

For a set of maps $S^n\to X$ to generate $\pi_n(X)$ as a module over $\mathbb{Z}[\pi_1(X)]$ means that all elements of $\pi_n(X)$ can be represented by sums of these maps along arbitrary paths in X, where we allow reversing orientations to get negatives and repetitions to get arbitrary integer multiples. Examples of finite CW complexes X for which $\pi_n(X)$ is not finitely generated as a module over $\mathbb{Z}[\pi_1(X)]$ were given in Exercise 38 in §4.2, provided $n\ge 3$. Finding such an example for $n=2$ seems to be more difficult. The rest of this section will be devoted to a somewhat complicated construction which does this, and is interesting for other reasons as well.

An Example of Nonfinite Generation

We will construct a finite CW complex having π_n not finitely generated as a $\mathbb{Z}[\pi_1]$-module, for a given integer $n \geq 2$. The complex will be a subcomplex of a $K(\pi,1)$ having interesting homological properties: It is an $(n+1)$-dimensional CW complex with H_{n+1} nonfinitely generated, but its n-skeleton is finite so H_i is finitely generated for $i \leq n$ and π is finitely presented if $n > 1$. The first such example was found in [Stallings 1963] for $n = 2$. Our construction will be essentially the n-dimensional generalization of this, but described in a more geometric way as in [Bestvina & Brady 1997], which provides a general technique for constructing many examples of this sort.

To begin, let X be the product of n copies of $S^1 \vee S^1$. Since $S^1 \vee S^1$ is the 1-skeleton of the torus $T^2 = S^1 \times S^1$ in its usual CW structure, X can be regarded as a subcomplex of the $2n$-dimensional torus T^{2n}, the product of $2n$ circles. Define $f : T^{2n} \to S^1$ by $f(\theta_1, \cdots, \theta_{2n}) = \theta_1 + \cdots + \theta_{2n}$ where the coordinates $\theta_i \in S^1$ are viewed as angles measured in radians. The space $Z = X \cap f^{-1}(0)$ will provide the example we are looking for. As we shall see, Z is a finite CW complex of dimension $n-1$, with $\pi_{n-1}(Z)$ nonfinitely generated as a module over $\pi_1(Z)$ if $n \geq 3$. We will also see that $\pi_i(Z) = 0$ for $1 < i < n-1$.

The induced homomorphism $f_* : \pi_1(T^{2n}) \to \pi_1(S^1) = \mathbb{Z}$ sends each generator coming from an S^1 factor to 1. Let $\widetilde{T}^{2n} \to T^{2n}$ be the covering space corresponding to the kernel of f_*. This is a normal covering space since it corresponds to a normal subgroup, and the deck transformation group is $\mathbb{Z}$. The subcomplex of $\widetilde{T}^{2n}$ projecting to X is a normal covering space $\widetilde{X} \to X$ with the same group of deck transformations. Since $\pi_1(X)$ is the product of n free groups on two generators, $\widetilde{X}$ is the covering space of X corresponding to the kernel of the homomorphism $\pi_1(X) \to \mathbb{Z}$ sending each of the two generators of each free factor to 1. Since X is a $K(\pi,1)$, so is $\widetilde{X}$. For example, when $n = 1$, $\widetilde{X}$ is the union of two helices on the infinite cylinder $\widetilde{T}^2$:

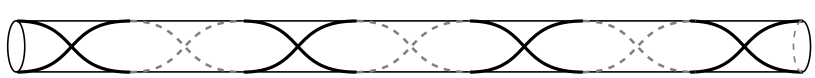

The map f lifts to a map $\widetilde{f} : \widetilde{T}^{2n} \to \mathbb{R}$, and Z lifts homeomorphically to a subspace $Z \subset \widetilde{X}$, namely $\widetilde{f}^{-1}(0) \cap \widetilde{X}$. We will show:

(∗) $\widetilde{X}$ is homotopy equivalent to a space Y obtained from Z by attaching an infinite sequence of n-cells.

Assuming this is true, it follows that $H_n(Y)$ is not finitely generated since in the exact sequence $H_n(Z) \to H_n(Y) \to H_n(Y,Z) \to H_{n-1}(Z)$ the first term is zero and the last term is finitely generated, Z being a finite CW complex of dimension $n-1$, while the third term is an infinite sum of $\mathbb{Z}$'s, one for each n-cell of Y. If $\pi_{n-1}(Z)$ were finitely generated as a $\pi_1(Z)$-module, then by attaching finitely many n-cells to Z we could make it $(n-1)$-connected since it is already $(n-2)$-connected as the $(n-1)$-skeleton of the $K(\pi,1)$ Y. Then by attaching cells of dimension greater than n we could build a $K(\pi,1)$ with finite n-skeleton. But this contradicts the fact that $H_n(Y)$ is not finitely generated.

To begin the verification of $(*)$, consider the torus T^m. The standard cell structure on T^m lifts to a cubical cell structure on the universal cover $\mathbb{R}^m$, with vertices the integer lattice points $\mathbb{Z}^m$. The function f lifts to a linear projection $L:\mathbb{R}^m \to \mathbb{R}$, $L(x_1,\cdots,x_m) = x_1 + \cdots + x_m$. The planes in $L^{-1}(\mathbb{Z})$ cut the cubes of $\mathbb{R}^m$ into convex polyhedra which we call *slabs*. There are m slabs in each m-dimensional cube. The boundary of a slab in $L^{-1}[i,i+1]$ consists of lateral faces that are slabs for lower-dimensional cubes, together with a lower face in $L^{-1}(i)$ and an upper face in $L^{-1}(i+1)$. In each cube there are two exceptional slabs whose lower or upper face degenerates to a point. These are the slabs containing the vertices of the cube where L has its maximum and minimum values. A slab deformation retracts onto the union of its lower and lateral faces, provided that the slab has an upper face that is not just a point. Slabs of the latter type are m-simplices, and we will refer to them as *cones* in what follows. These are the slabs containing the vertex of a cube on which L takes its maximal value. The lateral faces of a cone are also cones, of lower dimension.

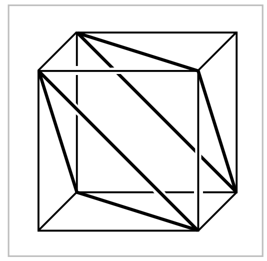

The slabs, together with all their lower-dimensional faces, give a CW structure on $\mathbb{R}^m$ with the planes of $L^{-1}(\mathbb{Z})$ as subcomplexes. These structures are preserved by the deck transformations of the cover $\mathbb{R}^m \to T^m$ so there is an induced CW structure in the quotient T^m, with $f^{-1}(0)$ as a subcomplex.

If X is any subcomplex of T^m in its original cubical cell structure, then the slab CW structure on T^m restricts to a CW structure on X. In particular, we obtain a CW structure on $Z = X \cap f^{-1}(0)$. Likewise we get a lifted CW structure on the cover $\widetilde{X} \subset \widetilde{T}^m$. Let $\widetilde{X}[i,j] = \widetilde{X} \cap \widetilde{f}^{-1}[i,j]$. The deformation retractions of noncone slabs onto their lateral and lower faces give rise to a deformation retraction of $\widetilde{X}[i,i+1]$ onto $\widetilde{X}[i] \cup C_i$ where C_i consists of all the cones in $\widetilde{X}[i,i+1]$. These cones are attached along their lower faces, and they all have the same vertex in $\widetilde{X}[i+1]$, so C_i is itself a cone in the usual sense, attached to $\widetilde{X}[i]$ along its base.

For the particular X we are interested in, we claim that each C_i is an n-disk attached along its boundary sphere. When $n = 1$ this is evident from the earlier picture of $\widetilde{X}$ as the union of two helices on a cylinder. For larger n we argue by induction. Passing from n to $n+1$ replaces X by two copies of $X \times S^1$ intersecting in X, one copy for each of the additional S^1 factors of T^{2n+2}. Replacing X by $X \times S^1$ changes C_i to its join with a point in the base of the new C_i. The two copies of this join then yield the suspension of C_i attached along the suspension of the base.

The same argument shows that $\widetilde{X}[-i-1,-i]$ deformation retracts onto $\widetilde{X}[-i]$ with an n-cell attached. We build the space Y and a homotopy equivalence $g:Y\to\widetilde{X}$ by an inductive procedure, starting with $Y_0 = Z$. Assuming that Y_i and a homotopy equivalence $g_i:Y_i\to\widetilde{X}[-i,i]$ have already been defined, we form Y_{i+1} by attaching two n-cells by the maps obtained from the attaching maps of the two n-cells in $\widetilde{X}[-i-1,i+1]-\widetilde{X}[-i,i]$ by composing with a homotopy inverse to g_i. This allows g_i to be extended to a homotopy equivalence $g_{i+1}:Y_{i+1}\to\widetilde{X}[-i-1,i+1]$. Taking the union over i gives $g:Y\to\widetilde{X}$. One can check this is a homotopy equivalence by seeing that it induces isomorphisms on all homotopy groups, using the standard compactness argument. This finishes the verification of $(*)$.

It is interesting to see what the complex Z looks like in the case $n=3$, when Z is 2-dimensional and has π_2 nonfinitely generated over $\mathbb{Z}[\pi_1(Z)]$. In this case X is the product of three $S^1\vee S^1$'s, so X is the union of the eight 3-tori obtained by choosing one of the two S^1 summands in each $S^1\vee S^1$ factor. We denote these 3-tori $S^1_\pm\times S^1_\pm\times S^1_\pm$. Viewing each of these 3-tori as the cube in the previous figure with opposite faces identified, we see that Z is the union of the eight 2-tori formed by the two sloping triangles in each cube. Two of these 2-tori intersect along a circle when the corresponding 3-tori of X intersect along a 2-torus. This happens when the triples of $\pm$'s for the two 3-tori differ in exactly one entry. The pattern of intersection of the eight 2-tori of Z can thus be described combinatorially via the 1-skeleton of the cube, with vertices $(\pm1,\pm1,\pm1)$. There is a torus of Z for each vertex of the cube, and two tori intersect along a circle when the corresponding vertices of the cube are the endpoints of an edge of the cube. All eight tori contain the single 0-cell of Z.

To obtain a model of Z itself, consider a regular octahedron inscribed in the cube with vertices $(\pm1,\pm1,\pm1)$. If we identify each pair of opposite edges of the octahedron, each pair of opposite triangular faces becomes a torus. However, there are only four pairs of opposite faces, so we get only four tori this way, not eight. To correct this problem, regard each triangular face of the octahedron as two copies of the same triangle, distinguished from each other by a choice of normal direction, an arrow attached to the triangle pointing either inside the octahedron or outside it, that is, either toward the nearest vertex of the surrounding cube or toward the opposite vertex of the cube. Then each pair of opposite triangles of the octahedron having normal vectors pointing toward the same vertex of the cube determines a torus, when opposite edges are identified as before. Each edge of the original octahedron is also replaced by two edges oriented either toward the interior or exterior of the octahedron. The vertices of the octahedron may be left unduplicated since they will all be identified to a single point anyway. With this scheme, the two tori corresponding to the vertices at the ends

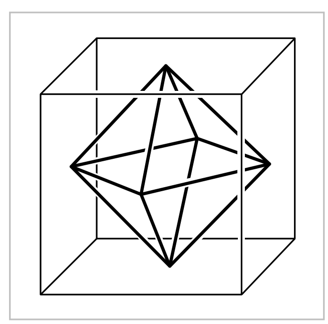

of an edge of the cube then intersect along a circle, as they should, and other pairs of tori intersect only at the 0-cell of Z.

This model of Z has the advantage of displaying the symmetry group of the cube, a group of order 48, as a symmetry group of Z, corresponding to the symmetries of X permuting the three $S^1 \vee S^1$ factors and the two S^1's of each $S^1 \vee S^1$. Undoubtedly Z would be very pretty to look at if we lived in a space with enough dimensions to see all of it at one glance.

It might be interesting to see an explicit set of maps $S^2 \to Z$ generating $\pi_2(Z)$ as a $\mathbb{Z}[\pi_1]$-module. One might also ask whether there are simpler examples of these nonfinite generation phenomena.

Exercises

1. Show directly that if X is a topological group with identity element x_0, then any two maps $f, g: (Z, z_0) \to (X, x_0)$ which are homotopic are homotopic through basepoint-preserving maps.

2. Show that under the map $\langle X, Y\rangle \to \operatorname{Hom}(\pi_n(X, x_0), \pi_n(Y, y_0))$, $[f] \mapsto f_*$, the action of $\pi_1(Y, y_0)$ on $\langle X, Y\rangle$ corresponds to composing with the action on $\pi_n(Y, y_0)$, that is, $(\gamma f)_* = \beta_\gamma f_*$. Deduce a bijection of $[X, K(\pi,1)]$ with the set of orbits of $\operatorname{Hom}(\pi_1(X), \pi)$ under composition with inner automorphisms of π. In particular, if π is abelian then $[X, K(\pi,1)] = \langle X, K(\pi,1)\rangle = \operatorname{Hom}(\pi_1(X), \pi)$.

3. For a space X let $\operatorname{Aut}(X)$ denote the group of homotopy classes of homotopy equivalences $X \to X$. Show that for a CW complex $K(\pi,1)$, $\operatorname{Aut}(K(\pi,1))$ is isomorphic to the group of outer automorphisms of π, that is, automorphisms modulo inner automorphisms.

4. With the notation of the preceding problem, show that $\operatorname{Aut}(\bigvee_n S^k) \approx GL_n(\mathbb{Z})$ for $k > 1$, where $\bigvee_n S^k$ denotes the wedge sum of n copies of S^k and $GL_n(\mathbb{Z})$ is the group of $n \times n$ matrices with entries in $\mathbb{Z}$ having an inverse matrix of the same form. [$GL_n(\mathbb{Z})$ is the automorphism group of $\mathbb{Z}^n \approx \pi_k(\bigvee_n S^k) \approx H_k(\bigvee_n S^k)$.]

5. This problem involves the spaces constructed in the latter part of this section.
(a) Compute the homology groups of the complex Z in the case $n = 3$, when Z is 2-dimensional.
(b) Letting $\widetilde{X}_n$ denote the n-dimensional complex $\widetilde{X}$, show that $\widetilde{X}_n$ can be obtained inductively from $\widetilde{X}_{n-1}$ as the union of two copies of the mapping torus of the generating deck transformation $\widetilde{X}_{n-1} \to \widetilde{X}_{n-1}$, with copies of $\widetilde{X}_{n-1}$ in these two mapping tori identified. Thus there is a fiber bundle $\widetilde{X}_n \to S^1 \vee S^1$ with fiber $\widetilde{X}_{n-1}$.
(c) Use part (b) to find a presentation for $\pi_1(\widetilde{X}_n)$, and show this presentation reduces to a finite presentation if $n > 2$ and a presentation with a finite number of generators if $n = 2$. In the latter case, deduce that $\pi_1(\widetilde{X}_2)$ has no finite presentation from the fact that $H_2(\widetilde{X}_2)$ is not finitely generated.

4.B The Hopf Invariant

In §2.2 we used homology to distinguish different homotopy classes of maps $S^n \to S^n$ via the notion of degree. We will show here that cup product can be used to do something similar for maps $S^{2n-1} \to S^n$. Originally this was done by Hopf using more geometric constructions, before the invention of cohomology and cup products.

In general, given a map $f : S^m \to S^n$ with $m \geq n$, we can form a CW complex C_f by attaching a cell e^{m+1} to S^n via f. The homotopy type of C_f depends only on the homotopy class of f, by Proposition 0.18. Thus for maps $f, g : S^m \to S^n$, any invariant of homotopy type that distinguishes C_f from C_g will show that f is not homotopic to g. For example, if $m = n$ and f has degree d, then from the cellular chain complex of C_f we see that $H_n(C_f) \approx \mathbb{Z}_{|d|}$, so the homology of C_f detects the degree of f, up to sign. When $m > n$, however, the homology of C_f consists of $\mathbb{Z}$'s in dimensions 0, n, and $m+1$, independent of f. The same is true of cohomology groups, but cup products have a chance of being nontrivial in $H^*(C_f)$ when $m = 2n-1$. In this case, if we choose generators $\alpha \in H^n(C_f)$ and $\beta \in H^{2n}(C_f)$, then the multiplicative structure of $H^*(C_f)$ is determined by a relation $\alpha^2 = H(f)\beta$ for an integer $H(f)$ called the **Hopf invariant** of f. The sign of $H(f)$ depends on the choice of the generator β, but this can be specified by requiring β to correspond to a fixed generator of $H^{2n}(D^{2n}, \partial D^{2n})$ under the map $H^{2n}(C_f) \approx H^{2n}(C_f, S^n) \to H^{2n}(D^{2n}, \partial D^{2n})$ induced by the characteristic map of the cell e^{2n}, which is determined by f. We can then change the sign of $H(f)$ by composing f with a reflection of S^{2n-1}, of degree -1. If $f \simeq g$, then under the homotopy equivalence $C_f \simeq C_g$ the chosen generators β for $H^{2n}(C_f)$ and $H^{2n}(C_g)$ correspond, so $H(f)$ depends only on the homotopy class of f.

If f is a constant map then $C_f = S^n \vee S^{2n}$ and $H(f) = 0$ since C_f retracts onto S^n. Also, $H(f)$ is always zero for odd n since in this case $\alpha^2 = -\alpha^2$ by the commutativity property of cup product, hence $\alpha^2 = 0$.

Three basic examples of maps with nonzero Hopf invariant are the maps defining the three Hopf bundles in Examples 4.45, 4.46, and 4.47. The first of these Hopf maps is the attaching map $f : S^3 \to S^2$ for the 4-cell of CP^2. This has $H(f) = 1$ since $H^*(\mathbb{C}\mathrm{P}^2; \mathbb{Z}) \approx \mathbb{Z}[\alpha]/(\alpha^3)$ by Theorem 3.19. Similarly, $\mathbb{H}\mathrm{P}^2$ gives rise to a map $S^7 \to S^4$ of Hopf invariant 1. In the case of the octonionic projective plane $\mathbb{O}\mathrm{P}^2$, which is built from the map $S^{15} \to S^8$ defined in Example 4.47, we can deduce that $H^*(\mathbb{O}\mathrm{P}^2; \mathbb{Z}) \approx \mathbb{Z}[\alpha]/(\alpha^3)$ either from Poincaré duality as in Example 3.40 or from Exercise 5 for §4.D.

It is a fundamental theorem of [Adams 1960] that a map $f : S^{2n-1} \to S^n$ of Hopf invariant 1 exists only when $n = 2, 4, 8$. This has a number of very interesting consequences, for example:

- $\mathbb{R}^n$ is a division algebra only for $n = 1, 2, 4, 8$.
- S^n is an H-space only for $n = 0, 1, 3, 7$.
- S^n has n linearly independent tangent vector fields only for $n = 0, 1, 3, 7$.
- The only fiber bundles $S^p \to S^q \to S^r$ occur when $(p,q,r) = (0,1,1)$, $(1,3,2)$, $(3,7,4)$, and $(7,15,8)$.

The first and third assertions were in fact proved shortly before Adams' theorem in [Kervaire 1958] and [Milnor 1958] as applications of a theorem of Bott that $\pi_{2n}U(n) \approx \mathbb{Z}_{n!}$. A full discussion of all this, and a proof of Adams' theorem, is given in [VBKT].

Though maps of Hopf invariant 1 are rare, there are maps $S^{2n-1} \to S^n$ of Hopf invariant 2 for all even n. Namely, consider the space $J_2(S^n)$ constructed in §3.2. This has a CW structure with three cells, of dimensions 0, n, and $2n$, so $J_2(S^n)$ has the form C_f for some $f: S^{2n-1} \to S^n$. We showed that if n is even, the square of a generator of $H^n(J_2(S^n);\mathbb{Z})$ is twice a generator of $H^{2n}(J_2(S^n);\mathbb{Z})$, so $H(f) = \pm 2$.

From this example we can get maps of any even Hopf invariant when n is even via the following fact.

Proposition 4B.1. *The Hopf invariant $H: \pi_{2n-1}(S^n) \to \mathbb{Z}$ is a homomorphism.*

Proof: For $f,g: S^{2n-1} \to S^n$, let us compare C_{f+g} with the space $C_{f\vee g}$ obtained from S^n by attaching two $2n$-cells via f and g. There is a natural quotient map $q: C_{f+g} \to C_{f\vee g}$ collapsing the equatorial disk of the $2n$-cell of C_{f+g} to a point. The induced cellular chain map q_* sends e^{2n}_{f+g} to $e^{2n}_f + e^{2n}_g$. In cohomology this implies that $q^*(\beta_f) = q^*(\beta_g) = \beta_{f+g}$ where β_f, β_g, and β_{f+g} are the cohomology classes dual to the $2n$-cells. Letting α_{f+g} and $\alpha_{f\vee g}$ be the cohomology classes corresponding to the n-cells, we have $q^*(\alpha_{f\vee g}) = \alpha_{f+g}$ since q is a homeomorphism on the n-cells. By restricting to the subspaces C_f and C_g of $C_{f\vee g}$ we see that $\alpha^2_{f\vee g} = H(f)\beta_f + H(g)\beta_g$. Thus $\alpha^2_{f+g} = q^*(\alpha^2_{f\vee g}) = H(f)q^*(\beta_f) + H(g)q^*(\beta_g) = \big(H(f) + H(g)\big)\beta_{f+g}$. ☐

Corollary 4B.2. *$\pi_{2n-1}(S^n)$ contains a $\mathbb{Z}$ direct summand when n is even.*

Proof: Either H or $H/2$ is a surjective homomorphism $\pi_{2n-1}(S^n) \to \mathbb{Z}$. ☐

Exercises

1. Show that the Hopf invariant of a composition $S^{2n-1} \xrightarrow{f} S^{2n-1} \xrightarrow{g} S^n$ is given by $H(gf) = (\deg f)H(g)$, and for a composition $S^{2n-1} \xrightarrow{f} S^n \xrightarrow{g} S^n$ the Hopf invariant satisfies $H(gf) = (\deg g)^2 H(f)$.

2. Show that if $S^k \longrightarrow S^m \xrightarrow{p} S^n$ is a fiber bundle, then $m = 2n-1$, $k = n-1$, and, when $n > 1$, $H(p) = \pm 1$. [Show that C_p is a manifold and apply Poincaré duality.]

4.C Minimal Cell Structures

We can apply the homology version of Whitehead's theorem, Corollary 4.33, to show that a simply-connected CW complex with finitely generated homology groups is always homotopy equivalent to a CW complex having the minimum number of cells consistent with its homology, namely, one n-cell for each $\mathbb{Z}$ summand of H_n and a pair of cells of dimension n and $n+1$ for each $\mathbb{Z}_k$ summand of H_n.

Proposition 4C.1. *Given a simply-connected CW complex X and a decomposition of each of its homology groups $H_n(X)$ as a direct sum of cyclic groups with specified generators, then there is a CW complex Z and a cellular homotopy equivalence $f:Z\to X$ such that each cell of Z is either:*

(a) *a 'generator' n-cell e^n_α, which is a cycle in cellular homology mapped by f to a cellular cycle representing the specified generator α of one of the cyclic summands of $H_n(X)$; or*

(b) *a 'relator' $(n+1)$-cell e^{n+1}_α, with cellular boundary equal to a multiple of the generator n-cell e^n_α, in the case that α has finite order.*

In the nonsimply-connected case this result can easily be false, counterexamples being provided by acyclic spaces and the space $X=(S^1\vee S^n)\cup e^{n+1}$ constructed in Example 4.35, which has the same homology as S^1 but which must have cells of dimension greater than 1 in order to have π_n nontrivial.

Proof: We build Z inductively over skeleta, starting with Z^1 a point since X is simply-connected. For the inductive step, suppose we have constructed $f:Z^n\to X$ inducing an isomorphism on H_i for $i<n$ and a surjection on H_n. For the mapping cylinder M_f we then have $H_i(M_f,Z^n)=0$ for $i\le n$ and $H_{n+1}(M_f,Z^n)\approx\pi_{n+1}(M_f,Z^n)$ by the Hurewicz theorem. To construct Z^{n+1} we use the following diagram:

$$\begin{array}{ccccccccc}
H_{n+1}(X) & & \pi_{n+1}(M_f,Z^n) & & & & H_n(X) & & \\
\wr & & \wr & & & & \wr & & \\
H_{n+1}(M_f) & \longrightarrow & H_{n+1}(M_f,Z^n) & \longrightarrow & H_n(Z^n) & \longrightarrow & H_n(M_f) & \longrightarrow & 0 \\
\uparrow & & \uparrow & & \| & & \uparrow & & \\
H_{n+1}(Z^{n+1}) & \longrightarrow & H_{n+1}(Z^{n+1},Z^n) & \longrightarrow & H_n(Z^n) & \longrightarrow & H_n(Z^{n+1}) & \longrightarrow & 0
\end{array}$$

By induction we know the map $H_n(Z^n)\to H_n(M_f)\approx H_n(X)$ exactly, namely, Z^n has generator n-cells, which are cellular cycles mapping to the given generators of $H_n(X)$, along with relator n-cells that do not contribute to $H_n(Z^n)$. Thus $H_n(Z^n)$ is free with basis the generator n-cells, and the kernel of $H_n(Z^n)\to H_n(X)$ is free with basis given by certain multiples of some of the generator n-cells. Choose 'relator' elements ρ_i in $H_{n+1}(M_f,Z^n)$ mapping to this basis for the kernel, and let the 'generator' elements $\gamma_i\in H_{n+1}(M_f,Z^n)$ be the images of the chosen generators of $H_{n+1}(M_f)\approx H_{n+1}(X)$.

Via the Hurewicz isomorphism $H_{n+1}(M_f,Z^n)\approx\pi_{n+1}(M_f,Z^n)$, the homology classes ρ_i and γ_i are represented by maps $r_i,g_i:(D^{n+1},S^n)\to(M_f,Z^n)$. We form

Z^{n+1} from Z^n by attaching $(n+1)$-cells via the restrictions of the maps r_i and g_i to S^n. The maps r_i and g_i themselves then give an extension of the inclusion $Z^n \hookrightarrow M_f$ to a map $Z^{n+1} \to M_f$, whose composition with the retraction $M_f \to X$ is the extended map $f: Z^{n+1} \to X$. This gives us the lower row of the preceding diagram, with commutative squares. By construction, the subgroup of $H_{n+1}(Z^{n+1}, Z^n)$ generated by the relator $(n+1)$-cells maps injectively to $H_n(Z^n)$, with image the kernel of $H_n(Z^n) \to H_n(X)$, so $f_*: H_n(Z^{n+1}) \to H_n(X)$ is an isomorphism. The elements of $H_{n+1}(Z^{n+1}, Z^n)$ represented by the generator $(n+1)$-cells map to the y_i's, hence map to zero in $H_n(Z^n)$, so by exactness of the second row these generator $(n+1)$-cells are cellular cycles representing elements of $H_{n+1}(Z^{n+1})$ mapped by f_* to the given generators of $H_{n+1}(X)$. In particular, $f_*: H_{n+1}(Z^{n+1}) \to H_{n+1}(X)$ is surjective, and the induction step is finished.

Doing this for all n, we produce a CW complex Z and a map $f: Z \to X$ with the desired properties. □

Example 4C.2. Suppose X is a simply-connected CW complex such that for some $n \geq 2$, the only nonzero reduced homology groups of X are $H_n(X)$, which is finitely generated, and $H_{n+1}(X)$, which is finitely generated and free. Then the proposition says that X is homotopy equivalent to a CW complex Z obtained from a wedge sum of n-spheres by attaching $(n+1)$-cells. The attaching maps of these cells are determined up to homotopy by the cellular boundary map $H_{n+1}(Z^{n+1}, Z^n) \to H_n(Z^n)$ since $\pi_n(Z^n) \approx H_n(Z^n)$. So the attaching maps are either trivial, in the case of generator $(n+1)$-cells, or they represent some multiple of an inclusion of one of the wedge summands, in the case of a relator $(n+1)$-cell. Hence Z is the wedge sum of spheres S^n and S^{n+1} together with Moore spaces $M(\mathbb{Z}_m, n)$ of the form $S^n \cup e^{n+1}$. In particular, the homotopy type of X is uniquely determined by its homology groups.

Proposition 4C.3. *Let X be a simply-connected space homotopy equivalent to a CW complex, such that the only nontrivial reduced homology groups of X are $H_2(X) \approx \mathbb{Z}^m$ and $H_4(X) \approx \mathbb{Z}$. Then the homotopy type of X is uniquely determined by the cup product ring $H^*(X;\mathbb{Z})$. In particular, this applies to any simply-connected closed 4-manifold.*

Proof: By the previous proposition we may assume X is a complex X_φ obtained from a wedge sum $\bigvee_j S_j^2$ of m 2-spheres S_j^2 by attaching a cell e^4 via a map $\varphi: S^3 \to \bigvee_j S_j^2$. As shown in Example 4.52, $\pi_3(\bigvee_j S_j^2)$ is free with basis the Hopf maps $\eta_j: S^3 \to S_j^2$ and the Whitehead products $[i_j, i_k]$, $j < k$, where i_j is the inclusion $S_j^2 \hookrightarrow \bigvee_j S_j^2$. Since a homotopy of φ does not change the homotopy type of X_φ, we may assume φ is a linear combination $\sum_j a_j\eta_j + \sum_{j<k} a_{jk}[i_j, i_k]$. We need to see how the coefficients a_j and a_{jk} determine the cup product $H^2(X;\mathbb{Z}) \times H^2(X;\mathbb{Z}) \to H^4(X;\mathbb{Z})$.

This cup product can be represented by an $m \times m$ symmetric matrix (b_{jk}) where the cup product of the cohomology classes dual to the j^{th} and k^{th} 2-cells is b_{jk}

times the class dual to the 4-cell. We claim that $b_{jk} = a_{jk}$ for $j < k$ and $b_{jj} = a_j$. If φ is one of the generators η_i or $[i_j, i_k]$ this is clear, since if $\varphi = \eta_j$ then X_φ is the wedge sum of $\mathbb{C}P^2$ with $m-1$ 2-spheres, while if $\varphi = [i_j, i_k]$ then X_φ is the wedge sum of $S_j^2 \times S_k^2$ with $m-2$ 2-spheres. The claim is also true when φ is $-\eta_j$ or $-[i_j, i_k]$ since changing the sign of φ amounts to composing φ with a reflection of S^3, and this changes the generator of $H^4(X_\varphi;\mathbb{Z})$ to its negative. The general case now follows by induction from the assertion that the matrix (b_{jk}) for $X_{\varphi+\psi}$ is the sum of the corresponding matrices for X_φ and X_ψ. This assertion can be proved as follows. By attaching two 4-cells to $\bigvee_j S_j^2$ by φ and ψ we obtain a complex $X_{\varphi,\psi}$ which we can view as $X_\varphi \cup X_\psi$. There is a quotient map $q : X_{\varphi+\psi} \to X_{\varphi,\psi}$ that is a homeomorphism on the 2-skeleton and collapses the closure of an equatorial 3-disk in the 4-cell of $X_{\varphi+\psi}$ to a point. The induced map $q^* : H^4(X_{\varphi,\psi}) \to H^4(X_{\varphi+\psi})$ sends each of the two generators corresponding to the 4-cells of $X_{\varphi,\psi}$ to a generator, and the assertion follows.

Now suppose X_φ and X_ψ have isomorphic cup product rings. This means bases for $H^*(X_\varphi;\mathbb{Z})$ and $H^*(X_\psi;\mathbb{Z})$ can be chosen so that the matrices specifying the cup product $H^2 \times H^2 \to H^4$ with respect to these bases are the same. The preceding proposition says that any choice of basis can be realized as the dual basis to a cell structure on a CW complex homotopy equivalent to the given complex. Therefore we may assume the matrices (b_{jk}) for X_φ and X_ψ are the same. By what we have shown in the preceding paragraph, this means φ and ψ are homotopic, hence X_φ and X_ψ are homotopy equivalent.

For the statement about simply-connected closed 4-manifolds, Corollaries A.8 and A.9 and Proposition A.11 in the Appendix say that such a manifold M has the homotopy type of a CW complex with finitely generated homology groups. Then Poincaré duality and the universal coefficient theorem imply that the only nontrivial homology groups $H_i(M)$ are $\mathbb{Z}$ for $i = 0, 4$ and $\mathbb{Z}^m$ for $i = 2$, for some $m \geq 0$. □

This result and the example preceding it are special cases of a homotopy classification by Whitehead of simply-connected CW complexes with positive-dimensional cells in three adjacent dimensions n, $n+1$, and $n+2$; see [Baues 1996] for a full treatment of this.

4.D Cohomology of Fiber Bundles

While the homotopy groups of the three spaces in a fiber bundle fit into a long exact sequence, the relation between their homology or cohomology groups is much more complicated. The Künneth formula shows that there are some subtleties even for a product bundle, and for general bundles the machinery of spectral sequences, developed in [SSAT], is required. In this section we will describe a few special sorts of fiber bundles where more elementary techniques suffice. As applications we calculate the cohomology rings of some important spaces closely related to Lie groups. In particular we find a number of spaces with exterior and polynomial cohomology rings.

The Leray–Hirsch Theorem

This theorem will be the basis for all the other results in this section. It gives hypotheses sufficient to guarantee that a fiber bundle has cohomology very much like that of a product bundle.

Theorem 4D.1. *Let $F \to E \xrightarrow{p} B$ be a fiber bundle such that, for some commutative coefficient ring R:*

(a) *$H^n(F;R)$ is a finitely generated free R-module for each n.*

(b) *There exist classes $c_j \in H^{k_j}(E;R)$ whose restrictions $i^*(c_j)$ form a basis for $H^*(F;R)$ in each fiber F, where $i:F \to E$ is the inclusion.*

Then the map $\Phi : H^*(B;R) \otimes_R H^*(F;R) \to H^*(E;R)$, $\sum_{ij} b_i \otimes i^*(c_j) \mapsto \sum_{ij} p^*(b_i) \smile c_j$, is an isomorphism of R-modules.

The conclusion can be restated as saying that $H^*(E;R)$ is a free $H^*(B;R)$-module with basis $\{c_j\}$, where we view $H^*(E;R)$ as a module over the ring $H^*(B;R)$ by defining scalar multiplication by $bc = p^*(b) \smile c$ for $b \in H^*(B;R)$ and $c \in H^*(E;R)$.

In the case of a product $E = B \times F$ with $H^*(F;R)$ free over R, we can pull back a basis for $H^*(F;R)$ via the projection $E \to F$ to obtain the classes c_j. Thus the Leray–Hirsch theorem generalizes the version of the Künneth formula involving cup products, Theorem 3.15, at least as far as the additive structure and the module structure over $H^*(B;R)$ are concerned. However, the Leray–Hirsch theorem does not assert that the isomorphism $H^*(E;R) \approx H^*(B;R) \otimes_R H^*(F;R)$ is a ring isomorphism, and in fact this need not be true, for example for the Klein bottle viewed as a bundle with fiber and base S^1, where the Leray–Hirsch theorem applies with $\mathbb{Z}_2$ coefficients.

An example of a bundle where the classes c_j do not exist is the Hopf bundle $S^1 \to S^3 \to S^2$, since $H^*(S^3) \not\approx H^*(S^2) \otimes H^*(S^1)$.

Proof: We first prove the result for finite-dimensional CW complexes B by induction on their dimension. The case that B is 0-dimensional is trivial. For the induction step, suppose B has dimension n, and let $B' \subset B$ be the subspace obtained by deleting a point x_α from the interior of each n-cell e^n_α of B. Let $E' = p^{-1}(B')$. Then we have a commutative diagram, with coefficients in R understood:

$$\begin{array}{ccccccccc}
\cdots \longrightarrow & H^*(B,B') \otimes_R H^*(F) & \longrightarrow & H^*(B) \otimes_R H^*(F) & \longrightarrow & H^*(B') \otimes_R H^*(F) & \longrightarrow \cdots \\
& \big\downarrow \Phi & & \big\downarrow \Phi & & \big\downarrow \Phi & \\
\cdots \longrightarrow & H^*(E,E') & \longrightarrow & H^*(E) & \longrightarrow & H^*(E') & \longrightarrow \cdots
\end{array}$$

The map Φ on the left is defined exactly as in the absolute case, using the relative cup product $H^*(E,E') \otimes_R H^*(E) \to H^*(E,E')$. The first row of the diagram is exact since tensoring with a free module preserves exactness. The second row is of course exact also. The commutativity of the diagram follows from the evident naturality of Φ in the case of the two squares shown. For the other square involving coboundary maps, if we start with an element $b \otimes i^*(c_j) \in H^*(B') \otimes_R H^*(F)$ and map this horizontally we get $\delta b \otimes i^*(c_j)$ which maps vertically to $p^*(\delta b) \smile c_j$, whereas if we first map vertically we get $p^*(b) \smile c_j$ which maps horizontally to $\delta(p^*(b) \smile c_j) = \delta p^*(b) \smile c_j = p^*(\delta b) \smile c_j$ since $\delta c_j = 0$.

The space B' deformation retracts onto the skeleton B^{n-1}, and the following lemma implies that the inclusion $p^{-1}(B^{n-1}) \hookrightarrow E'$ is a weak homotopy equivalence, hence induces an isomorphism on all cohomology groups:

Lemma 4D.2. *Given a fiber bundle $p : E \to B$ and a subspace $A \subset B$ such that (B, A) is k-connected, then $(E, p^{-1}(A))$ is also k-connected.*

Proof: For a map $g : (D^i, \partial D^i) \to (E, p^{-1}(A))$ with $i \le k$, there is by hypothesis a homotopy $f_t : (D^i, \partial D^i) \to (B, A)$ of $f_0 = pg$ to a map f_1 with image in A. The homotopy lifting property then gives a homotopy $g_t : (D^i, \partial D^i) \to (E, p^{-1}(A))$ of g to a map with image in $p^{-1}(A)$. □

The theorem for finite-dimensional B will now follow by induction on n and the five-lemma once we show that the left-hand Φ in the diagram is an isomorphism.

By the fiber bundle property there are open disk neighborhoods $U_\alpha \subset e^n_\alpha$ of the points x_α such that the bundle is a product over each U_α. Let $U = \bigcup_\alpha U_\alpha$ and let $U' = U \cap B'$. By excision we have $H^*(B, B') \approx H^*(U, U')$, and $H^*(E, E') \approx H^*(p^{-1}(U), p^{-1}(U'))$. This gives a reduction to the problem of showing that the map $\Phi : H^*(U, U') \otimes_R H^*(F) \to H^*(U \times F, U' \times F)$ is an isomorphism. For this we can either appeal to the relative Künneth formula in Theorem 3.18 or we can argue again by induction, applying the five-lemma to the diagram with (B, B') replaced by (U, U'), induction implying that the theorem holds for U and U' since they deformation retract onto complexes of dimensions 0 and $n-1$, respectively, and by the lemma we can restrict to the bundles over these complexes.

Next there is the case that B is an infinite-dimensional CW complex. Since (B, B^n) is n-connected, the lemma implies that the same is true of $(E, p^{-1}(B^n))$. Hence in the commutative diagram at the right the horizontal maps are isomorphisms below dimension n. Then the fact that the right-hand Φ is an isomorphism, as we have already shown, implies that the left-hand Φ is an isomorphism below dimension n. Since n is arbitrary, this gives the theorem for all CW complexes B.

$$\begin{array}{ccc} H^*(B) \otimes_R H^*(F) & \longrightarrow & H^*(B^n) \otimes_R H^*(F) \\ \big\downarrow \Phi & & \big\downarrow \Phi \\ H^*(E) & \longrightarrow & H^*(p^{-1}(B^n)) \end{array}$$

To extend to the case of arbitrary base spaces B we need the notion of a **pullback bundle** which is used quite frequently in bundle theory. Given a fiber bundle

$p:E \to B$ and a map $f:A \to B$, let $f^*(E) = \{(a,e) \in A \times E \mid f(a) = p(e)\}$, so there is a commutative diagram as at the right, where the two maps from $f^*(E)$ are $(a,e) \mapsto a$ and $(a,e) \mapsto e$. It is a simple exercise to verify that the projection $f^*(E) \to A$ is a fiber bundle with the same fiber as $E \to B$, since a local trivialization of $E \to B$ over $U \subset B$ gives rise to a local trivialization of $f^*(E) \to A$ over $f^{-1}(U)$.

$$\begin{array}{ccc} f^*(E) & \longrightarrow & E \\ \downarrow & & \downarrow p \\ A & \xrightarrow{\ f\ } & B \end{array}$$

If $f:A \to B$ is a CW approximation to an arbitrary base space B, then $f^*(E) \to E$ induces an isomorphism on homotopy groups by the five-lemma applied to the long exact sequences of homotopy groups for the two bundles $E \to B$ and $f^*(E) \to A$ with fiber F. Hence $f^*(E) \to E$ is also an isomorphism on cohomology. The classes c_j pull back to classes in $H^*(f^*(E);R)$ which still restrict to a basis in each fiber, and so the naturality of Φ reduces the theorem for $E \to B$ to the case of $f^*(E) \to A$. □

Corollary 4D.3. (a) $H^*(U(n);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_1, x_3, \cdots, x_{2n-1}]$, *the exterior algebra on generators* x_i *of odd dimension* i.
(b) $H^*(SU(n);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_3, x_5, \cdots, x_{2n-1}]$.
(c) $H^*(Sp(n);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_3, x_7, \cdots, x_{4n-1}]$.

These are ring isomorphisms, and the proof will involve bundles where the isomorphism in the Leray–Hirsch theorem happens to be an isomorphism of rings.

Proof: For (a), assume inductively that the result holds for $U(n-1)$. From the bundle $U(n-1) \to U(n) \to S^{2n-1}$ we see that the pair $(U(n), U(n-1))$ is $(2n-2)$-connected, so $H^i(U(n);\mathbb{Z}) \to H^i(U(n-1);\mathbb{Z})$ is an isomorphism for $i \le 2n-3$ and the classes $x_1, \cdots, x_{2n-3} \in H^*(U(n-1);\mathbb{Z})$ given by induction are the restrictions of classes $c_1, \cdots, c_{2n-3} \in H^*(U(n);\mathbb{Z})$. The products of distinct x_i's form an additive basis for $H^*(U(n-1);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_1, \cdots, x_{2n-3}]$, and these products are restrictions of the corresponding products of c_i's, so the Leray–Hirsch theorem applies to give an additive basis for $H^*(U(n);\mathbb{Z})$ consisting of all products of distinct elements $x_1 = c_1, \cdots, x_{2n-3} = c_{2n-3}$ and a new generator x_{2n-1} coming from $H^{2n-1}(S^{2n-1};\mathbb{Z})$. By commutativity of cup product this is the exterior algebra $\Lambda_{\mathbb{Z}}[x_1, \cdots, x_{2n-1}]$.

The same proof works for $Sp(n)$ using the bundle $Sp(n-1) \to Sp(n) \to S^{4n-1}$. In the case of $SU(n)$ one uses the bundle $SU(n-1) \to SU(n) \to S^{2n-1}$. Since $SU(1)$ is the trivial group, the bundle $SU(1) \to SU(2) \to S^3$ shows that $SU(2) = S^3$, so the first generator is x_3. □

It is illuminating to look more closely at how the homology and cohomology of $O(n)$, $U(n)$, and $Sp(n)$ are related to their bundle structures. For $U(n)$ one has the sequence of bundles

$$\begin{array}{ccccccccccc} S^1 = U(1) & \hookrightarrow & U(2) & \hookrightarrow & U(3) & \hookrightarrow & \cdots & \hookrightarrow & U(n-1) & \hookrightarrow & U(n) \\ & & \downarrow & & \downarrow & & & & \downarrow & & \downarrow \\ & & S^3 & & S^5 & & \cdots & & S^{2n-3} & & S^{2n-1} \end{array}$$

If all these were product bundles, $U(n)$ would be homeomorphic to the product $S^1\times S^3\times\cdots\times S^{2n-1}$. In actuality the bundles are nontrivial, but the homology and cohomology of $U(n)$ are the same as for this product of spheres, including the cup product structure. For $Sp(n)$ the situation is quite similar, with the corresponding product of spheres $S^3\times S^7\times\cdots\times S^{4n-1}$. For $O(n)$ the corresponding sequence of bundles is

$$\begin{array}{ccccccccccc} S^0 = O(1) & \hookrightarrow & O(2) & \hookrightarrow & O(3) & \hookrightarrow & \cdots & \hookrightarrow & O(n-1) & \hookrightarrow & O(n) \\ & & \downarrow & & \downarrow & & & & \downarrow & & \downarrow \\ & & S^1 & & S^2 & & \cdots & & S^{n-2} & & S^{n-1} \end{array}$$

The calculations in §3.D show that $H_*(O(n);\mathbb{Z}_2)\approx H_*(S^0\times S^1\times\cdots\times S^{n-1};\mathbb{Z}_2)$, but with $\mathbb{Z}$ coefficients this no longer holds. Instead, consider the coarser sequence of bundles

$$\begin{array}{ccccccccccc} S^0 = O(1) & \hookrightarrow & O(3) & \hookrightarrow & O(5) & \hookrightarrow & \cdots & \hookrightarrow & O(2k-1) & \hookrightarrow & O(2k) \\ & & \downarrow & & \downarrow & & & & \downarrow & & \downarrow \\ & & V_2(\mathbb{R}^3) & & V_2(\mathbb{R}^5) & & \cdots & & V_2(\mathbb{R}^{2k-1}) & & S^{2k-1} \end{array}$$

where the last bundle $O(2k)\to S^{2k-1}$ is omitted if $n=2k-1$. As we remarked at the end of §3.D in the case of $SO(n)$, the integral homology and cohomology groups of $O(n)$ are the same as if these bundles were products, but the cup product structure for $O(n)$ with $\mathbb{Z}_2$ coefficients is not the same as in this product.

Cohomology of Grassmannians

Here is an important application of the Leray–Hirsch theorem, generalizing the calculation of the cohomology rings of projective spaces:

Theorem 4D.4. *If $G_n(\mathbb{C}^\infty)$ is the Grassmann manifold of n-dimensional vector subspaces of $\mathbb{C}^\infty$, then $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ is a polynomial ring $\mathbb{Z}[c_1,\cdots,c_n]$ on generators c_i of dimension $2i$. Similarly, $H^*(G_n(\mathbb{R}^\infty);\mathbb{Z}_2)$ is a polynomial ring $\mathbb{Z}_2[w_1,\cdots,w_n]$ on generators w_i of dimension i, and $H^*(G_n(\mathbb{H}^\infty);\mathbb{Z})\approx\mathbb{Z}[q_1,\cdots,q_n]$ with q_i of dimension $4i$.*

The plan of the proof is to apply the Leray–Hirsch theorem to a fiber bundle $F\longrightarrow E\xrightarrow{p} G_n(\mathbb{C}^\infty)$ where E has the same cohomology ring as the product of n copies of $\mathbb{C}\mathrm{P}^\infty$, a polynomial ring $\mathbb{Z}[x_1,\cdots,x_n]$ with each x_i 2-dimensional. The induced map $p^*:H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})\to H^*(E;\mathbb{Z})$ will be injective, and we will show that its image consists of the symmetric polynomials in $\mathbb{Z}[x_1,\cdots,x_n]$, the polynomials invariant under permutations of the variables x_i. It is a classical theorem in algebra that the symmetric polynomials themselves form a polynomial ring $\mathbb{Z}[\sigma_1,\cdots,\sigma_n]$ where σ_i is a certain symmetric polynomial of degree i, namely the sum of all products of i distinct x_j's. This gives the result for $G_n(\mathbb{C}^\infty)$, and the same argument will also apply in the real and quaternionic cases.

Proof: Define an ***n*-flag** in $\mathbb{C}^k$ to be an ordered n-tuple of orthogonal 1-dimensional vector subspaces of $\mathbb{C}^k$. Equivalently, an n-flag could be defined as a chain of vector subspaces $V_1 \subset \cdots \subset V_n$ of $\mathbb{C}^k$ where V_i has dimension i. Why either of these objects should be called a 'flag' is not exactly clear, but that is the traditional name. The set of all n-flags in $\mathbb{C}^k$ forms a subspace $F_n(\mathbb{C}^k)$ of the product of n copies of $\mathbb{C}\mathrm{P}^{k-1}$. There is a natural fiber bundle

$$F_n(\mathbb{C}^n) \longrightarrow F_n(\mathbb{C}^k) \xrightarrow{p} G_n(\mathbb{C}^k)$$

where p sends an n-tuple of orthogonal lines to the n-plane it spans. The local triviality property can be verified just as was done for the analogous Stiefel bundle $V_n(\mathbb{C}^n) \to V_n(\mathbb{C}^k) \to G_n(\mathbb{C}^k)$ in Example 4.53. The case $k = \infty$ is covered by the same argument, and this case will be the bundle $F \to E \to G_n(\mathbb{C}^\infty)$ alluded to in the paragraph preceding the proof.

The first step in the proof is to show that $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}) \approx \mathbb{Z}[x_1,\cdots,x_n]$ where x_i is the pullback of a generator of $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ under the map $F_n(\mathbb{C}^\infty) \to \mathbb{C}\mathrm{P}^\infty$ projecting an n-flag onto its i^{th} line. This can be seen by considering the fiber bundle

$$\mathbb{C}\mathrm{P}^\infty \longrightarrow F_n(\mathbb{C}^\infty) \xrightarrow{p} F_{n-1}(\mathbb{C}^\infty)$$

where p projects an n-flag onto the $(n-1)$-flag obtained by ignoring its last line. The local triviality property can be verified by the argument in Example 4.54. The Leray–Hirsch theorem applies since the powers of x_n restrict to a basis for $H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ in the fibers $\mathbb{C}\mathrm{P}^\infty$, each fiber being the space of lines in a vector subspace $\mathbb{C}^\infty$ of the standard $\mathbb{C}^\infty$. The elements x_i for $i < n$ are the pullbacks via p of elements of $H^*(F_{n-1}(\mathbb{C}^\infty);\mathbb{Z})$ defined in the same way. By induction $H^*(F_{n-1}(\mathbb{C}^\infty);\mathbb{Z})$ is a polynomial ring on these elements. From the Leray–Hirsch theorem we conclude that the products of powers of the x_i's for $1 \le i \le n$ form an additive basis for $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$, hence this ring is the polynomial ring on the x_i's.

There is a corresponding result for $F_n(\mathbb{C}^k)$, that $H^*(F_n(\mathbb{C}^k);\mathbb{Z})$ is free with basis the monomials $x_1^{i_1}\cdots x_n^{i_n}$ with $i_j \le k-j$ for each j. This is proved in exactly the same way, using induction on n and the fiber bundle $\mathbb{C}\mathrm{P}^{k-n} \to F_n(\mathbb{C}^k) \to F_{n-1}(\mathbb{C}^k)$. Thus the cohomology groups of $F_n(\mathbb{C}^k)$ are isomorphic to those of $\mathbb{C}\mathrm{P}^{k-1} \times \cdots \times \mathbb{C}\mathrm{P}^{k-n}$.

After these preliminaries we can start the main argument, using the fiber bundle $F_n(\mathbb{C}^n) \longrightarrow F_n(\mathbb{C}^\infty) \xrightarrow{p} G_n(\mathbb{C}^\infty)$. The preceding calculations show that the Leray–Hirsch theorem applies, so $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$ is a free module over $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ with basis the monomials $x_1^{i_1}\cdots x_n^{i_n}$ with $i_j \le n-j$ for each j. In particular, since 1 is among the basis elements, the homomorphism p^* is injective and its image is a direct summand of $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$. It remains to show that the image of p^* is exactly the symmetric polynomials.

To show that the image of p^* is contained in the symmetric polynomials, consider a map $\pi : F_n(\mathbb{C}^\infty) \to F_n(\mathbb{C}^\infty)$ permuting the lines in each n-flag according to a given permutation of the numbers $1, \cdots, n$. The induced map π^* on $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}) \approx \mathbb{Z}[x_1, \cdots, x_n]$ is the corresponding permutation of the variables x_i. Since permuting the lines in an n-flag has no effect on the n-plane they span, we have $p\pi = p$, hence $\pi^* p^* = p^*$, which says that polynomials in the image of p^* are invariant under permutations of the variables.

As remarked earlier, the symmetric polynomials in $\mathbb{Z}[x_1, \cdots, x_n]$ form a polynomial ring $\mathbb{Z}[\sigma_1, \cdots, \sigma_n]$ where σ_i has degree i. We have shown that the image of p^* is a direct summand, so to show that p^* maps onto the symmetric polynomials it will suffice to show that the graded rings $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ and $\mathbb{Z}[\sigma_1, \cdots, \sigma_n]$ have the same rank in each dimension, where the rank of a finitely generated free abelian group is the number of $\mathbb{Z}$ summands.

For a graded free $\mathbb{Z}$-module $A = \bigoplus_i A_i$, define its **Poincaré series** to be the formal power series $p_A(t) = \sum_i a_i t^i$ where a_i is the rank of A_i, which we assume to be finite for all i. The basic formula we need is that $p_{A\otimes B}(t) = p_A(t)\, p_B(t)$, which is immediate from the definition of the graded tensor product.

In the case at hand all nonzero cohomology is in even dimensions, so let us simplify notation by taking A_i to be the $2i$-dimensional cohomology of the space in question. Since the Poincaré series of $\mathbb{Z}[x]$ is $\sum_i t^i = (1-t)^{-1}$, the Poincaré series of $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$ is $(1-t)^{-n}$. For $H^*(F_n(\mathbb{C}^n);\mathbb{Z})$ the Poincaré series is

$$(1+t)(1+t+t^2)\cdots(1+t+\cdots+t^{n-1}) = \prod_{i=1}^{n} \frac{1-t^i}{1-t} = (1-t)^{-n} \prod_{i=1}^{n}(1-t^i)$$

From the additive isomorphism $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z}) \approx H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}) \otimes H^*(F_n(\mathbb{C}^n);\mathbb{Z})$ we see that the Poincaré series $p(t)$ of $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ satisfies

$$p(t)(1-t)^{-n} \prod_{i=1}^{n}(1-t^i) = (1-t)^{-n} \qquad \text{and hence} \qquad p(t) = \prod_{i=1}^{n}(1-t^i)^{-1}$$

This is exactly the Poincaré series of $\mathbb{Z}[\sigma_1, \cdots, \sigma_n]$ since σ_i has degree i. As noted before, this implies that the image of p^* is all the symmetric polynomials.

This finishes the proof for $G_n(\mathbb{C}^\infty)$. The same arguments apply in the other two cases, using $\mathbb{Z}_2$ coefficients throughout in the real case and replacing 'rank' by 'dimension' for $\mathbb{Z}_2$ vector spaces. ☐

These calculations show that the isomorphism $H^*(E;R) \approx H^*(B;R) \otimes_R H^*(F;R)$ of the Leray–Hirsch theorem is not generally a ring isomorphism, for if it were, then the polynomial ring $H^*(F_n(\mathbb{C}^\infty);\mathbb{Z})$ would contain a copy of $H^*(F_n(\mathbb{C}^n);\mathbb{Z})$ as a subring, but in the latter ring some power of every positive-dimensional element is zero since $H^k(F_n(\mathbb{C}^n);\mathbb{Z}) = 0$ for sufficiently large k.

The Gysin Sequence

Besides the Leray–Hirsch theorem, which deals with fiber bundles that are cohomologically like products, there is another special class of fiber bundles for which an elementary analysis of their cohomology structure is possible. These are fiber bundles $S^{n-1}\to E\xrightarrow{p} B$ satisfying an orientability hypothesis that will always hold if B is simply-connected or if we take cohomology with $\mathbb{Z}_2$ coefficients. For such bundles we will show there is an exact sequence, called the **Gysin sequence**,

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

where e is a certain 'Euler class' in $H^n(B;R)$. Since $H^i(B;R)=0$ for $i<0$, the initial portion of the Gysin sequence gives isomorphisms $p^*:H^i(B;R)\xrightarrow{\approx} H^i(E;R)$ for $i<n-1$, and the more interesting part of the sequence begins

$$0\longrightarrow H^{n-1}(B;R)\xrightarrow{p^*} H^{n-1}(E;R)\longrightarrow H^0(B;R)\xrightarrow{\smile e} H^n(B;R)\xrightarrow{p^*} H^n(E;R)\longrightarrow\cdots$$

In the case of a product bundle $E=S^{n-1}\times B$ there is a section, a map $s:B\to E$ with $ps=1\!\!1$, so the Gysin sequence breaks up into split short exact sequences

$$0\longrightarrow H^i(B;R)\xrightarrow{p^*} H^i(S^{n-1}\times B;R)\longrightarrow H^{i-n+1}(B;R)\longrightarrow 0$$

which agrees with the Künneth formula $H^*(S^{n-1}\times B;R)\approx H^*(S^{n-1};R)\otimes_R H^*(B;R)$. The splitting holds whenever the bundle has a section, even if it is not a product.

For example, consider the bundle $S^{n-1}\to V_2(\mathbb{R}^{n+1})\xrightarrow{p} S^n$. Points of $V_2(\mathbb{R}^{n+1})$ are pairs (v_1,v_2) of orthogonal unit vectors in $\mathbb{R}^{n+1}$, and $p(v_1,v_2)=v_1$. If we think of v_1 as a point of S^n and v_2 as a unit vector tangent to S^n at v_1, then $V_2(\mathbb{R}^{n+1})$ is exactly the bundle of unit tangent vectors to S^n. A section of this bundle is a field of unit tangent vectors to S^n, and such a vector field exists iff n is odd by Theorem 2.28. The fact that the Gysin sequence splits when there is a section then says that $V_2(\mathbb{R}^{n+1})$ has the same cohomology as the product $S^{n-1}\times S^n$ if n is odd, at least when $n>1$ so that the base space S^n is simply-connected and the orientability hypothesis is satisfied. When n is even, the calculations at the end of §3.D show that $H^*(V_2(\mathbb{R}^{n+1});\mathbb{Z})$ consists of $\mathbb{Z}$'s in dimensions 0 and $2n-1$ and a $\mathbb{Z}_2$ in dimension n. The latter group appears in the Gysin sequence as

$$\begin{array}{ccccccc} H^0(S^n) & \xrightarrow{\smile e} & H^n(S^n) & \longrightarrow & H^n(V_2(\mathbb{R}^{n+1})) & \longrightarrow & H^1(S^n)\\ \| & & \| & & \| & & \|\\ \mathbb{Z} & & \mathbb{Z} & & \mathbb{Z}_2 & & 0\end{array}$$

hence the Euler class e must be twice a generator of $H^n(S^n)$ in the case that n is even. When n is odd it must be zero in order for the Gysin sequence to split.

This example illustrates a theorem in differential topology that explains why the Euler class has this name: The Euler class of the unit tangent bundle of a closed orientable smooth n-manifold M is equal to the Euler characteristic $\chi(M)$ times a generator of $H^n(M;\mathbb{Z})$.

Whenever a bundle $S^{n-1}\to E\xrightarrow{p} B$ has a section, the Euler class e must be zero from exactness of $H^0(B)\xrightarrow{\smile e} H^n(B)\xrightarrow{p^*} H^n(E)$ since p^* is injective if there is a section. Thus the Euler class can be viewed as an obstruction to the existence of a section: If the Euler class is nonzero, there can be no section. This qualitative statement can be made more precise by bringing in the machinery of obstruction theory, as explained in [Milnor & Stasheff 1974] or [VBKT].

Before deriving the Gysin sequence let us look at some examples of how it can be used to compute cup products.

Example 4D.5. Consider a bundle $S^{n-1} \longrightarrow E \xrightarrow{p} B$ with E contractible, for example the bundle $S^1 \to S^\infty \to \mathbb{C}P^\infty$ or its real or quaternionic analogs. The long exact sequence of homotopy groups for the bundle shows that B is $(n-1)$-connected. Thus if $n > 1$, B is simply-connected and we have a Gysin sequence for cohomology with $\mathbb{Z}$ coefficients. For $n = 1$ we take $\mathbb{Z}_2$ coefficients. If $n > 1$ then since E is contractible, the Gysin sequence implies that $H^i(B;\mathbb{Z}) = 0$ for $0 < i < n$ and that $\smile e : H^i(B;\mathbb{Z}) \to H^{i+n}(B;\mathbb{Z})$ is an isomorphism for $i \geq 0$. It follows that $H^*(B;\mathbb{Z})$ is the polynomial ring $\mathbb{Z}[e]$. When $n = 1$ the map $p^* : H^{n-1}(B;\mathbb{Z}_2) \to H^{n-1}(E;\mathbb{Z}_2)$ in the Gysin sequence is surjective, so we see that $\smile e : H^i(B;\mathbb{Z}_2) \to H^{i+n}(B;\mathbb{Z}_2)$ is again an isomorphism for all $i \geq 0$, and hence $H^*(B;\mathbb{Z}_2) \approx \mathbb{Z}_2[e]$. Thus the Gysin sequence gives a new derivation of the cup product structure in projective spaces. Also, since polynomial rings $\mathbb{Z}[e]$ are realizable as $H^*(X;\mathbb{Z})$ only when e has dimension 2 or 4, as we show in Corollary 4L.10, we can conclude that there exist bundles $S^{n-1} \to E \to B$ with E contractible only when n is 1, 2, or 4.

Example 4D.6. For the Grassmann manifold $G_n = G_n(\mathbb{R}^\infty)$ we have $\pi_1(G_n) \approx \pi_0 O(n) \approx \mathbb{Z}_2$, so the universal cover of G_n gives a bundle $S^0 \to \widetilde{G}_n \to G_n$. One can view $\widetilde{G}_n$ as the space of oriented n-planes in $\mathbb{R}^\infty$, which is obviously a 2-sheeted covering space of G_n, hence the universal cover since it is path-connected, being the quotient $V_n(\mathbb{R}^\infty)/SO(n)$ of the contractible space $V_n(\mathbb{R}^\infty)$. A portion of the Gysin sequence for the bundle $S^0 \to \widetilde{G}_n \to G_n$ is $H^0(G_n;\mathbb{Z}_2) \xrightarrow{\smile e} H^1(G_n;\mathbb{Z}_2) \longrightarrow H^1(\widetilde{G}_n;\mathbb{Z}_2)$. This last group is zero since $\widetilde{G}_n$ is simply-connected, and $H^1(G_n;\mathbb{Z}_2) \approx \mathbb{Z}_2$ since $H^*(G_n;\mathbb{Z}_2) \approx \mathbb{Z}_2[w_1, \cdots, w_n]$ as we showed earlier in this section, so $e = w_1$ and the map $\smile e : H^*(G_n;\mathbb{Z}_2) \to H^*(G_n;\mathbb{Z}_2)$ is injective. The Gysin sequence then breaks up into short exact sequences $0 \to H^i(G_n;\mathbb{Z}_2) \xrightarrow{\smile e} H^{i+1}(G_n;\mathbb{Z}_2) \to H^{i+1}(\widetilde{G}_n;\mathbb{Z}_2) \to 0$, from which it follows that $H^*(\widetilde{G}_n;\mathbb{Z}_2)$ is the quotient ring $\mathbb{Z}_2[w_1, \cdots, w_n]/(w_1) \approx \mathbb{Z}_2[w_2, \cdots, w_n]$.

Example 4D.7. The complex analog of the bundle in the preceding example is a bundle $S^1 \to \widetilde{G}_n(\mathbb{C}^\infty) \to G_n(\mathbb{C}^\infty)$ with $\widetilde{G}_n(\mathbb{C}^\infty)$ 2-connected. This can be constructed in the following way. There is a determinant homomorphism $U(n) \to S^1$ with kernel $SU(n)$, the unitary matrices of determinant 1, so S^1 is the coset space $U(n)/SU(n)$, and by restricting the action of $U(n)$ on $V_n(\mathbb{C}^\infty)$ to $SU(n)$ we obtain the second row of the commutative diagram at the right. The second row is a fiber bundle by the usual argument of choosing continuously varying orthonormal bases in n-planes near a

$$\begin{array}{ccccc} U(n) & \longrightarrow & V_n(\mathbb{C}^\infty) & \longrightarrow & G_n(\mathbb{C}^\infty) \\ \downarrow & & \downarrow & & \| \\ S^1 & \longrightarrow & V_n(\mathbb{C}^\infty)/SU(n) & \longrightarrow & G_n(\mathbb{C}^\infty) \end{array}$$

given n-plane. One sees that the space $\widetilde{G}_n(\mathbb{C}^\infty) = V_n(\mathbb{C}^\infty)/SU(n)$ is 2-connected by looking at the relevant portion of the diagram of homotopy groups associated to these two bundles:

$$\begin{array}{ccccccccc} & & 0 & \longrightarrow & \pi_2(G_n) & \xrightarrow[\approx]{\partial} & \pi_1(U(n)) & \longrightarrow & 0 \\ & & & & \| & & \downarrow\approx & & \\ 0 & \longrightarrow & \pi_2(\widetilde{G}_n) & \longrightarrow & \pi_2(G_n) & \xrightarrow{\partial} & \pi_1(S^1) & \longrightarrow & \pi_1(\widetilde{G}_n) \longrightarrow 0 \end{array}$$

The second vertical map is an isomorphism since S^1 embeds in $U(n)$ as the subgroup $U(1)$. Since the boundary map in the upper row is an isomorphism, so also is the boundary map in the lower row, and then exactness implies that $\widetilde{G}_n$ is 2-connected.

The Gysin sequence for $S^1 \to \widetilde{G}_n(\mathbb{C}^\infty) \to G_n(\mathbb{C}^\infty)$ can be analyzed just as in the preceding example. Part of the sequence is $H^0(G_n;\mathbb{Z}) \xrightarrow{\smile e} H^2(G_n;\mathbb{Z}) \longrightarrow H^2(\widetilde{G}_n;\mathbb{Z})$, and this last group is zero since $\widetilde{G}_n$ is 2-connected, so e must be a generator of $H^2(G_n;\mathbb{Z}) \approx \mathbb{Z}$. Since $H^*(G_n;\mathbb{Z})$ is a polynomial algebra $\mathbb{Z}[c_1,\cdots,c_n]$, we must have $e = \pm c_1$, so the map $\smile e : H^*(G_n;\mathbb{Z}) \to H^*(G_n;\mathbb{Z})$ is injective, the Gysin sequence breaks up into short exact sequences, and $H^*(\widetilde{G}_n;\mathbb{Z})$ is the quotient ring $\mathbb{Z}[c_1,\cdots,c_n]/(c_1) \approx \mathbb{Z}[c_2,\cdots,c_n]$.

The spaces $\widetilde{G}_n$ in the last two examples are often denoted $BSO(n)$ and $BSU(n)$, expressing the fact that they are related to the groups $SO(n)$ and $SU(n)$ via bundles $SO(n) \to V_n(\mathbb{R}^\infty) \to BSO(n)$ and $SU(n) \to V_n(\mathbb{C}^\infty) \to BSU(n)$ with contractible total spaces V_n. There is no quaternion analog of $BSO(n)$ and $BSU(n)$ since for $n = 2$ this would give a space with cohomology ring $\mathbb{Z}[x]$ on an 8-dimensional generator, which is impossible by Corollary 4L.10.

Now we turn to the derivation of the Gysin sequence, which follows a rather roundabout route:

(1) Deduce a relative version of the Leray–Hirsch theorem from the absolute case.

(2) Specialize this to the case of bundles with fiber a disk, yielding a basic result called the Thom isomorphism.

(3) Show this applies to all orientable disk bundles.

(4) Deduce the Gysin sequence by plugging the Thom isomorphism into the long exact sequence of cohomology groups for the pair consisting of a disk bundle and its boundary sphere bundle.

(1) A **fiber bundle pair** consists of a fiber bundle $p: E \to B$ with fiber F, together with a subspace $E' \subset E$ such that $p: E' \to B$ is a bundle with fiber a subspace $F' \subset F$, with local trivializations for E' obtained by restricting local trivializations for E. For example, if $E \to B$ is a bundle with fiber D^n and $E' \subset E$ is the union of the boundary spheres of the fibers, then (E, E') is a fiber bundle pair since local trivializations of E restrict to local trivializations of E', in view of the fact that homeomorphisms from an n-disk to an n-disk restrict to homeomorphisms between their boundary spheres, boundary and interior points of D^n being distinguished by the local homology groups $H_n(D^n, D^n - \{x\};\mathbb{Z})$.

Theorem 4D.8. *Suppose that* $(F,F')\to(E,E')\xrightarrow{p}B$ *is a fiber bundle pair such that* $H^*(F,F';R)$ *is a free* R*-module, finitely generated in each dimension. If there exist classes* $c_j\in H^*(E,E';R)$ *whose restrictions form a basis for* $H^*(F,F';R)$ *in each fiber* (F,F')*, then* $H^*(E,E';R)$*, as a module over* $H^*(B;R)$*, is free with basis* $\{c_j\}$*.*

The module structure is defined just as in the absolute case by $bc=p^*(b)\smile c$, but now we use the relative cup product $H^*(E;R)\times H^*(E,E';R)\to H^*(E,E';R)$.

Proof: Construct a bundle $\hat{E}\to B$ from E by attaching the mapping cylinder M of $p:E'\to B$ to E by identifying the subspaces $E'\subset E$ and $E'\subset M$. Thus the fibers $\hat{F}$ of $\hat{E}$ are obtained from the fibers F by attaching cones CF' on the subspaces $F'\subset F$. Regarding B as the subspace of $\hat{E}$ at one end of the mapping cylinder M, we have $H^*(\hat{E},M;R)\approx H^*(\hat{E}-B,M-B;R)\approx H^*(E,E';R)$ via excision and the obvious deformation retraction of $\hat{E}-B$ onto E. The long exact sequence of a triple gives $H^*(\hat{E},M;R)\approx H^*(\hat{E},B;R)$ since M deformation retracts to B. All these isomorphisms are $H^*(B;R)$-module isomorphisms. Since B is a retract of $\hat{E}$ via the bundle projection $\hat{E}\to B$, we have a splitting $H^*(\hat{E};R)\approx H^*(\hat{E},B;R)\oplus H^*(B;R)$ as $H^*(B;R)$-modules. Let $\hat{c}_j\in H^*(\hat{E};R)$ correspond to $c_j\in H^*(E,E';R)\approx H^*(\hat{E},B;R)$ in this splitting. The classes $\hat{c}_j$ together with 1 restrict to a basis for $H^*(\hat{F};R)$ in each fiber $\hat{F}=F\cup CF'$, so the absolute form of the Leray–Hirsch theorem implies that $H^*(\hat{E};R)$ is a free $H^*(B;R)$-module with basis $\{1,\hat{c}_j\}$. It follows that $\{c_j\}$ is a basis for the free $H^*(B;R)$-module $H^*(E,E';R)$. □

(2) Now we specialize to the case of a fiber bundle pair $(D^n,S^{n-1})\longrightarrow(E,E')\xrightarrow{p}B$. An element $c\in H^n(E,E';R)$ whose restriction to each fiber (D^n,S^{n-1}) is a generator of $H^n(D^n,S^{n-1};R)\approx R$ is called a **Thom class** for the bundle. We are mainly interested in the cases $R=\mathbb{Z}$ and $\mathbb{Z}_2$, but R could be any commutative ring with identity, in which case a 'generator' is an element with a multiplicative inverse, so all elements of R are multiples of the generator. A Thom class with $\mathbb{Z}$ coefficients gives rise to a Thom class with any other coefficient ring R under the homomorphism $H^n(E,E';\mathbb{Z})\to H^n(E,E';R)$ induced by the homomorphism $\mathbb{Z}\to R$ sending 1 to the identity element of R.

Corollary 4D.9. *If the disk bundle* $(D^n,S^{n-1})\longrightarrow(E,E')\xrightarrow{p}B$ *has a Thom class* $c\in H^n(E,E';R)$*, then the map* $\Phi:H^i(B;R)\to H^{i+n}(E,E';R)$*,* $\Phi(b)=p^*(b)\smile c$*, is an isomorphism for all* $i\ge 0$*, and* $H^i(E,E';R)=0$ *for* $i<n$*.* □

The isomorphism Φ is called the **Thom isomorphism.** The corollary can be made into a statement about absolute cohomology by defining the **Thom space** $T(E)$ to be the quotient E/E'. Each disk fiber D^n of E becomes a sphere S^n in $T(E)$, and all these spheres coming from different fibers are disjoint except for the common basepoint $x_0=E'/E'$. A Thom class can be regarded as an element of $H^n(T(E),x_0;R)\approx$

$H^n(T(E);R)$ that restricts to a generator of $H^n(S^n;R)$ in each 'fiber' S^n in $T(E)$, and the Thom isomorphism becomes $H^i(B;R) \approx \widetilde{H}^{n+i}(T(E);R)$.

(3) The major remaining step in the derivation of the Gysin sequence is to relate the existence of a Thom class for a disk bundle $D^n \to E \to B$ to a notion of orientability of the bundle. First we define orientability for a sphere bundle $S^{n-1} \to E' \to B$. In the proof of Proposition 4.61 we described a procedure for lifting paths γ in B to homotopy equivalences L_γ between the fibers above the endpoints of γ. We did this for fibrations rather than fiber bundles, but the method applies equally well to fiber bundles whose fiber is a CW complex since the homotopy lifting property was used only for the fiber and for the product of the fiber with I. In the case of a sphere bundle $S^{n-1} \to E' \to B$, if γ is a loop in B then L_γ is a homotopy equivalence from the fiber S^{n-1} over the basepoint of γ to itself, and we define the sphere bundle to be **orientable** if L_γ induces the identity map on $H^{n-1}(S^{n-1};\mathbb{Z})$ for each loop γ in B.

For example, the Klein bottle, regarded as a bundle over S^1 with fiber S^1, is nonorientable since as we follow a path looping once around the base circle, the corresponding fiber circles sweep out the full Klein bottle, ending up where they started but with orientation reversed. The same reasoning shows that the torus, viewed as a circle bundle over S^1, is orientable. More generally, any sphere bundle that is a product is orientable since the maps L_γ can be taken to be the identity for all loops γ. Also, sphere bundles over simply-connected base spaces are orientable since $\gamma \simeq \eta$ implies $L_\gamma \simeq L_\eta$, hence all L_γ's are homotopic to the identity when all loops γ are nullhomotopic.

One could define orientability for a disk bundle $D^n \to E \to B$ by relativizing the previous definition, constructing lifts L_γ which are homotopy equivalences of the fiber pairs (D^n, S^{n-1}). However, since $H^n(D^n, S^{n-1};\mathbb{Z})$ is canonically isomorphic to $H^{n-1}(S^{n-1};\mathbb{Z})$ via the coboundary map in the long exact sequence of the pair, it is simpler and amounts to the same thing just to define E to be orientable if its boundary sphere subbundle E' is orientable.

Theorem 4D.10. *Every disk bundle has a Thom class with $\mathbb{Z}_2$ coefficients, and orientable disk bundles have Thom classes with $\mathbb{Z}$ coefficients.*

An exercise at the end of the section is to show that the converse of the last statement is also true: A disk bundle is orientable if it has a Thom class with $\mathbb{Z}$ coefficients.

Proof: The case of a non-CW base space B reduces to the CW case by pulling back over a CW approximation to B, as in the Leray–Hirsch theorem, applying the five-lemma to say that the pullback bundle has isomorphic homotopy groups, hence isomorphic absolute and relative cohomology groups. From the definition of the pullback bundle it is immediate that the pullback of an orientable sphere bundle is orientable. There is also no harm in assuming the base CW complex B is connected. We will show:

(∗) If the disk bundle $D^n \to E \to B$ is orientable and B is a connected CW complex, then the restriction map $H^i(E,E';\mathbb{Z}) \to H^i(D^n_x,S^{n-1}_x;\mathbb{Z})$ is an isomorphism for all fibers D^n_x, $x \in B$, and for all $i \le n$.

For $\mathbb{Z}_2$ coefficients we will see that $(*)$ holds without any orientability hypothesis. Hence with either $\mathbb{Z}$ or $\mathbb{Z}_2$ coefficients, a generator of $H^n(E,E') \approx H^n(D^n_x,S^{n-1}_x)$ is a Thom class.

If the disk bundle $D^n \to E \to B$ is orientable, then if we choose an isomorphism $H^n(D^n_x,S^{n-1}_x;\mathbb{Z}) \approx \mathbb{Z}$ for one fiber D^n_x, this determines such isomorphisms for all fibers by composing with the isomorphisms L^*_γ, which depend only on the endpoints of γ. Having made such a choice, then if $(*)$ is true, we have a preferred isomorphism $H^n(E,E';\mathbb{Z}) \approx \mathbb{Z}$ which restricts to the chosen isomorphism $H^n(D^n_x,S^{n-1}_x;\mathbb{Z}) \approx \mathbb{Z}$ for each fiber. This is because for a path γ from x to y, the inclusion $(D^n_x,S^{n-1}_x) \hookrightarrow (E,E')$ is homotopic to the composition of L_γ with the inclusion $(D^n_y,S^{n-1}_y) \hookrightarrow (E,E')$. We will use this preferred isomorphism $H^n(E,E';\mathbb{Z}) \approx \mathbb{Z}$ in the inductive proof of $(*)$ given below. In the case of $\mathbb{Z}_2$ coefficients, there can be only one isomorphism of a group with $\mathbb{Z}_2$ so no choices are necessary and orientability is irrelevant. We will prove $(*)$ in the $\mathbb{Z}$ coefficient case, leaving it to the reader to replace all $\mathbb{Z}$'s in the proof by $\mathbb{Z}_2$'s to obtain a proof in the $\mathbb{Z}_2$ case.

Suppose first that the CW complex B has finite dimension k. Let $U \subset B$ be the subspace obtained by deleting one point from the interior of each k-cell of B, and let $V \subset B$ be the union of the open k-cells. Thus $B = U \cup V$. For a subspace $A \subset B$ let $E_A \to A$ and $E'_A \to A$ be the disk and sphere bundles obtained by taking the subspaces of E and E' projecting to A. Consider the following portion of a Mayer–Vietoris sequence, with $\mathbb{Z}$ coefficients implicit from now on:

$$H^n(E,E') \longrightarrow H^n(E_U,E'_U) \oplus H^n(E_V,E'_V) \xrightarrow{\ \Psi\ } H^n(E_{U\cap V},E'_{U\cap V})$$

The first map is injective since the preceding term in the sequence is zero by induction on k, since $U \cap V$ deformation retracts onto a disjoint union of $(k-1)$-spheres and we can apply Lemma 4D.2 to replace $E_{U\cap V}$ by the part of E over this union of $(k-1)$-spheres. By exactness we then have an isomorphism $H^n(E,E') \approx \operatorname{Ker}\Psi$. Similarly, by Lemma 4D.2 and induction each of the terms $H^n(E_U,E'_U)$, $H^n(E_V,E'_V)$, and $H^n(E_{U\cap V},E'_{U\cap V})$ is a product of $\mathbb{Z}$'s, with one $\mathbb{Z}$ factor for each component of the spaces involved, projection onto the $\mathbb{Z}$ factor being given by restriction to any fiber in the component. Elements of $\operatorname{Ker}\Psi$ are pairs $(\alpha,\beta) \in H^n(E_U,E'_U) \oplus H^n(E_V,E'_V)$ having the same restriction to $H^n(E_{U\cap V},E'_{U\cap V})$. Since B is connected, this means that all the $\mathbb{Z}$ coordinates of α and β in the previous direct product decompositions must be equal, since between any two components of U or V one can interpolate a finite sequence of components of U and V alternately, each component in the sequence having nontrivial intersection with its neighbors. Thus $\operatorname{Ker}\Psi$ is a copy of $\mathbb{Z}$, with restriction to a fiber being the isomorphism $H^n(E,E') \approx \mathbb{Z}$.

To finish proving $(*)$ for finite-dimensional B it remains to see that $H^i(E,E') = 0$ for $i < n$, but this follows immediately by looking at an earlier stage of the Mayer–Vietoris sequence, where the two terms adjacent to $H^i(E,E')$ vanish by induction.

Proving $(*)$ for an infinite-dimensional CW complex B reduces to the finite-dimensional case as in the Leray–Hirsch theorem since we are only interested in cohomology in a finite range of dimensions. $\square$

(4) Now we can derive the Gysin sequence for a sphere bundle $S^{n-1} \to E \xrightarrow{p} B$. Consider the mapping cylinder M_p, which is a disk bundle $D^n \to M_p \xrightarrow{p} B$ with E as its boundary sphere bundle. Assuming that a Thom class $c \in H^n(M_p, E; R)$ exists, as is the case if E is orientable or if $R = \mathbb{Z}_2$, then the long exact sequence of cohomology groups for the pair (M_p, E) gives the first row of the following commutative diagram, with R coefficients implicit:

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

The maps Φ are the Thom isomorphism, and the vertical map p^* is an isomorphism since M_p deformation retracts onto B. The **Euler class** $e \in H^n(B;R)$ is defined to be $(p^*)^{-1}j^*(c)$, c being a Thom class. The square containing the map $\smile e$ commutes since for $b \in H^{i-n}(B;R)$ we have $j^*\Phi(b) = j^*(p^*(b) \smile c) = p^*(b) \smile j^*(c)$, which equals $p^*(b \smile e) = p^*(b) \smile p^*(e)$ since $p^*(e) = j^*(c)$. Another way of defining e is as the class corresponding to $c \smile c$ under the Thom isomorphism, since $\Phi(e) = p^*(e) \smile c = j^*(c) \smile c = c \smile c$.

Finally, the lower row of the diagram is by definition the Gysin sequence. $\square$

To conclude this section we will use the following rather specialized application of the Gysin sequence to compute a few more examples of spaces with polynomial cohomology.

Proposition 4D.11. *Suppose that $S^{2k-1} \longrightarrow E \xrightarrow{p} B$ is an orientable sphere bundle such that $H^*(E;R)$ is a polynomial ring $R[x_1, \cdots, x_\ell]$ on even-dimensional generators x_i. Then $H^*(B;R) = R[y_1, \cdots, y_\ell, e]$ where e is the Euler class of the bundle and $p^*(y_i) = x_i$ for each i.*

Proof: Consider the three terms $H^i(B;R) \xrightarrow{\smile e} H^{i+2k}(B;R) \longrightarrow H^{i+2k}(E;R)$ of the Gysin sequence. If i is odd, the third term is zero since E has no odd-dimensional cohomology. Hence the map $\smile e$ is surjective, and by induction on dimension this implies that $H^*(B;R)$ is zero in odd dimensions. This means the Gysin sequence reduces to short exact sequences

$$0 \longrightarrow H^{2i}(B;R) \xrightarrow{\smile e} H^{2i+2k}(B;R) \xrightarrow{p^*} H^{2i+2k}(E;R) \longrightarrow 0$$

Since p^* is surjective, we can choose elements $y_j \in H^*(B;R)$ with $p^*(y_j) = x_j$. It remains to check that $H^*(B;R) = R[y_1, \cdots, y_\ell, e]$, which is elementary algebra: Given $b \in H^*(B;R)$, $p^*(b)$ must be a polynomial $f(x_1, \cdots, x_\ell)$, so $b - f(y_1, \cdots, y_\ell)$ is in the kernel of p^* and exactness gives an equation $b - f(y_1, \cdots, y_\ell) = b' \smile e$ for some $b' \in H^*(B;R)$. Since b' has lower dimension than b, we may assume by induction that b' is a polynomial in $y_1, \cdots, y_\ell, e$. Hence $b = f(y_1, \cdots, y_\ell) + b' \smile e$ is also a polynomial in $y_1, \cdots, y_\ell, e$. Thus the natural map $R[y_1, \cdots, y_\ell, e] \to H^*(B;R)$ is surjective. To see that it is injective, suppose there is a polynomial relation $f(y_1, \cdots, y_\ell, e) = 0$ in $H^*(B;R)$. Applying p^*, we get $f(x_1, \cdots, x_\ell, 0) = 0$ since $p^*(y_i) = x_i$ and $p^*(e) = 0$ from the short exact sequence. The relation $f(x_1, \cdots, x_\ell, 0) = 0$ takes place in the polynomial ring $R[x_1, \cdots, x_\ell]$, so $f(y_1, \cdots, y_\ell, 0) = 0$ in $R[y_1, \cdots, y_\ell, e]$, hence $f(y_1, \cdots, y_\ell, e)$ must be divisible by e, say $f = ge$ for some polynomial g. The relation $f(y_1, \cdots, y_\ell, e) = 0$ in $H^*(B;R)$ then has the form $g(y_1, \cdots, y_\ell, e) \smile e = 0$. Since $\smile e$ is injective, this gives a polynomial relation $g(y_1, \cdots, y_\ell, e) = 0$ with g having lower degree than f. By induction we deduce that g must be the zero polynomial, hence also f. □

Example 4D.12. Let us apply this to give another proof that $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z})$ is a polynomial ring $\mathbb{Z}[c_1, \cdots, c_n]$ with $|c_i| = 2i$. We use two fiber bundles:

$$S^{2n-1} \longrightarrow E \longrightarrow G_n(\mathbb{C}^\infty) \qquad\qquad S^\infty \longrightarrow E \longrightarrow G_{n-1}(\mathbb{C}^\infty)$$

The total space E in both cases is the space of pairs (P, v) where P is an n-plane in $\mathbb{C}^\infty$ and v is a unit vector in P. In the first bundle the map $E \to G_n(\mathbb{C}^\infty)$ is $(P, v) \mapsto P$, with fiber S^{2n-1}, and for the second bundle the map $E \to G_{n-1}(\mathbb{C}^\infty)$ sends (P, v) to the $(n-1)$-plane in P orthogonal to v, with fiber S^∞ consisting of all the unit vectors in $\mathbb{C}^\infty$ orthogonal to a given $(n-1)$-plane. Local triviality for the two bundles is verified in the usual way. Since S^∞ is contractible, the map $E \to G_{n-1}(\mathbb{C}^\infty)$ induces isomorphisms on homotopy groups, hence also on cohomology. By induction on n we then have $H^*(E;\mathbb{Z}) \approx \mathbb{Z}[c_1, \cdots, c_{n-1}]$. The first bundle is orientable since $G_n(\mathbb{C}^\infty)$ is simply-connected, so the proposition gives $H^*(G_n(\mathbb{C}^\infty);\mathbb{Z}) \approx \mathbb{Z}[c_1, \cdots, c_n]$ for $c_n = e$.

The same argument works in the quaternionic case. For a version of this argument in the real case see §3.3 of [VBKT].

Before giving our next example, let us observe that the Gysin sequence with a fixed coefficient ring R is valid for any orientable fiber bundle $F \longrightarrow E \xrightarrow{p} B$ whose fiber is a CW complex F with $H^*(F;R) \approx H^*(S^{n-1};R)$. Orientability is defined just as before in terms of induced maps $L_\gamma^*: H^{n-1}(F;R) \to H^{n-1}(F;R)$. No changes are needed in the derivation of the Gysin sequence to get this more general case, if the associated 'disk' bundle is again taken to be the mapping cylinder $CF \to M_p \to B$.

Example 4D.13. We have computed the cohomology of $\widetilde{G}_n(\mathbb{R}^\infty)$ with $\mathbb{Z}_2$ coefficients, finding it to be a polynomial ring on generators in dimensions 2 through n, and now we compute the cohomology with $\mathbb{Z}_p$ coefficients for p an odd prime. The answer will again be a polynomial algebra, but this time on even-dimensional generators, depending on the parity of n. Consider first the case that n is odd, say $n = 2k+1$. There are two fiber bundles

$$V_2(\mathbb{R}^{2k+1}) \longrightarrow E \longrightarrow \widetilde{G}_{2k+1}(\mathbb{R}^\infty) \qquad\qquad V_2(\mathbb{R}^\infty) \longrightarrow E \longrightarrow \widetilde{G}_{2k-1}(\mathbb{R}^\infty)$$

where E is the space of triples (P, v_1, v_2) with P an oriented $(2k+1)$-plane in $\mathbb{R}^\infty$ and v_1 and v_2 two orthogonal unit vectors in P. The projection map in the first bundle is $(P, v_1, v_2) \mapsto P$, and for the second bundle the projection sends (P, v_1, v_2) to the oriented $(2k-1)$-plane in P orthogonal to v_1 and v_2, with the orientation specified by saying for example that v_1, v_2 followed by a positively oriented basis for the orthogonal $(2k-1)$-plane is a positively oriented basis for P. Both bundles are orientable since their base spaces $\widetilde{G}_n(\mathbb{R}^\infty)$ are simply-connected, from the bundle $SO(n) \to V_n(\mathbb{R}^\infty) \to \widetilde{G}_n(\mathbb{R}^\infty)$.

The fiber $V_2(\mathbb{R}^\infty)$ of the second bundle is contractible, so E has the same cohomology as $\widetilde{G}_{2k-1}(\mathbb{R}^\infty)$. The fiber of the first bundle has the same $\mathbb{Z}_p$ cohomology as S^{4k-1} if p is odd, by the calculation at the end of §3.D. So if we assume inductively that $H^*(\widetilde{G}_{2k-1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_{k-1}]$ with $|p_i| = 4i$, then Proposition 4D.11 above implies that $H^*(\widetilde{G}_{2k+1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_k]$ where $p_k = e$ has dimension $4k$. The induction can start with $\widetilde{G}_1(\mathbb{R}^\infty)$ which is just S^∞ since an oriented line in $\mathbb{R}^\infty$ contains a unique unit vector in the positive direction.

To handle the case of $\widetilde{G}_n(\mathbb{R}^\infty)$ with $n = 2k$ even, we proceed just as in Example 4D.12, considering the bundles

$$S^{2k-1} \longrightarrow E \longrightarrow \widetilde{G}_{2k}(\mathbb{R}^\infty) \qquad\qquad S^\infty \longrightarrow E \longrightarrow \widetilde{G}_{2k-1}(\mathbb{R}^\infty)$$

By the case n odd we have $H^*(\widetilde{G}_{2k-1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_{k-1}]$ with $|p_i| = 4i$, so the corollary implies that $H^*(\widetilde{G}_{2k}(\mathbb{R}^\infty); \mathbb{Z}_p)$ is a polynomial ring on these generators and also a generator in dimension $2k$.

Summarizing, for p an odd prime we have shown:

$$H^*(\widetilde{G}_{2k+1}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_k], \quad |p_i| = 4i$$
$$H^*(\widetilde{G}_{2k}(\mathbb{R}^\infty); \mathbb{Z}_p) \approx \mathbb{Z}_p[p_1, \cdots, p_{k-1}, e], \quad |p_i| = 4i, \;\; |e| = 2k$$

The same result holds also with $\mathbb{Q}$ coefficients. In fact, our proof applies for any coefficient ring in which 2 has a multiplicative inverse, since all that is needed is that $H^*(V_2(\mathbb{R}^{2k+1}); R) \approx H^*(S^{4k-1}; R)$. For a calculation of the cohomology of $\widetilde{G}_n(\mathbb{R}^\infty)$ with $\mathbb{Z}$ coefficients, see [VBKT]. It turns out that all torsion elements have order 2, and modulo this torsion the integral cohomology is again a polynomial ring on the generators p_i and e. Similar results hold also for the cohomology of the unoriented Grassmann manifold $G_n(\mathbb{R}^\infty)$, but with the generator e replaced by p_k when $n = 2k$.

Exercises

1. By Exercise 35 in §4.2 there is a bundle $S^2 \to \mathbb{C}P^3 \to S^4$. Let $S^2 \to E_k \to S^4$ be the pullback of this bundle via a degree k map $S^4 \to S^4$, $k > 1$. Use the Leray-Hirsch theorem to show that $H^*(E_k;\mathbb{Z})$ is additively isomorphic to $H^*(\mathbb{C}P^3;\mathbb{Z})$ but has a different cup product structure in which the square of a generator of $H^2(E_k;\mathbb{Z})$ is k times a generator of $H^4(E_k;\mathbb{Z})$.

2. Apply the Leray–Hirsch theorem to the bundle $S^1 \to S^\infty/\mathbb{Z}_p \to \mathbb{C}P^\infty$ to compute $H^*(K(\mathbb{Z}_p,1);\mathbb{Z}_p)$ from $H^*(\mathbb{C}P^\infty;\mathbb{Z}_p)$.

3. Use the Leray–Hirsch theorem as in Corollary 4D.3 to compute $H^*(V_n(\mathbb{C}^k);\mathbb{Z}) \approx \Lambda_{\mathbb{Z}}[x_{2k-2n+1}, x_{2k-2n+3}, \cdots, x_{2k-1}]$ and similarly in the quaternionic case.

4. For the flag space $F_n(\mathbb{C}^n)$ show that $H^*(F_n(\mathbb{C}^n);\mathbb{Z}) \approx \mathbb{Z}[x_1,\cdots,x_n]/(\sigma_1,\cdots,\sigma_n)$ where σ_i is the i^{th} elementary symmetric polynomial.

5. Use the Gysin sequence to show that for a fiber bundle $S^k \to S^m \xrightarrow{p} S^n$ we must have $k = n-1$ and $m = 2n-1$. Then use the Thom isomorphism to show that the Hopf invariant of p must be ± 1. [Hence $n = 1,2,4,8$ by Adams' theorem.]

6. Show that if M is a manifold of dimension $2n$ for which there exists a fiber bundle $S^1 \to S^{2n+1} \to M$, then M is simply-connected and $H^*(M;\mathbb{Z}) \approx H^*(\mathbb{C}P^n;\mathbb{Z})$ as rings. Conversely, if M is simply-connected and $H^*(M;\mathbb{Z}) \approx H^*(\mathbb{C}P^n;\mathbb{Z})$ as rings, show there is a bundle $S^1 \to E \to M$ where $E \simeq S^{2n+1}$. [When $n > 1$ there are examples where M is not homeomorphic to $\mathbb{C}P^n$.]

7. Show that if a disk bundle $D^n \to E \to B$ has a Thom class with $\mathbb{Z}$ coefficients, then it is orientable.

8. If E is the product bundle $B \times D^n$ with B a CW complex, show that the Thom space $T(E)$ is the n-fold reduced suspension $\Sigma^n(B_+)$, where B_+ is the union of B with a disjoint basepoint, and that the Thom isomorphism specializes to the suspension isomorphism $\widetilde{H}^i(B;R) \approx \widetilde{H}^{n+i}(\Sigma^n B;R)$ given by the reduced cross product in §3.2.

9. Show that the inclusion $T^n \hookrightarrow U(n)$ of the n-torus of diagonal matrices is homotopic to the map $T^n \to U(1) \hookrightarrow U(n)$ sending an n-tuple of unit complex numbers $(z_1,\cdots,z_n)$ to the 1×1 matrix $(z_1 \cdots z_n)$. Do the same for the diagonal subgroup of $Sp(n)$. [Hint: Diagonal matrices in $U(n)$ are compositions of scalar multiplication in n lines in $\mathbb{C}^n$, and $\mathbb{C}P^{n-1}$ is connected.]

10. Fill in the details of the following argument to show that every $n\times n$ matrix A with entries in $\mathbb{H}$ has an eigenvalue in $\mathbb{H}$. (The usual argument over $\mathbb{C}$ involving roots of the characteristic polynomial does not work due to the lack of a good quaternionic determinant function.) For $t \in [0,1]$ and $\lambda \in S^3 \subset \mathbb{H}$, consider the matrix $t\lambda I + (1-t)A$. If A has no eigenvalues, this is invertible for all t. Thus the map $S^3 \to GL_n(\mathbb{H})$, $\lambda \mapsto \lambda I$, is nullhomotopic. But by the preceding problem and Exercise 10(b) in §3.C, this map represents n times a generator of $\pi_3 GL_n(\mathbb{H})$.

4.E The Brown Representability Theorem

In Theorem 4.58 in §4.3 we showed that Ω-spectra define cohomology theories, and now we will prove the converse statement that all cohomology theories on the CW category arise in this way from Ω-spectra.

Theorem 4E.1. *Every reduced cohomology theory on the category of basepointed CW complexes and basepoint-preserving maps has the form $h^n(X) = \langle X, K_n \rangle$ for some Ω-spectrum $\{K_n\}$.*

We will also see that the spaces K_n are unique up to homotopy equivalence.

This theorem gives another proof that ordinary cohomology is representable as maps into Eilenberg–MacLane spaces, since for the spaces K_n in an Ω-spectrum representing $\widetilde{H}^*(-;G)$ we have $\pi_i(K_n) = \langle S^i, K_n \rangle = \widetilde{H}^n(S^i;G)$, so K_n is a $K(G,n)$.

Before getting into the proof of the theorem let us observe that cofibration sequences, as constructed in §4.3, allow us to recast the definition of a reduced cohomology theory in a slightly more concise form: A reduced cohomology theory on the category $\mathcal{C}$ whose objects are CW complexes with a chosen basepoint 0-cell and whose morphisms are basepoint-preserving maps is a sequence of functors h^n, $n \in \mathbb{Z}$, from $\mathcal{C}$ to abelian groups, together with natural isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$ for all X in $\mathcal{C}$, such that the following axioms hold for each h^n:

(i) If $f \simeq g : X \to Y$ in the basepointed sense, then $f^* = g^* : h^n(Y) \to h^n(X)$.

(ii) For each inclusion $A \hookrightarrow X$ in $\mathcal{C}$ the sequence $h^n(X/A) \to h^n(X) \to h^n(A)$ is exact.

(iii) For a wedge sum $X = \bigvee_\alpha X_\alpha$ with inclusions $i_\alpha : X_\alpha \hookrightarrow X$, the product map $\prod_\alpha i_\alpha^* : h^n(X) \to \prod_\alpha h^n(X_\alpha)$ is an isomorphism.

To see that these axioms suffice to define a cohomology theory, the main thing to note is that the cofibration sequence $A \to X \to X/A \to \Sigma A \to \cdots$ allows us to construct the long exact sequence of a pair, just as we did in the case of the functors $h^n(X) = \langle X, K_n \rangle$. In the converse direction, if we have natural long exact sequences of pairs, then by applying these to pairs of the form (CX, X) we get natural isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$. Note that these natural isomorphisms coming from coboundary maps of pairs (CX, X) uniquely determine the coboundary maps for all pairs (X, A) via the diagram at the right, where the maps from $h^n(A)$ are coboundary maps of pairs and the diagram commutes by naturality of these coboundary maps. The isomorphism comes from a deformation retraction of CX onto CA. It is easy to check that these processes for converting one definition of a cohomology theory into the other are inverses of each other.

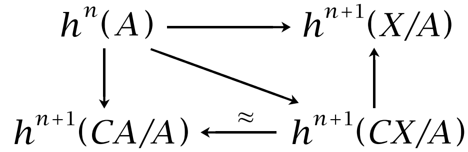

Most of the work in representing cohomology theories by Ω-spectra will be in realizing a single functor h^n of a cohomology theory as $\langle -, K_n \rangle$ for some space K_n.

So let us consider what properties the functor $h(X) = \langle X, K\rangle$ has, where K is a fixed space with basepoint. First of all, it is a contravariant functor from the category of basepointed CW complexes to the category of pointed sets, that is, sets with a distinguished element, the homotopy class of the constant map in the present case. Morphisms in the category of pointed sets are maps preserving the distinguished element. We have already seen in §4.3 that $h(X)$ satisfies the three axioms (i)–(iii). A further property is the following **Mayer–Vietoris axiom**:

- Suppose the CW complex X is the union of subcomplexes A and B containing the basepoint. Then if $a \in h(A)$ and $b \in h(B)$ restrict to the same element of $h(A\cap B)$, there exists an element $x \in h(X)$ whose restrictions to A and B are the given elements a and b.

Here and in what follows we use the term 'restriction' to mean the map induced by inclusion. In the case that $h(X) = \langle X, K\rangle$, this axiom is an immediate consequence of the homotopy extension property. The functors h^n in any cohomology theory also satisfy this axiom since there are Mayer–Vietoris exact sequences in any cohomology theory, as we observed in §2.3 in the analogous setting of homology theories.

Theorem 4E.2. *If h is a contravariant functor from the category of connected basepointed CW complexes to the category of pointed sets, satisfying the homotopy axiom (i), the Mayer–Vietoris axiom, and the wedge axiom (iii), then there exists a connected CW complex K and an element $u \in h(K)$ such that the transformation $T_u : \langle X, K\rangle \to h(X)$, $T_u(f) = f^*(u)$, is a bijection for all X.*

Such a pair (K, u) is called **universal** for the functor h. It is automatic from the definition that the space K in a universal pair (K, u) is unique up to homotopy equivalence. For if (K', u') is also universal for h, then, using the notation $f : (K, u) \to (K', u')$ to mean $f : K \to K'$ with $f^*(u') = u$, universality implies that there are maps $f : (K, u) \to (K', u')$ and $g : (K', u') \to (K, u)$ that are unique up to homotopy. Likewise the compositions $gf : (K, u) \to (K, u)$ and $fg : (K', u') \to (K', u')$ are unique up to homotopy, hence are homotopic to the identity maps.

Before starting the proof of this theorem we make a few preliminary comments on the axioms.

(1) The wedge axiom implies that $h(point)$ is trivial. To see this, just use the fact that for any X we have $X \vee point = X$, so the map $h(X) \times h(point) \to h(X)$ induced by inclusion of the first summand is a bijection, but this map is the projection $(a, b) \mapsto a$, hence $h(point)$ must have only one element.

(2) Axioms (i), (iii), and the Mayer–Vietoris axiom imply axiom (ii). Namely, (ii) is equivalent to exactness of $h(A) \leftarrow h(X) \leftarrow h(X \cup CA)$, where CA is the reduced cone since we are in the basepointed category. The inclusion Im ⊂ Ker holds since the composition $A \to X \cup CA$ is nullhomotopic, so the induced map factors through $h(point) = 0$.

To obtain the opposite inclusion Ker $\subset$ Im, decompose $X \cup CA$ into two subspaces Y and Z by cutting along a copy of A halfway up the cone CA, so Y is a smaller copy of CA and Z is the reduced mapping cylinder of the inclusion $A \hookrightarrow X$. Given an element $x \in h(X)$, this extends to an element $z \in h(Z)$ since Z deformation retracts to X. If x restricts to the trivial element of $h(A)$, then z restricts to the trivial element of $h(Y \cap Z)$. The latter element extends to the trivial element of $h(Y)$, so the Mayer–Vietoris axiom implies there is an element of $h(X \cup CA)$ restricting to z in $h(Z)$ and hence to x in $h(X)$.

(3) If h satisfies axioms (i) and (iii) then $h(\Sigma Y)$ is a group and $T_u : \langle \Sigma Y, K\rangle \to h(\Sigma Y)$ is a homomorphism for all suspensions ΣY and all pairs (K, u). [See the Corrections.]

The proof of Theorem 4E.2 will use two lemmas. To state the first, consider pairs (K, u) with K a basepointed connected CW complex and $u \in h(K)$, where h satisfies the hypotheses of the theorem. Call such a pair (K, u) ***n*-universal** if the homomorphism $T_u : \pi_i(K) \to h(S^i)$, $T_u(f) = f^*(u)$, is an isomorphism for $i < n$ and surjective for $i = n$. Call (K, u) π_*-**universal** if it is n-universal for all n.

Lemma 4E.3. *Given any pair (Z, z) with Z a connected CW complex and $z \in h(Z)$, there exists a π_*-universal pair (K, u) with Z a subcomplex of K and $u \mid Z = z$.*

Proof: We construct K from Z by an inductive process of attaching cells. To begin, let $K_1 = Z \bigvee_\alpha S^1_\alpha$ where α ranges over the elements of $h(S^1)$. By the wedge axiom there exists $u_1 \in h(K_1)$ with $u_1 \mid Z = z$ and $u_1 \mid S^1_\alpha = \alpha$, so (K_1, u_1) is 1-universal.

For the inductive step, suppose we have already constructed (K_n, u_n) with u_n n-universal, $Z \subset K_n$, and $u_n \mid Z = z$. Represent each element α in the kernel of $T_{u_n} : \pi_n(K_n) \to h(S^n)$ by a map $f_\alpha : S^n \to K_n$. Let $f = \bigvee_\alpha f_\alpha : \bigvee_\alpha S^n_\alpha \to K_n$. The reduced mapping cylinder M_f deformation retracts to K_n, so we can regard u_n as an element of $h(M_f)$, and this element restricts to the trivial element of $h(\bigvee_\alpha S^n_\alpha)$ by the definition of f. The exactness property of h then implies that for the reduced mapping cone $C_f = M_f / \bigvee_\alpha S^n_\alpha$ there is an element $w \in h(C_f)$ restricting to u_n on K_n. Note that C_f is obtained from K_n by attaching cells e^{n+1}_α by the maps f_α. To finish the construction of K_{n+1}, set $K_{n+1} = C_f \bigvee_\beta S^{n+1}_\beta$ where β ranges over $h(S^{n+1})$. By the wedge axiom, there exists $u_{n+1} \in h(K_{n+1})$ restricting to w on C_f and β on S^{n+1}_β.

To see that (K_{n+1}, u_{n+1}) is $(n+1)$-universal, consider the commutative diagram displayed at the right. Since K_{n+1} is obtained from K_n by attaching $(n+1)$-cells, the upper map is an isomorphism for $i < n$ and a surjection for $i = n$. By induction the same is true for T_{u_n}, hence it is also true for $T_{u_{n+1}}$. The kernel of $T_{u_{n+1}}$ is trivial for $i = n$ since an element of this kernel pulls back to $\mathrm{Ker}\, T_{u_n} \subset \pi_n(K_n)$, by surjectivity of the upper map when $i = n$, and we attached cells to K_n by maps representing all elements of $\mathrm{Ker}\, T_{u_n}$. Also, $T_{u_{n+1}}$ is surjective for $i = n + 1$ by construction.

$$\begin{array}{ccc} \pi_i(K_n) & \longrightarrow & \pi_i(K_{n+1}) \\ & {\scriptstyle T_{u_n}} \searrow \quad \swarrow {\scriptstyle T_{u_{n+1}}} & \\ & h(S^i) & \end{array}$$

Now let $K = \bigcup_n K_n$. We apply a mapping telescope argument as in the proofs of Lemma 2.34 and Theorem 3F.8 to show there is an element $u \in h(K)$ restricting to u_n on K_n, for all n. The mapping telescope of the inclusions $K_1 \hookrightarrow K_2 \hookrightarrow \cdots$ is the subcomplex $T = \bigcup_i K_i \times [i, i+1]$ of $K \times [1, \infty)$. We take '$\times$' to be the reduced product here, with *basepoint* $\times$ *interval* collapsed to a point. The natural projection $T \to K$ is a homotopy equivalence since $K \times [1, \infty)$ deformation retracts onto T, as we showed in the proof of Lemma 2.34. Let $A \subset T$ be the union of the subcomplexes $K_i \times [i, i+1]$ for i odd and let B be the corresponding union for i even. Thus $A \cup B = T$, $A \cap B = \bigvee_i K_i$, $A \simeq \bigvee_i K_{2i-1}$, and $B \simeq \bigvee_i K_{2i}$. By the wedge axiom there exist $a \in h(A)$ and $b \in h(B)$ restricting to u_i on each K_i. Then using the fact that $u_{i+1} | K_i = u_i$, the Mayer–Vietoris axiom implies that a and b are the restrictions of an element $t \in h(T)$. Under the isomorphism $h(T) \approx h(K)$, t corresponds to an element $u \in h(K)$ restricting to u_n on K_n for all n.

To verify that (K, u) is π_*-universal we use the commutative diagram at the right. For $n > i+1$ the upper map is an isomorphism and T_{u_n} is surjective with trivial kernel, so the same is true of T_u. □

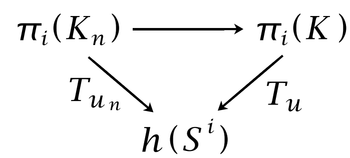

Lemma 4E.4. *Let (K, u) be a π_*-universal pair and let (X, A) be a basepointed CW pair. Then for each $x \in h(X)$ and each map $f : A \to K$ with $f^*(u) = x | A$ there exists a map $g : X \to K$ extending f with $g^*(u) = x$.*

Schematically, this is saying that the diagonal arrow in the diagram at the right always exists, where the map i is inclusion.

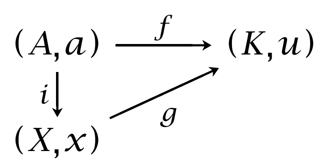

Proof: Replacing K by the reduced mapping cylinder of f reduces us to the case that f is the inclusion of a subcomplex. Let Z be the union of X and K with the two copies of A identified. By the Mayer–Vietoris axiom, there exists $z \in h(Z)$ with $z | X = x$ and $z | K = u$. By the previous lemma, we can embed (Z, z) in a π_*-universal pair (K', u'). The inclusion $(K, u) \hookrightarrow (K', u')$ induces an isomorphism on homotopy groups since both u and u' are π_*-universal, so K' deformation retracts onto K. This deformation retraction induces a homotopy rel A of the inclusion $X \hookrightarrow K'$ to a map $g : X \to K$. The relation $g^*(u) = x$ holds since $u' | K = u$ and $u' | X = x$. □

Proof of Theorem 4E.2: It suffices to show that a π_*-universal pair (K, u) is universal. Applying the preceding lemma with A a point shows that $T_u : \langle X, K \rangle \to h(X)$ is surjective. To show injectivity, suppose $T_u(f_0) = T_u(f_1)$, that is, $f_0^*(u) = f_1^*(u)$. We apply the preceding lemma with $(X \times I, X \times \partial I)$ playing the role of (X, A), using the maps f_0 and f_1 on $X \times \partial I$ and taking x to be $p^* f_0^*(u) = p^* f_1^*(u)$ where p is the projection $X \times I \to X$. Here $X \times I$ should be the reduced product, with *basepoint* $\times I$ collapsed to a point. The lemma then gives a homotopy from f_0 to f_1. □

Proof of Theorem 4E.1: Since suspension is an isomorphism in any reduced cohomology theory, and the suspension of any CW complex is connected, it suffices to restrict attention to connected CW complexes. Each functor h^n satisfies the homotopy, wedge, and Mayer–Vietoris axioms, as we noted earlier, so the preceding theorem gives CW complexes K_n with $h^n(X) = \langle X, K_n \rangle$. It remains to show that the natural isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$ correspond to weak homotopy equivalences $K_n \to \Omega K_{n+1}$. The natural isomorphism $h^n(X) \approx h^{n+1}(\Sigma X)$ corresponds to a natural bijection $\langle X, K_n \rangle \approx \langle \Sigma X, K_{n+1} \rangle = \langle X, \Omega K_{n+1} \rangle$ which we call Φ. The naturality of this bijection gives, for any map $f : X \to K_n$, a commutative diagram as at the right. Let $\varepsilon_n = \Phi(\mathbb{1}) : K_n \to \Omega K_{n+1}$. Then using commutativity we have $\Phi(f) = \Phi f^*(\mathbb{1}) = f^*\Phi(\mathbb{1}) = f^*(\varepsilon_n) = \varepsilon_n f$, which says that the map $\Phi : \langle X, K_n \rangle \to \langle X, \Omega K_{n+1} \rangle$ is composition with ε_n. Since Φ is a bijection, if we take X to be S^i, we see that ε_n induces an isomorphism on π_i for all i, so ε_n is a weak homotopy equivalence and we have an Ω-spectrum.

$$\begin{array}{ccc} \langle K_n, K_n \rangle & \xrightarrow{f^*} & \langle X, K_n \rangle \\ \big\downarrow \Phi & & \big\downarrow \Phi \\ \langle K_n, \Omega K_{n+1} \rangle & \xrightarrow{f^*} & \langle X, \Omega K_{n+1} \rangle \end{array}$$

There is one final thing to verify, that the bijection $h^n(X) = \langle X, K_n \rangle$ is a group isomorphism, where $\langle X, K_n \rangle$ has the group structure that comes from identifying it with $\langle X, \Omega K_{n+1} \rangle = \langle \Sigma X, K_{n+1} \rangle$. Via the natural isomorphism $h^n(X) \approx h^{n+1}(\Sigma X)$ this is equivalent to showing the bijection $h^{n+1}(\Sigma X) = \langle \Sigma X, K_{n+1} \rangle$ preserves group structure. For maps $f, g : \Sigma X \to K$, the relation $T_u(f+g) = T_u(f) + T_u(g)$ means $(f+g)^*(u) = f^*(u) + g^*(u)$, and this holds since $(f+g)^* = f^* + g^* : h(K) \to h(\Sigma X)$ by Lemma 4.60. □

4.F Spectra and Homology Theories

We have seen in §4.3 and the preceding section how cohomology theories have a homotopy-theoretic interpretation in terms of Ω-spectra, and it is natural to look for a corresponding description of homology theories. In this case we do not already have a homotopy-theoretic description of ordinary homology to serve as a starting point. But there is another homology theory we have encountered which does have a very homotopy-theoretic flavor:

Proposition 4F.1. *Stable homotopy groups $\pi_n^s(X)$ define a reduced homology theory on the category of basepointed CW complexes and basepoint-preserving maps.*

Proof: In the preceding section we reformulated the axioms for a cohomology theory so that the exactness axiom asserts just the exactness of $h^n(X/A) \to h^n(X) \to h^n(A)$ for CW pairs (X, A). In order to derive long exact sequences, the reformulated axioms require also that natural suspension isomorphisms $h^n(X) \approx h^{n+1}(\Sigma X)$ be specified as part of the cohomology theory. The analogous reformulation of the axioms for a homology theory is valid as well, by the same argument, and we shall use this in what follows.

For stable homotopy groups, suspension isomorphisms $\pi_n^s(X) \approx \pi_{n+1}^s(\Sigma X)$ are automatic, so it remains to verify the three axioms. The homotopy axiom is apparent. The exactness of a sequence $\pi_n^s(A) \to \pi_n^s(X) \to \pi_n^s(X/A)$ follows from exactness of $\pi_n(A) \to \pi_n(X) \to \pi_n(X,A)$ together with the isomorphism $\pi_n(X,A) \approx \pi_n(X/A)$ which holds under connectivity assumptions that are achieved after sufficiently many suspensions. The wedge sum axiom $\pi_n^s(\bigvee_\alpha X_\alpha) \approx \bigoplus_\alpha \pi_n^s(X_\alpha)$ reduces to the case of finitely many summands by the usual compactness argument, and the case of finitely many summands reduces to the case of two summands by induction. Then we have isomorphisms $\pi_{n+i}(\Sigma^i X \vee \Sigma^i Y) \approx \pi_{n+i}(\Sigma^i X \times \Sigma^i Y) \approx \pi_{n+i}(\Sigma^i X) \bigoplus \pi_{n+i}(\Sigma^i Y)$, the first of these isomorphisms holding when $n + i < 2i - 1$, or $i > n + 1$, since $\Sigma^i X \vee \Sigma^i Y$ is the $(2i-1)$-skeleton of $\Sigma^i X \times \Sigma^i Y$. Passing to the limit over increasing i, we get the desired isomorphism $\pi_n^s(X \vee Y) \approx \pi_n^s(X) \bigoplus \pi_n^s(Y)$. ☐

A modest generalization of this homology theory can be obtained by defining $h_n(X) = \pi_n^s(X \wedge K)$ for a fixed complex K. Verifying the homology axioms reduces to the case of stable homotopy groups themselves by basic properties of smash product:

- $h_n(X) \approx h_{n+1}(\Sigma X)$ since $\Sigma(X \wedge K) = (\Sigma X) \wedge K$, both spaces being $S^1 \wedge X \wedge K$.
- The exactness axiom holds since $(X \wedge K)/(A \wedge K) = (X/A) \wedge K$, both spaces being quotients of $X \times K$ with $A \times K \cup X \times \{k_0\}$ collapsed to a point.
- The wedge axiom follows from distributivity: $(\bigvee_\alpha X_\alpha) \wedge K = \bigvee_\alpha (X_\alpha \wedge K)$.

The coefficients of this homology theory are $h_n(S^0) = \pi_n^s(S^0 \wedge K) = \pi_n^s(K)$. Suppose for example that K is an Eilenberg–MacLane space $K(G,n)$. Because $K(G,n)$ is $(n-1)$-connected, its stable homotopy groups are the same as its unstable homotopy groups below dimension $2n$. Thus if we shift dimensions by defining $h_i(X) = \pi_{i+n}^s(X \wedge K(G,n))$ we obtain a homology theory whose coefficient groups below dimension n are the same as ordinary homology with coefficients in G. It follows as in Theorem 4.59 that this homology theory agrees with ordinary homology for CW complexes of dimension less than $n-1$.

This dimension restriction could be removed if there were a 'stable Eilenberg–MacLane space' whose stable homotopy groups were zero except in one dimension. However, this is a lot to ask for, so instead one seeks to form a limit of the groups $\pi_{i+n}^s(X \wedge K(G,n))$ as n goes to infinity. The spaces $K(G,n)$ for varying n are related by weak homotopy equivalences $K(G,n) \to \Omega K(G,n+1)$. Since suspension plays such a large role in the current discussion, let us consider instead the corresponding map $\Sigma K(G,n) \to K(G,n+1)$, or to write this more concisely, $\Sigma K_n \to K_{n+1}$ This induces a map $\pi_{i+n}^s(X \wedge K_n) = \pi_{i+n+1}^s(X \wedge \Sigma K_n) \to \pi_{i+n+1}^s(X \wedge K_{n+1})$. Via these maps, it then makes sense to consider the direct limit as n goes to infinity, the group $h_i(X) = \varinjlim \pi^s_{i+n}(X \wedge K_n)$. This gives a homology theory since direct limits preserve exact sequences so the exactness axiom holds, and direct limits preserve isomorphisms so the suspension isomorphism and the wedge axiom hold. The coefficient groups of this homology theory are the same as for ordinary homology with G coefficients since $h_i(S^0) = \varinjlim \pi^s_{i+n}(K_n)$ is zero unless $i = 0$, when it is G. Hence this homology theory coincides with ordinary homology by Theorem 4.59.

To place this result in its natural generality, define a **spectrum** to be a sequence of CW complexes K_n together with basepoint-preserving maps $\Sigma K_n \to K_{n+1}$. This generalizes the notion of an Ω-spectrum, where the maps $\Sigma K_n \to K_{n+1}$ come from weak homotopy equivalences $K_n \to \Omega K_{n+1}$. Another obvious family of examples is **suspension spectra**, where one starts with an arbitrary CW complex X and defines $K_n = \Sigma^n X$ with $\Sigma K_n \to K_{n+1}$ the identity map.

The homotopy groups of a spectrum K are defined to be $\pi_i(K) = \varinjlim \pi_{i+n}(K_n)$ where the direct limit is computed using the compositions

$$\pi_{i+n}(K_n) \xrightarrow{\Sigma} \pi_{i+n+1}(\Sigma K_n) \longrightarrow \pi_{i+n+1}(K_{n+1})$$

with the latter map induced by the given map $\Sigma K_n \to K_{n+1}$. Thus in the case of the suspension spectrum of a space X, the homotopy groups of the spectrum are the same as the stable homotopy groups of X. For a general spectrum K we could also describe $\pi_i(K)$ as $\varinjlim \pi^s_{i+n}(K_n)$ since the composition $\pi_{i+n}(K_n) \to \pi_{i+n+j}(K_{n+j})$ factors through $\pi_{i+n+j}(\Sigma^j K_n)$. So the homotopy groups of a spectrum are 'stable homotopy groups' essentially by definition.

Returning now to the context of homology theories, if we are given a spectrum K and a CW complex X, then we have a spectrum $X \wedge K$ with $(X \wedge K)_n = X \wedge K_n$, using the obvious maps $\Sigma(X \wedge K_n) = X \wedge \Sigma K_n \to X \wedge K_{n+1}$. The groups $\pi_i(X \wedge K)$ are the groups $\varinjlim \pi^s_{i+n}(X \wedge K_n)$ considered earlier in the case of an Eilenberg–MacLane spectrum, and the arguments given there show:

Proposition 4F.2. *For a spectrum K, the groups $h_i(X) = \pi_i(X \wedge K)$ form a reduced homology theory. When K is the Eilenberg–MacLane spectrum with $K_n = K(G, n)$, this homology theory is ordinary homology, so $\pi_i(X \wedge K) \approx \widetilde{H}_i(X; G)$.* □

If one wanted to associate a cohomology theory to an arbitrary spectrum K, one's first inclination would be to set $h^i(X) = \varinjlim \langle \Sigma^n X, K_{n+i} \rangle$, the direct limit with respect to the compositions

$$\langle \Sigma^n X, K_{n+i} \rangle \xrightarrow{\Sigma} \langle \Sigma^{n+1} X, \Sigma K_{n+i} \rangle \longrightarrow \langle \Sigma^{n+1} X, K_{n+i+1} \rangle$$

For example, in the case of the sphere spectrum $S = \{S^n\}$ this definition yields the **stable cohomotopy groups** $\pi^i_s(X) = \varinjlim \langle \Sigma^n X, S^{n+i} \rangle$. Unfortunately the definition $h^i(X) = \varinjlim \langle \Sigma^n X, K_{n+i} \rangle$ runs into problems with the wedge sum axiom since the direct limit of a product need not equal the product of the direct limits. For finite wedge sums there is no difficulty, so we do have a cohomology theory for finite CW complexes. But for general CW complexes a different definition is needed. The simplest thing to do is to associate to each spectrum K an Ω-spectrum K' and let $h^n(X) = \langle X, K'_n \rangle$. We obtain K' from K by setting $K'_n = \varinjlim \Omega^i K_{n+i}$, the mapping telescope of the sequence $K_n \to \Omega K_{n+1} \to \Omega^2 K_{n+2} \to \cdots$. The Ω-spectrum structure is given by equivalences

$$K'_n = \varinjlim \Omega^i K_{n+i} \simeq \varinjlim \Omega^{i+1} K_{n+i+1} \xrightarrow{\kappa} \Omega \varinjlim \Omega^i K_{n+i+1} = \Omega K'_{n+1}$$

The first homotopy equivalence comes from deleting the first term of the sequence $K_n \to \Omega K_{n+1} \to \Omega^2 K_{n+2} \to \cdots$, which has negligible effect on the mapping telescope. The next map κ is a special case of the natural weak equivalence $\varinjlim \Omega Z_n \to \Omega \varinjlim Z_n$ that holds for any sequence $Z_1 \to Z_2 \to \cdots$. Strictly speaking, we should let K'_n be a CW approximation to the mapping telescope $\varinjlim \Omega^i K_{n+i}$ in order to obtain a spectrum consisting of CW complexes, in accordance with our definition of a spectrum.

In case one starts with a suspension spectrum $K_n = \Sigma^n K$ it is not necessary to take mapping telescopes since one can just set $K'_n = \bigcup_i \Omega^i \Sigma^{i+n} K = \bigcup_i \Omega^i \Sigma^i K_n$, the union with respect to the natural inclusions $\Omega^i \Sigma^i K_n \subset \Omega^{i+1} \Sigma^{i+1} K_n$. The union $\bigcup_i \Omega^i \Sigma^i X$ is usually abbreviated to $\Omega^\infty \Sigma^\infty X$. Another common notation for this union is QX. Thus $\pi_i(QX) = \pi_i^s(X)$, so Q is a functor converting stable homotopy groups into ordinary homotopy groups.

It follows routinely from the definitions that the homology theory defined by a spectrum is the same as the homology theory defined by the associated Ω-spectrum. One may ask whether every homology theory is defined by a spectrum, as we showed for cohomology. The answer is yes if one replaces the wedge axiom by a stronger **direct limit axiom**: $h_i(X) = \varinjlim h_i(X_\alpha)$, the direct limit over the finite subcomplexes X_α of X. The homology theory defined by a spectrum satisfies this axiom, and the converse is proved in [Adams 1971].

Spectra have become the preferred language for describing many stable phenomena in algebraic topology. The increased flexibility of spectra is not without its price, however, since a number of concepts that are elementary for spaces become quite a bit more subtle for spectra, such as the proper definition of a map between spectra, or the smash product of two spectra. For the reader who wants to learn more about this language a good starting point is [Adams 1974].

Exercises

1. Assuming the first two axioms for a homology theory on the CW category, show that the direct limit axiom implies the wedge sum axiom. Show that the converse also holds for countable CW complexes.

2. For CW complexes X and Y consider the suspension sequence

$$\langle X, Y \rangle \xrightarrow{\Sigma} \langle \Sigma X, \Sigma Y \rangle \xrightarrow{\Sigma} \langle \Sigma^2 X, \Sigma^2 Y \rangle \longrightarrow \cdots$$

Show that if X is a finite complex, these maps eventually become isomorphisms. [Use induction on the number of cells of X and the five-lemma.]

3. Show that for any sequence $Z_1 \to Z_2 \to \cdots$, the natural map $\varinjlim \Omega Z_n \to \Omega \varinjlim Z_n$ is a weak homotopy equivalence, where the direct limits mean mapping telescopes.

4.G Gluing Constructions

It is a common practice in algebraic topology to glue spaces together to form more complicated spaces. In this section we describe two general procedures for making such constructions. The first is fairly straightforward but also rather rigid, lacking some homotopy invariance properties an algebraic topologist would like to see. The second type of gluing construction avoids these drawbacks by systematic use of mapping cylinders. We have already seen many special cases of both types of constructions, and having a general framework covering all these special cases should provide some conceptual clarity.

A **diagram of spaces** consists of an oriented graph Γ with a space X_v for each vertex v of Γ and a map $f_e : X_v \to X_w$ for each edge e of Γ from a vertex v to a vertex w, the words 'from' and 'to' referring to the given orientation of e. Commutativity of the diagram is not assumed. Denoting such a diagram of spaces simply by X, we define a space $\sqcup X$ to be the quotient of the disjoint union of all the spaces X_v associated to vertices of Γ under the identifications $x \sim f_e(x)$ for all maps f_e associated to edges of Γ. To give a name to this construction, let us call $\sqcup X$ the *amalgamation* of the diagram X. Here are some examples:

- If the diagram of spaces has the simple form $X_0 \xleftarrow{f} A \hookrightarrow X_1$ then $\sqcup X$ is the space $X_0 \sqcup_f X_1$ obtained from X_0 by attaching X_1 along A via f.
- A sequence of inclusions $X_0 \hookrightarrow X_1 \hookrightarrow \cdots$ determines a diagram of spaces X for which $\sqcup X$ is $\bigcup_i X_i$ with the weak topology. This holds more generally when the spaces X_i are indexed by any directed set.
- From a cover $\mathcal{U} = \{X_i\}$ of a space X by subspaces X_i we can form a diagram of spaces $X_{\mathcal{U}}$ whose vertices are the nonempty finite intersections $X_{i_1} \cap \cdots \cap X_{i_n}$ with distinct indices i_j, and whose edges are the various inclusions obtained by omitting some of the subspaces in such an intersection, for example the inclusions $X_i \cap X_j \hookrightarrow X_i$. Then $\sqcup X_{\mathcal{U}}$ equals X as a set, though possibly with a different topology. If the cover is an open cover, or if X is a CW complex and the X_i's are subcomplexes, then the topology will be the original topology on X.
- An action of a group G on a space X determines a diagram of spaces X_G, with X itself as the only space and with maps the homeomorphisms $g : X \to X$, $g \in G$, given by the action. In this case $\sqcup X_G$ is the orbit space X/G.

- A Δ-complex X can be viewed as a diagram of spaces X_Δ where each simplex of X gives a vertex space X_v which is a simplex of the same dimension, and the edge maps are the inclusions of faces into the simplices that contain them. Then $\sqcup X_\Delta = X$.

It can very easily happen that for a diagram of spaces X the amalgamation $\sqcup X$ is rather useless because so much collapsing has occurred that little of the original diagram remains. For example, consider a diagram X of the form $X_0 \leftarrow X_0 \times X_1 \rightarrow X_1$ whose maps are the projections onto the two factors. In this case $\sqcup X$ is simply a point. To correct for problems like this, and to get a notion with nicer homotopy-theoretic properties, we introduce the homotopy version of $\sqcup X$, which we shall denote ΔX and call the *realization* of X. Here we again start with the disjoint union of all the vertex spaces X_v, but instead of passing to a quotient space of this disjoint union, we enlarge it by filling in a mapping cylinder M_f for each map f of the diagram, identifying the two ends of this cylinder with the appropriate X_v's. In the case of the projection diagram $X_0 \leftarrow X_0 \times X_1 \rightarrow X_1$, the union of the two mapping cylinders is the same as the quotient of $X_0 \times X_1 \times I$ with $X_0 \times X_1 \times \{0\}$ collapsed to X_0 and $X_0 \times X_1 \times \{1\}$ collapsed to X_1. Thus ΔX is the join $X_0 * X_1$ defined in Chapter 0.

We have seen a number of other special cases of the construction ΔX. For a diagram consisting of just one map $f : X_0 \rightarrow X_1$ one gets of course the mapping cylinder M_f itself. For a diagram $X_0 \xleftarrow{f} X_1 \xrightarrow{g} X_2$ the realization ΔX is a double mapping cylinder. In case X_2 is a point this is the mapping cone of f. When the diagram has just one space and one map from this space to itself, then ΔX is the mapping torus. For a diagram consisting of two maps $f, g : X_0 \rightarrow X_1$ the space ΔX was studied in Example 2.48. Mapping telescopes are the case of a sequence of maps $X_0 \rightarrow X_1 \rightarrow \cdots$. In §1.B we considered general diagrams in which the spaces are $K(G,1)$'s.

There is a natural generalization of ΔX in which one starts with a Δ-complex Γ and a diagram of spaces associated to the 1-skeleton of Γ such that the maps corresponding to the edges of each n-simplex of Γ, $n > 1$, form a commutative diagram. We call this data a **complex of spaces**. If X is a complex of spaces, then for each n-simplex of Γ we have a sequence of maps $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots \xrightarrow{f_n} X_n$, and we define the *iterated mapping cylinder* $M(f_1, \cdots, f_n)$ to be the usual mapping cylinder for $n = 1$, and inductively for $n > 1$, the mapping cylinder of the composition $M(f_1, \cdots, f_{n-1}) \rightarrow X_{n-1} \xrightarrow{f_n} X_n$ where the first map is the canonical projection of a mapping cylinder onto its target end. There is a natural projection $M(f_1, \cdots, f_n) \rightarrow \Delta^n$, and over each face of Δ^n one has the iterated mapping cylinder for the maps associated to the edges in this face. For example when $n = 2$ one has the three mapping cylinders $M(f_1)$, $M(f_2)$, and $M(f_2 f_1)$ over the three edges of Δ^2. All these iterated mapping cylinders over the various simplices of Γ thus fit together to form a space ΔX with a canonical projection $\Delta X \rightarrow \Gamma$. We again call ΔX the realization of the complex of spaces X, and we call Γ the *base* of X or ΔX.

Some of our earlier examples of diagrams of spaces can be regarded in a natural way as complexes of spaces:

- For a cover $\mathcal{U} = \{X_i\}$ of a space X the diagram of spaces $X_{\mathcal{U}}$ whose vertices are the finite intersections of X_i's and whose edges are inclusions is a complex of spaces with n-simplices the n-fold inclusions. The base Γ for this complex of spaces is the barycentric subdivision of the nerve of the cover. Recall from the end of §3.3 that the nerve of a cover is the simplicial complex with n-simplices the nonempty $(n+1)$-fold intersections of sets in the cover.
- The diagram of spaces X_G associated to an action of a group G on a space X is a complex of spaces, with n-simplices corresponding to the n-fold compositions $X \xrightarrow{g_1} X \xrightarrow{g_2} \cdots \xrightarrow{g_n} X$. The base Δ-complex Γ is the $K(G,1)$ called BG in §1.B. This was the orbit space of a free action of G on a contractible Δ-complex EG. Checking through the definitions, one sees that the space ΔX_G in this case can be regarded as the quotient of $X \times EG$ under the diagonal action of G, $g(x,y) = (g(x), g(y))$. This is the space we called the Borel construction in §3.G, with the notation $X \times_G EG$.

By a map $f : X \to Y$ of complexes of spaces over the same base Γ we mean a collection of maps $f_v : X_v \to Y_v$ for all the vertices of Γ, with commutative squares over all edges of Γ. There is then an induced map $\Delta f : \Delta X \to \Delta Y$.

Proposition 4G.1. *If all the maps f_v making up a map of complexes of spaces $f : X \to Y$ are homotopy equivalences, then so is the map $\Delta f : \Delta X \to \Delta Y$.*

Proof: The mapping cylinders $M(f_v)$ form a complex of spaces $M(f)$ over the same base Γ, and the space $\Delta M(f)$ is the mapping cylinder $M(\Delta f)$. This deformation retracts onto ΔY, so it will suffice to show that it also deformation retracts onto ΔX.

Let $M^n(\Delta f)$ be the part of $M(\Delta f)$ lying over the n-skeleton of Γ. We claim that $M^n(\Delta f) \cup \Delta X$ deformation retracts onto $M^{n-1}(\Delta f) \cup \Delta X$. It is enough to show this when $\Gamma = \Delta^n$. In this case f is a map from $X_0 \to \cdots \to X_n$ to $Y_0 \to \cdots \to Y_n$. By Corollary 0.20 it suffices to show that the inclusion $M^{n-1}(\Delta f) \cup \Delta X \hookrightarrow M(\Delta f)$ is a homotopy equivalence and the pair $(M(\Delta f), M^{n-1}(\Delta f) \cup \Delta X)$ satisfies the homotopy extension property. The latter assertion is evident from Example 0.15 since a mapping cylinder neighborhood is easily constructed for this pair. For the other condition, note that by induction on the dimension of Γ we may assume that $M^{n-1}(\Delta f)$ deformation retracts onto the part of ΔX over $\partial \Delta^n$. Also, the inclusion $\Delta X \hookrightarrow M(\Delta f)$ is a homotopy equivalence since it is equivalent to the map $X_n \to Y_n$, which is a homotopy equivalence by hypothesis. So Corollary 0.20 applies, and the claim that $M^n(\Delta f) \cup \Delta X$ deformation retracts onto $M^{n-1}(\Delta f) \cup \Delta X$ is proved.

Letting n vary, the infinite concatenation of these deformation retractions in the t-intervals $[1/2^{n+1}, 1/2^n]$ gives a deformation retraction of $M(\Delta f)$ onto ΔX. $\square$

There is a canonical map $\Delta X \to \sqcup X$ induced by retracting each mapping cylinder onto its target end. In some cases this is a homotopy equivalence, for example, for a diagram $X_0 \leftarrow A \hookrightarrow X_1$ where the pair (X_1, A) has the homotopy extension property. Another example is a sequence of inclusions $X_0 \hookrightarrow X_1 \hookrightarrow \cdots$ for which the pairs (X_n, X_{n-1}) satisfy the homotopy extension property, by the argument involving mapping telescopes in the proof of Lemma 2.34. However, without some conditions on the maps it need not be true that $\Delta X \to \sqcup X$ is a homotopy equivalence, as the earlier example of the projections $X_0 \leftarrow X_0 \times X_1 \to X_1$ shows. Even with inclusion maps one need not have $\Delta X \simeq \sqcup X$ if the base Γ is not contractible. A trivial example is the diagram consisting of the two spaces Δ^0 and Δ^1 and two maps $f_0, f_1 : \Delta^0 \to \Delta^1$ that happen to have the same image.

Thus one can expect the map $\Delta X \to \sqcup X$ to be a homotopy equivalence only in special circumstances. Here is one such situation:

Proposition 4G.2. *When $X_{\mathcal{U}}$ is the complex of spaces associated to an open cover $\mathcal{U} = \{X_i\}$ of a paracompact space X, the map $p : \Delta X_{\mathcal{U}} \to \sqcup X_{\mathcal{U}} = X$ is a homotopy equivalence.*

Proof: The realization $\Delta X_{\mathcal{U}}$ can also be described as the quotient space of the disjoint union of all the products $X_{i_0} \cap \cdots \cap X_{i_n} \times \Delta^n$, as the subscripts range over sets of $n+1$ distinct indices and $n \geq 0$, with the identifications over the faces of Δ^n using inclusions $X_{i_0} \cap \cdots \cap X_{i_n} \hookrightarrow X_{i_0} \cap \cdots \cap \widehat{X}_{i_j} \cap \cdots \cap X_{i_n}$. From this viewpoint, points of $\Delta X_{\mathcal{U}}$ in a given 'fiber' $p^{-1}(x)$ can be written as finite linear combinations $\sum_i t_i x_i$ where $\sum_i t_i = 1$ and x_i is x regarded as a point of X_i, for those X_i's that contain x.

Since X is paracompact there is a partition of unity subordinate to the cover $\mathcal{U}$. This is a family of maps $\varphi_\alpha : X \to [0,1]$ satisfying three conditions: The support of each φ_α is contained in some $X_{i(\alpha)}$, only finitely many φ_α's are nonzero near each point of X, and $\sum_\alpha \varphi_\alpha = 1$. Define a section $s : X \to \Delta X_{\mathcal{U}}$ of p by setting $s(x) = \sum_\alpha \varphi_\alpha(x) x_{i(\alpha)}$. The figure shows the case $X = S^1$ with a cover by two arcs, the heavy line indicating the image of s. In the general case the section s embeds X as a retract of $\Delta X_{\mathcal{U}}$, and it is a deformation retract since points in fibers $p^{-1}(x)$ can move linearly along line segments to $s(x)$. ☐

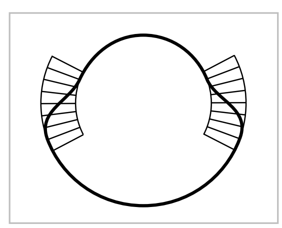

Corollary 4G.3. *If $\mathcal{U}$ is an open cover of a paracompact space X such that every nonempty intersection of finitely many sets in $\mathcal{U}$ is contractible, then X is homotopy equivalent to the nerve $N\mathcal{U}$.*

Proof: The proposition gives a homotopy equivalence $X \simeq \Delta X_{\mathcal{U}}$. Since the nonempty finite intersections of sets in $\mathcal{U}$ are contractible, the earlier proposition implies that the map $\Delta X_{\mathcal{U}} \to \Gamma$ induced by sending each intersection to a point is a homotopy equivalence. Since Γ is the barycentric subdivision of $N\mathcal{U}$, the result follows. ☐

Let us conclude this section with a few comments about terminology. For some diagrams of spaces such as sequences $X_1 \to X_2 \to \cdots$ the amalgamation $\sqcup X$ can be regarded as the direct limit of the vertex spaces X_v with respect to the edge maps f_e. Following this cue, the space $\sqcup X$ is commonly called the direct limit for arbitrary diagrams, even finite ones. If one views $\sqcup X$ as a direct limit, then ΔX becomes a sort of homotopy direct limit. For reasons that are explained in the next section, direct limits are often called 'colimits'. This has given rise to the rather unfortunate name of 'hocolim' for ΔX, short for 'homotopy colimit'. In preference to this we have chosen the term 'realization', both for its intrinsic merits and because ΔX is closely related to what is called the geometric realization of a simplicial space.

Exercises

1. Show that for a sequence of maps $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots$, the infinite iterated mapping cylinder $M(f_1, f_2, \cdots)$, which is the union of the finite iterated mapping cylinders $M(f_1, \cdots, f_n)$, deformation retracts onto the mapping telescope.

2. Show that if X is a complex of spaces in which all the maps are homeomorphisms, then the projection $\Delta X \to \Gamma$ is a fiber bundle.

3. What is the nerve of the cover of a simplicial complex by the open stars of its vertices? [See Lemma 2C.2.]

4. Show that Proposition 4G.2 and its corollary hold also for CW complexes and covers by families of subcomplexes. [CW complexes are paracompact; see [VBKT].]

4.H Eckmann-Hilton Duality

There is a very nice duality principle in homotopy theory, called Eckmann–Hilton duality in its more refined and systematic aspects, but which in its most basic form involves the simple idea of reversing the direction of all arrows in a given construction. For example, if in the definition of a fibration as a map satisfying the homotopy lifting property we reverse the direction of all the arrows, we obtain the dual notion of a **cofibration**. This is a map $i: A \to B$ satisfying the following property: Given $\widetilde{g}_0: B \to X$ and a homotopy $g_t: A \to X$ such that $\widetilde{g}_0 i = g_0$, there exists a homotopy $\widetilde{g}_t: B \to X$ such that $\widetilde{g}_t i = g_t$. In the special case that i is the inclusion of a subspace, this is the homotopy extension property, and the next proposition says that this is indeed the general case. So a cofibration is the same as an inclusion satisfying the homotopy extension property.

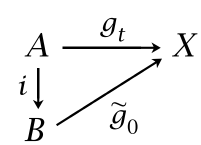

Proposition 4H.1. *If $i: A \to B$ is a cofibration, then i is injective, and in fact a homeomorphism onto its image.*

Proof: Consider the mapping cylinder M_i, the quotient of $A\times I \amalg B$ in which $(a,1)$ is identified with $i(a)$. Let $g_t : A \to M_i$ be the homotopy mapping $a \in A$ to the image of $(a, 1-t) \in A\times I$ in M_i, and let $\widetilde{g}_0$ be the inclusion $B \hookrightarrow M_i$. The cofibration property gives $\widetilde{g}_t : B \to M_i$ with $\widetilde{g}_t i = g_t$. Restricting to a fixed $t > 0$, this implies i is injective since g_t is. Furthermore, since g_t is a homeomorphism onto its image $A\times\{1-t\}$, the relation $\widetilde{g}_t i = g_t$ implies that the map $g_t^{-1}\widetilde{g}_t : i(A) \to A$ is a continuous inverse of $i : A \to i(A)$. ☐

Many constructions for fibrations have analogs for cofibrations, and vice versa. For example, for an arbitrary map $f : A \to B$ the inclusion $A \hookrightarrow M_f$ is readily seen to be a cofibration, so the analog of the factorization $A \hookrightarrow E_f \to B$ of f into a homotopy equivalence followed by a fibration is the factorization $A \hookrightarrow M_f \to B$ into a cofibration followed by a homotopy equivalence. Even the definition of M_f is in some way dual to the definition of E_f, since E_f can be defined as a pullback and M_f can be defined as a dual pushout. In general, the **pushout** of maps $f : Z \to X$ and $g : Z \to Y$ is defined as the quotient of $X \amalg Y$ under the identifications $f(z) \sim g(z)$.

$$\begin{array}{ccc} E_f & \longrightarrow & B^I \\ \downarrow & & \downarrow \\ A & \longrightarrow & B \end{array} \qquad\qquad \begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ A\times I & \longrightarrow & M_f \end{array}$$

Thus the pushout is a quotient of $X \amalg Y$, while the pullback of maps $X \to Z$ and $Y \to Z$ is a subobject of $X\times Y$, so we see here two instances of duality: a duality between disjoint union and product, and a duality between subobjects and quotients. The first of these is easily explained, since a collection of maps $X_\alpha \to X$ is equivalent to a map $\coprod_\alpha X_\alpha \to X$, while a collection of maps $X \to X_\alpha$ is equivalent to a map $X \to \prod_\alpha X_\alpha$. The notation $\coprod$ for the 'coproduct' was chosen to indicate that it is dual to $\prod$. If we were dealing with basepointed spaces and maps, the coproduct would be wedge sum. In the category of abelian groups the coproduct is direct sum.

The duality between subobjects and quotient objects is clear for abelian groups, where subobjects are kernels and quotient objects cokernels. The strict topological analog of a kernel is a fiber of a fibration. Dually, the topological analog of a cokernel is the **cofiber** B/A of a cofibration $A \hookrightarrow B$. If we make an arbitrary map $f : A \to B$ into a cofibration $A \hookrightarrow M_f$, the cofiber is the mapping cone $C_f = M_f/(A\times\{0\})$.

In the diagram showing E_f and M_f as pullback and pushout, there also appears to be some sort of duality involving the terms $A\times I$ and B^I. This leads us to ask whether $X\times I$ and X^I are in some way dual. Indeed, if we ignore topology and just think set-theoretically, this is an instance of the familiar product–coproduct duality since the product of copies of X indexed by I is X^I, all functions $I \to X$, while the coproduct of copies of X indexed by I is $X\times I$, the disjoint union of the sets $X\times\{t\}$ for $t \in I$. Switching back from the category of sets to the topological category, we can view X^I as a 'continuous product' of copies of X and $X\times I$ as a 'continuous coproduct'.

On a less abstract level, the fact that maps $A\times I \to B$ are the same as maps $A \to B^I$ indicates a certain duality between $A\times I$ and B^I. This leads to a duality between suspension and loopspace, since ΣA is a quotient of $A\times I$ and ΩB is a subspace of B^I. This duality is expressed in the adjoint relation $\langle \Sigma X, Y\rangle = \langle X, \Omega Y\rangle$ from §4.3. Combining this duality between Σ and Ω with the duality between fibers and cofibers, we see a duality relationship between the fibration and cofibration sequences of §4.3:

$$\cdots \longrightarrow \Omega F \longrightarrow \Omega E \longrightarrow \Omega B \longrightarrow F \longrightarrow E \longrightarrow B$$

$$A \to X \to X/A \to \Sigma A \to \Sigma X \to \Sigma(X/A) \to \cdots$$

Pushout and pullback constructions can be generalized to arbitrary diagrams. In the case of pushouts, this was done in §4.G where we associated a space $\sqcup X$ to a diagram of spaces X. This was the quotient of the coproduct $\coprod_v X_v$, with v ranging over vertices of the diagram, under the identifications $x \sim f_e(x)$ for all maps f_e associated to edges e of the diagram. The dual construction $\sqcap X$ would be the subspace of the product $\prod_v X_v$ consisting of tuples (x_v) with $f_e(x_v) = x_w$ for all maps $f_e : X_v \to X_w$ in the diagram. Perhaps more useful in algebraic topology is the homotopy variant of this notion obtained by dualizing the definition of ΔX in the previous section. This is the space ∇X consisting of all choices of a point x_v in each X_v and a path γ_e in the target space of each edge map $f_e : X_v \to X_w$, with $\gamma_e(0) = f(x_v)$ and $\gamma_e(1) = x_w$. The subspace with all paths constant is $\sqcap X$. In the case of a diagram $\cdots \to X_2 \to X_1$ such as a Postnikov tower this construction gives something slightly different from simply turning each successive map into a fibration via the usual pathspace construction, starting with $X_2 \to X_1$ and proceeding up the tower, as we did in §4.3. The latter construction is rather the dual of an iterated mapping cylinder, involving spaces of maps $\Delta^n \to X_v$ instead of simply pathspaces. One could use such mapping spaces to generalize the definition of ∇X from diagrams of spaces to complexes of spaces.

As special cases of the constructions $\sqcup X$ and $\sqcap X$ we have direct limits and inverse limits for diagrams $X_0 \to X_1 \to \cdots$ and $\cdots \to X_1 \to X_0$, respectively. Since inverse limit is related to product and direct limit to coproduct, it is common practice in some circles to use reverse logic and call inverse limit simply 'limit' and direct limit 'colimit'. The homotopy versions are then called 'holim' for ∇X and 'hocolim' for ΔX. This terminology is frequently used for more general diagrams as well.

Homotopy Groups with Coefficients

There is a somewhat deeper duality between homotopy groups and cohomology, which one can see in the fact that cohomology groups are homotopy classes of maps into a space with a single nonzero homotopy group, while homotopy groups are homotopy classes of maps from a space with a single nonzero cohomology group. This duality is in one respect incomplete, however, in that the cohomology statement holds for an arbitrary coefficient group, but we have not yet defined homotopy groups with coefficients. In view of the duality, one would be tempted to define $\pi_n(X;G)$ to be the set of basepoint-preserving homotopy classes of maps from the cohomology analog of a Moore space $M(G,n)$ to X. The cohomology analog of $M(G,n)$ would be a space

Y whose only nonzero cohomology group $\widetilde{H}^i(Y;\mathbb{Z})$ is G for $i = n$. Unfortunately, such a space does not exist for arbitrary G, for example for $G = \mathbb{Q}$, since we showed in Proposition 3F.12 that if the cohomology groups of a space are all countable, then they are all finitely generated.

As a first approximation to $\pi_n(X;G)$ let us consider $\langle M(G,n), X\rangle$, the set of basepoint-preserving homotopy classes of maps $M(G,n)\to X$. To give this set a more suggestive name, let us call it $\mu_n(X;G)$. We should assume $n > 1$ to guarantee that the homotopy type of $M(G,n)$ is well-defined, as shown in Example 4.34. For $n > 1$, $\mu_n(X;G)$ is a group since we can choose $M(G,n)$ to be the suspension of an $M(G,n-1)$. And if $n > 2$ then $\mu_n(X;G)$ is abelian since we can choose $M(G,n)$ to be a double suspension.

There is something like a universal coefficient theorem for these groups $\mu_n(X;G)$:

Proposition 4H.2. *For $n > 1$ there are natural short exact sequences*

$$0 \longrightarrow \mathrm{Ext}(G,\pi_{n+1}(X)) \longrightarrow \mu_n(X;G) \longrightarrow \mathrm{Hom}(G,\pi_n(X)) \longrightarrow 0.$$

The similarity with the universal coefficient theorem for cohomology is apparent, but with a reversal of the variables in the Ext and Hom terms, reflecting the fact that $\mu_n(X;G)$ is covariant as a functor of X and contravariant as a functor of G, just like the Ext and Hom terms.

Proof: Let $0\to R \xrightarrow{\;i\;} F \longrightarrow G\to 0$ be a free resolution of G. The inclusion map i is realized by a map $M(R,n)\to M(F,n)$, where $M(R,n)$ and $M(F,n)$ are wedges of S^n's corresponding to bases for F and R. Converting this map into a cofibration via the mapping cylinder, the cofiber is an $M(G,n)$, as one sees from the long exact sequence of homology groups. As in §4.3, the cofibration sequence

$$M(R,n)\to M(F,n)\to M(G,n)\to M(R,n+1)\to M(F,n+1)$$

gives rise to the exact sequence across the top of the following diagram:

$$\begin{array}{ccccccccc}
\mu_{n+1}(X;F) & \longrightarrow & \mu_{n+1}(X;R) & \longrightarrow & \mu_n(X;G) & \longrightarrow & \mu_n(X;F) & \longrightarrow & \mu_n(X;R) \\
\| & & \| & & & & \| & & \| \\
\mathrm{Hom}(F,\pi_{n+1}(X)) & \xrightarrow{\;i^*\;} & \mathrm{Hom}(R,\pi_{n+1}(X)) & & & & \mathrm{Hom}(F,\pi_n(X)) & \xrightarrow{\;i^*\;} & \mathrm{Hom}(R,\pi_n(X))
\end{array}$$

The four outer terms of the exact sequence can be identified with the indicated Hom terms since mapping a wedge sum of S^n's into X amounts to choosing an element of $\pi_n(X)$ for each wedge summand. The kernel and cokernel of i^* are $\mathrm{Hom}(G,-)$ and $\mathrm{Ext}(G,-)$ by definition, and so we obtain the short exact sequence we are looking for. Naturality will be left for the reader to verify. $\square$

Unlike in the universal coefficient theorems for homology and cohomology, the short exact sequence in this proposition does not split in general. For an example, take $G = \mathbb{Z}_2$ and $X = M(\mathbb{Z}_2,n)$ for $n \ge 2$, where the identity map of $M(\mathbb{Z}_2,n)$

defines an element of $\mu_n(M(\mathbb{Z}_2,n);\mathbb{Z}_2) = \langle M(\mathbb{Z}_2,n), M(\mathbb{Z}_2,n)\rangle$ having order 4, as we show in Example 4L.7, whereas the two outer terms in the short exact sequence can only contain elements of order 2 since $G = \mathbb{Z}_2$. This example shows also that $\mu_n(X;\mathbb{Z}_m)$ need not be a module over $\mathbb{Z}_m$, as homology and cohomology groups with $\mathbb{Z}_m$ coefficients are.

The proposition implies that the first nonzero $\mu_i(S^n;\mathbb{Z}_m)$ is $\mu_{n-1}(S^n;\mathbb{Z}_m) = \mathbb{Z}_m$, from the Ext term. This result would look more reasonable if we changed notation to replace the subscript $n-1$ by n. So let us make the definition

$$\pi_n(X;\mathbb{Z}_m) = \langle M(\mathbb{Z}_m,n-1),X\rangle = \mu_{n-1}(X;\mathbb{Z}_m)$$

There are two good reasons to expect this to be the right definition. The first is formal: $M(\mathbb{Z}_m,n-1)$ is a 'cohomology $M(\mathbb{Z}_m,n)$' since its only nontrivial cohomology group $\widetilde{H}^i(M(\mathbb{Z}_m,n-1);\mathbb{Z})$ is $\mathbb{Z}_m$ in dimension n. The second reason is more geometric: Elements of $\pi_n(X;\mathbb{Z}_m)$ should be homotopy classes of 'homotopy-theoretic cycles mod m', meaning maps $D^n \to X$ whose boundary is not necessarily a constant map as would be the case for $\pi_n(X)$, but rather whose boundary is m times a cycle $S^{n-1} \to X$. This is precisely what a map $M(\mathbb{Z}_m,n-1) \to X$ is, if we choose $M(\mathbb{Z}_m,n-1)$ to be S^{n-1} with a cell e^n attached by a map of degree m.

Besides the calculation $\pi_n(S^n;\mathbb{Z}_m) \approx \mathbb{Z}_m$, the proposition also yields an isomorphism $\pi_n(M(\mathbb{Z}_m,n);\mathbb{Z}_m) \approx \operatorname{Ext}(\mathbb{Z}_m,\mathbb{Z}_m) = \mathbb{Z}_m$. Both these results are in fact special cases of a Hurewicz-type theorem relating $\pi_n(X;\mathbb{Z}_m)$ and $H_n(X;\mathbb{Z}_m)$, which is proved in [Neisendorfer 1980].

Along with $\mathbb{Z}$ and $\mathbb{Z}_m$, another extremely useful coefficient group for homology and cohomology is $\mathbb{Q}$. We pointed out above the difficulty that there is no cohomology analog of $M(\mathbb{Q},n)$. The groups $\mu_n(X;\mathbb{Q})$ are also problematic. For example the proposition gives $\mu_{n-1}(S^n;\mathbb{Q}) \approx \operatorname{Ext}(\mathbb{Q},\mathbb{Z})$, which is a somewhat complicated uncountable group as we showed in §3.F. However, there is an alternative approach that turns out to work rather well. One defines rational homotopy groups simply as $\pi_n(X)\otimes\mathbb{Q}$, analogous to the isomorphism $H_n(X;\mathbb{Q}) \approx H_n(X;\mathbb{Z})\otimes\mathbb{Q}$ from the universal coefficient theorem for homology. See [SSAT] for more on this.

Homology Decompositions

Eckmann–Hilton duality can be extremely helpful as an organizational principle, reducing significantly what one has to remember, and providing valuable hints on how to proceed in various situations. To illustrate, let us consider what would happen if we dualized the notion of a Postnikov tower of principal fibrations, where a space is represented as an inverse limit of a sequence of fibers of maps to Eilenberg–MacLane spaces. In the dual representation, a space would be realized as a direct limit of a sequence of cofibers of maps from Moore spaces.

In more detail, suppose we are given a sequence of abelian groups G_n, $n \geq 1$, and we build a CW complex X with $H_n(X) \approx G_n$ for all n by constructing inductively an increasing sequence of subcomplexes $X_1 \subset X_2 \subset \cdots$ with $H_i(X_n) \approx G_i$ for $i \le n$ and $H_i(X_n) = 0$ for $i > n$, where:

(1) X_1 is a Moore space $M(G_1, 1)$.
(2) X_{n+1} is the mapping cone of a cellular map $h_n : M(G_{n+1}, n) \to X_n$ such that the induced map $h_{n*} : H_n(M(G_{n+1}, n)) \to H_n(X_n)$ is trivial.
(3) $X = \bigcup_n X_n$.

One sees inductively that X_{n+1} has the desired homology groups by comparing the long exact sequences of the pairs (X_{n+1}, X_n) and (CM, M) where $M = M(G_{n+1}, n)$ and CM is the cone $M \times I / M \times \{0\}$:

$$\begin{array}{ccccccccc} 0 \longrightarrow H_{n+1}(X_{n+1}) \longrightarrow & H_{n+1}(X_{n+1}, X_n) & \xrightarrow{\partial} & H_n(X_n) & \longrightarrow H_n(X_{n+1}) \longrightarrow 0 \\ & \uparrow \approx & & \uparrow h_{n*} & \\ & H_{n+1}(CM, M) & \xrightarrow[\approx]{\partial} & H_n(M) \approx G_{n+1} & \end{array}$$

The assumption that h_{n*} is trivial means that the boundary map in the upper row is zero, hence $H_{n+1}(X_{n+1}) \approx G_{n+1}$. The other homology groups of X_{n+1} are the same as those of X_n since $H_i(X_{n+1}, X_n) \approx H_i(CM, M)$ for all i by excision, and $H_i(CM, M) \approx \widetilde{H}_{i-1}(M)$ since CM is contractible.

In case all the maps h_n are trivial, X is the wedge sum of the Moore spaces $M(G_n, n)$ since in this case the mapping cone construction in (2) produces a wedge sum with the suspension of $M(G_{n+1}, n)$, a Moore space $M(G_{n+1}, n+1)$.

For a space Y, a homotopy equivalence $f : X \to Y$ where X is constructed as in (1)–(3) is called a **homology decomposition** of Y.

Theorem 4H.3. *Every simply-connected CW complex has a homology decomposition.*

Proof: Given a simply-connected CW complex Y, let $G_n = H_n(Y)$. Suppose inductively that we have constructed X_n via maps h_i as in (2), together with a map $f : X_n \to Y$ inducing an isomorphism on H_i for $i \le n$. The induction can start with X_1 a point since Y is simply-connected. To construct X_{n+1} we first replace Y by the mapping cylinder of $f : X_n \to Y$, converting f into an inclusion. By the Hurewicz theorem and the homology exact sequence of the pair (Y, X_n) we have $\pi_{n+1}(Y, X_n) \approx H_{n+1}(Y, X_n) \approx H_{n+1}(Y) = G_{n+1}$. We will use this isomorphism to construct a map $h_n : M(G_{n+1}, n) \to X_n$ and an extension $f : X_{n+1} \to Y$.

The standard construction of an $M(G_{n+1}, n)$ consists of a wedge of spheres S^n_α corresponding to generators g_α of G_{n+1}, with cells e^{n+1}_β attached according to certain linear combinations $r_\beta = \sum_\alpha n_{\alpha\beta} g_\alpha$ that are zero in G_{n+1}. Under the isomorphism $G_{n+1} \approx \pi_{n+1}(Y, X_n)$ each g_α corresponds to a basepoint-preserving map $f_\alpha : (CS^n, S^n) \to (Y, X_n)$ where CS^n is the cone on S^n. The restrictions of these f_α's to S^n define $h_n : \bigvee_\alpha S^n_\alpha \to X_n$, and the maps $f_\alpha : CS^n \to Y$ themselves give an extension of $f : X_n \to Y$ to the mapping cone of $h_n : \bigvee_\alpha S^n_\alpha \to X_n$. Each relation r_β gives a homotopy $F_\beta : (CS^n, S^n) \times I \to (Y, X_n)$ from $\sum_\alpha n_{\alpha\beta} f_\alpha$ to the constant map. We use

$F_\beta | S^n \times \{0\}$ to attach e^{n+1}_β, and then $F_\beta | S^n \times I$ gives h_n on e^{n+1}_β and F_β gives an extension of f over the cone on e^{n+1}_β.

This construction assures that $f_* : H_{n+1}(X_{n+1}, X_n) \to H_{n+1}(Y, X_n)$ is an isomorphism, so from the five-lemma applied to the long exact sequences of these pairs we deduce that $f_* : H_i(X_{n+1}) \to H_i(Y)$ is an isomorphism for $i \le n+1$. This finishes the induction step. We may assume the maps f_α and F_β are cellular, so $X = \bigcup_n X_n$ is a CW complex with subcomplexes X_n. Since $f : X \to Y$ is a homology isomorphism between simply-connected CW complexes, it is a homotopy equivalence. ◻

As an example, suppose that Y is a simply-connected CW complex having all its homology groups free. Then the Moore spaces used in the construction of X can be taken to be wedges of spheres, and so X_n is obtained from X_{n-1} by attaching an n-cell for each $\mathbb{Z}$ summand of $H_n(Y)$. The attaching maps may be taken to be cellular, making X into a CW complex whose cellular chain complex has trivial boundary maps. Similarly, finite cyclic summands of $H_n(Y)$ can be realized by wedge summands of the form $S^{n-1} \cup e^n$ in $M(H_n(Y), n-1)$, contributing an n-cell and an $(n+1)$-cell to X. This is Proposition 4C.1, but the present result is stronger because it tells us that a finite cyclic summand of H_n can be realized in one step by attaching the cone on a Moore space $M(\mathbb{Z}_k, n-1)$, rather than in two steps of attaching an n-cell and then an $(n+1)$-cell.

Exercises

1. Show that if $A \hookrightarrow X$ is a cofibration of compact Hausdorff spaces, then for any space Y, the map $Y^X \to Y^A$ obtained by restriction of functions is a fibration. [If $A \hookrightarrow X$ is a cofibration, so is $A \times Y \hookrightarrow X \times Y$ for any space Y.]

2. Consider a pushout diagram as at the right, where $B \sqcup_f X$ is B with X attached along A via f. Show that if $A \hookrightarrow X$ is a cofibration, so is $B \hookrightarrow B \sqcup_f X$.

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \downarrow & & \downarrow \\ X & \longrightarrow & B \sqcup_f X \end{array}$$

3. For fibrations $E_1 \to B$ and $E_2 \to B$, show that a fiber-preserving map $E_1 \to E_2$ that is a homotopy equivalence is in fact a fiber homotopy equivalence. [This is dual to Proposition 0.19.]

4. Define the dual of an iterated mapping cylinder precisely, in terms of maps from Δ^n, and use this to give a definition of ∇X, the dual of ΔX, for X a complex of spaces.

4.I Stable Splittings of Spaces

It sometimes happens that suspending a space has the effect of simplifying its homotopy type, as the suspension becomes homotopy equivalent to a wedge sum of smaller spaces. Much of the interest in such stable splittings comes from the fact that they provide a geometric explanation for algebraic splittings of homology and cohomology groups, as well as other algebraic invariants of spaces that are unaffected by suspension such as the cohomology operations studied in §4.L.

The simplest example of a stable splitting occurs for the torus $S^1 \times S^1$. Here the reduced suspension $\Sigma(S^1 \times S^1)$ is homotopy equivalent to $S^2 \vee S^2 \vee S^3$ since $\Sigma(S^1 \times S^1)$ is $S^2 \vee S^2$ with a 3-cell attached by the suspension of the attaching map of the 2-cell of the torus, but the latter attaching map is the commutator of the two inclusions $S^1 \hookrightarrow S^1 \vee S^1$, and the suspension of this commutator is trivial since it lies in the abelian group $\pi_2(S^2 \vee S^2)$.

By an easy geometric argument we will prove more generally:

Proposition 4I.1. *If X and Y are CW complexes, then $\Sigma(X \times Y) \simeq \Sigma X \vee \Sigma Y \vee \Sigma(X \wedge Y)$.*

For example, $\Sigma(S^m \times S^n) \simeq S^{m+1} \vee S^{n+1} \vee S^{m+n+1}$. In view of the cup product structure on $H^*(S^m \times S^n)$ there can be no such splitting of $S^m \times S^n$ before suspension.

Proof: Consider the join $X * Y$ defined in Chapter 0, consisting of all line segments joining points in X to points in Y. For our present purposes it is convenient to use the reduced version of the join, obtained by collapsing to a point the line segment joining the basepoints $x_0 \in X$ and $y_0 \in Y$. We will still denote this reduced join by $X * Y$. Consider the space obtained from $X * Y$ by attaching reduced cones CX and CY to the copies of X and Y at the two ends of $X * Y$. If we collapse each of these cones to a point, we get the reduced suspension $\Sigma(X \times Y)$.

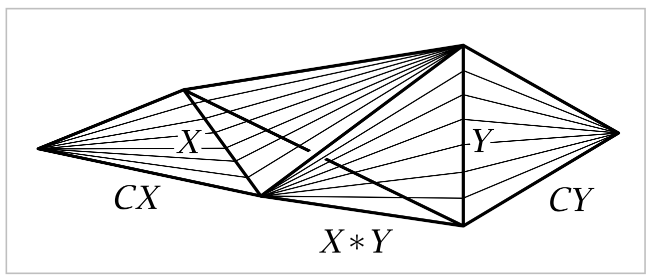

Since each cone is contractible, collapsing the cones gives a homotopy equivalence $X * Y \cup CX \cup CY \simeq \Sigma(X \times Y)$. Inside $X * Y$ there are also cones $x_0 * Y$ and $X * y_0$ intersecting in a point. Collapsing these cones converts $X * Y$ into $\Sigma(X \wedge Y)$ and $X * Y \cup CX \cup CY$ into $\Sigma(X \wedge Y) \vee \Sigma X \vee \Sigma Y$. ◻

This result can be applied inductively to obtain splittings for suspensions of products of more than two spaces, using the fact that reduced suspension is smash product with S^1, and smash product is associative and commutative. For example,

$$\Sigma(X \times Y \times Z) \simeq \Sigma X \vee \Sigma Y \vee \Sigma Z \vee \Sigma(X \wedge Y) \vee \Sigma(X \wedge Z) \vee \Sigma(Y \wedge Z) \vee \Sigma(X \wedge Y \wedge Z)$$

Our next example involves the reduced product $J(X)$ defined in §3.2. An interesting case is $J(S^n)$, which has a CW structure of the form $S^n \cup e^{2n} \cup e^{3n} \cup \cdots$. All the cells e^{in} for $i > 1$ are attached nontrivially since $H^*(J(S^n);\mathbb{Q})$ is a polynomial ring $\mathbb{Q}[x]$ for n even and a tensor product $\mathbb{Q}[x] \otimes \Lambda_{\mathbb{Q}}[y]$ for n odd. However, after we suspend to $\Sigma J(S^n)$, it is a rather surprising fact that all the attaching maps become trivial:

Proposition 4I.2. $\Sigma J(S^n) \simeq S^{n+1} \vee S^{2n+1} \vee S^{3n+1} \vee \cdots$. *More generally, if X is a connected CW complex then $\Sigma J(X) \simeq \bigvee_n \Sigma X^{\wedge n}$ where $X^{\wedge n}$ denotes the smash product of n copies of X.*

Proof: The space $J(X)$ is the union of an increasing sequence of subcomplexes $J_k(X)$ with $J_k(X)$ a quotient of the k-fold product $X^{\times k}$. The quotient $J_k(X)/J_{k-1}(X)$ is $X^{\wedge k}$. Thus we have maps

$$X^{\times k} \longrightarrow J_k(X) \longrightarrow X^{\wedge k} = J_k(X)/J_{k-1}(X)$$

By repeated application of the preceding proposition, $\Sigma X^{\wedge k}$ is a wedge summand of $\Sigma X^{\times k}$, up to homotopy equivalence. The proof shows moreover that there is a map $\Sigma X^{\wedge k} \to \Sigma X^{\times k}$ such that the composition $\Sigma X^{\wedge k} \to \Sigma X^{\times k} \to \Sigma X^{\wedge k}$ is homotopic to the identity. This composition factors as

$$\Sigma X^{\wedge k} \longrightarrow \Sigma X^{\times k} \longrightarrow \Sigma J_k(X) \longrightarrow \Sigma X^{\wedge k}$$

so we obtain a map $s_k : \Sigma X^{\wedge k} \to \Sigma J_k(X)$ such that $\Sigma X^{\wedge k} \xrightarrow{s_k} \Sigma J_k(X) \to \Sigma X^{\wedge k}$ is homotopic to the identity.

The map s_k induces a splitting of the long exact sequence of homology groups for the pair $(\Sigma J_k(X), \Sigma J_{k-1}(X))$. Hence the map $i \vee s_k : \Sigma J_{k-1}(X) \vee \Sigma X^{\wedge k} \to \Sigma J_k(X)$ induces an isomorphism on homology, where i denotes the inclusion map. It follows by induction that the map $\bigvee_{k=1}^n s_k : \bigvee_{k=1}^n \Sigma X^{\wedge k} \to \Sigma J_n(X)$ induces an isomorphism on homology for all finite n. This implies the corresponding statement for $n = \infty$ since $X^{\wedge n}$ is $(n-1)$-connected if X is connected. Thus we have a map $\bigvee_k \Sigma X^{\wedge k} \to \Sigma J(X)$ inducing an isomorphism on homology. By Whitehead's theorem this map is a homotopy equivalence since the spaces are simply-connected CW complexes. □

For our final example the stable splitting will be constructed using the group structure on $\langle \Sigma X, Y \rangle$, the set of basepointed homotopy classes of maps $\Sigma X \to Y$.

Proposition 4I.3. *For any prime power p^n the suspension $\Sigma K(\mathbb{Z}_{p^n}, 1)$ is homotopy equivalent to a wedge sum $X_1 \vee \cdots \vee X_{p-1}$ where X_i is a CW complex having $\widetilde{H}_*(X_i;\mathbb{Z})$ nonzero only in dimensions congruent to $2i \bmod 2p-2$.*

This result is best possible in a strong sense: No matter how many times any one of the spaces X_i is suspended, it never becomes homotopy equivalent to a nontrivial wedge sum. This will be shown in Example 4L.3 by studying cohomology operations in $H^*(K(\mathbb{Z}_{p^n}, 1); \mathbb{Z}_p)$. There is also a somewhat simpler K–theoretic explanation for this; see [VBKT].

Proof: Let $K = K(\mathbb{Z}_{p^n}, 1)$. The multiplicative group of nonzero elements in the field $\mathbb{Z}_p$ is cyclic, so let the integer r represent a generator. By Proposition 1B.9 there is a map $f : K \to K$ inducing multiplication by r on $\pi_1(K)$. We will need to know that f induces multiplication by r^i on $H_{2i-1}(K;\mathbb{Z}) \approx \mathbb{Z}_{p^n}$, and this can be seen as follows. Via the natural isomorphism $\pi_1(K) \approx H_1(K;\mathbb{Z})$ we know that f induces multiplication by r on $H_1(K;\mathbb{Z})$. Via the universal coefficient theorem, f also induces multiplication by r on $H^1(K;\mathbb{Z}_{p^n})$ and $H^2(K;\mathbb{Z}_{p^n})$. The cup product structure in $H^*(K;\mathbb{Z}_{p^n})$ computed in Examples 3.41 and 3E.2 then implies that f induces multiplication by r^i on $H^{2i-1}(K;\mathbb{Z}_{p^n})$, so the same is true for $H_{2i-1}(K;\mathbb{Z})$ by another application of the universal coefficient theorem.

For each integer $j \geq 0$ let $h_j : \Sigma K \to \Sigma K$ be the difference $\Sigma f - r^j \mathbb{1}$, so h_j induces multiplication by $r^i - r^j$ on $H_{2i}(\Sigma K;\mathbb{Z}) \approx \mathbb{Z}_{p^n}$. By the choice of r we know that p divides $r^i - r^j$ iff $i \equiv j \bmod p-1$. This means that the map induced by h_j on $\widetilde{H}_{2i}(\Sigma K;\mathbb{Z})$ has nontrivial kernel iff $i \equiv j \bmod p-1$. Therefore the composition $m_i = h_1 \circ \cdots h_{i-1} \circ h_{i+1} \cdots h_{p-1}$ induces an isomorphism on $\widetilde{H}_*(\Sigma K;\mathbb{Z})$ in dimensions congruent to $2i \bmod 2p-2$ and has a nontrivial kernel in other dimensions where the homology group is nonzero. When there is a nontrivial kernel, some power of m_i will induce the zero map since we are dealing with homomorphisms $\mathbb{Z}_{p^n} \to \mathbb{Z}_{p^n}$.

Let X_i be the mapping telescope of the sequence $\Sigma K \to \Sigma K \to \cdots$ where each map is m_i. Since homology commutes with direct limits, the inclusion of the first factor $\Sigma K \hookrightarrow X_i$ induces an isomorphism on $\widetilde{H}_*$ in dimensions congruent to $2i \bmod 2p-2$, and $\widetilde{H}_*(X_i;\mathbb{Z}) = 0$ in all other dimensions. The sum of these inclusions is a map $\Sigma K \to X_1 \vee \cdots \vee X_{p-1}$ inducing an isomorphism on all homology groups. Since these complexes are simply-connected, the result follows by Whitehead's theorem. $\square$

The construction of the spaces X_i as mapping telescopes produces rather large spaces, with infinitely many cells in each dimension. However, by Proposition 4C.1 each X_i is homotopy equivalent to a CW complex with the minimum configuration of cells consistent with its homology, namely, a 0-cell and a k-cell for each k congruent to $2i$ or $2i+1 \bmod 2p-2$.

Stable splittings of $K(G,1)$'s for finite groups G have been much studied and are a complicated and subtle business. To take the simplest noncyclic example, Proposition 4I.1 implies that $\Sigma K(\mathbb{Z}_2 \times \mathbb{Z}_2, 1)$ splits as the wedge sum of two copies of $\Sigma K(\mathbb{Z}_2, 1)$ and $\Sigma\big(K(\mathbb{Z}_2,1) \wedge K(\mathbb{Z}_2,1)\big)$, but the latter summand can be split further, according to a result in [Harris & Kuhn 1988] which says that for G the direct sum of k copies of $\mathbb{Z}_{p^n}$, $\Sigma K(G,1)$ splits canonically as the wedge sum of pieces having exactly $p^k - 1$ distinct homotopy types. Some of these summands occur more than once, as we see in the case of $\mathbb{Z}_2 \times \mathbb{Z}_2$.

Exercises

1. If a connected CW complex X retracts onto a subcomplex A, show that $\Sigma X \simeq \Sigma A \vee \Sigma(X/A)$. [One approach: Show the map $\Sigma r + \Sigma q : \Sigma X \to \Sigma A \vee \Sigma(X/A)$ induces an isomorphism on homology, where $r : X \to A$ is the retraction and $q : X \to X/A$ is the quotient map.]

2. Using the Künneth formula, show that $\Sigma K(\mathbb{Z}_m \times \mathbb{Z}_n, 1) \simeq \Sigma K(\mathbb{Z}_m, 1) \vee \Sigma K(\mathbb{Z}_n, 1)$ if m and n are relatively prime. Thus to determine stable splittings of $K(\mathbb{Z}_n, 1)$ it suffices to do the case that n is a prime power, as in Proposition 4I.3.

3. Extending Proposition 4I.3, show that the $(2k+1)$-skeleton of the suspension of a high-dimensional lens space with fundamental group of order p^n is homotopy equivalent to the wedge sum of the $(2k+1)$-skeleta of the spaces X_i, if these X_i's are chosen to have the minimum number of cells in each dimension, as described in the remarks following the proof.

4.J The Loopspace of a Suspension

Loopspaces appear at first glance to be hopelessly complicated objects, but if one is only interested in homotopy type, there are many cases when great simplifications are possible. One of the nicest of these cases is the loopspace of a sphere. We show in this section that ΩS^{n+1} has the weak homotopy type of the James reduced product $J(S^n)$ introduced in §3.2. More generally, we show that $\Omega\Sigma X$ has the weak homotopy type of $J(X)$ for every connected CW complex X. If one wants, one can strengthen 'weak homotopy type' to 'homotopy type' by quoting Milnor's theorem, mentioned in §4.3, that the loopspace of a CW complex has the homotopy type of a CW complex.

Part of the interest in $\Omega\Sigma X$ can be attributed to its close connection with the suspension homomorphism $\pi_i(X) \to \pi_{i+1}(\Sigma X)$. We will use the weak homotopy equivalence of $\Omega\Sigma X$ with $J(X)$ to give another proof that the suspension homomorphism is an isomorphism in dimensions up to approximately double the connectivity of X. In addition, we will obtain an exact sequence that measures the failure of the suspension map to be an isomorphism in dimensions between double and triple the connectivity of X. An easy application of this, together with results proved elsewhere in the book, will be to compute $\pi_{n+1}(S^n)$ and $\pi_{n+2}(S^n)$ for all n.

As a rough first approximation to ΩS^{n+1} there is a natural inclusion of S^n into ΩS^{n+1} obtained by regarding S^{n+1} as the reduced suspension ΣS^n, the quotient $(S^n \times I)/(S^n \times \partial I \cup \{e\} \times I)$ where e is the basepoint of S^n, then associating to each point $x \in S^n$ the loop $\lambda(x)$ in ΣS^n given by $t \mapsto (x, t)$. The figure shows what a few such loops look like. However, we cannot expect this inclusion $S^n \hookrightarrow \Omega S^{n+1}$ to be a homotopy equivalence since ΩS^{n+1} is an H-space but S^n is only an H-space when $n = 1, 3, 7$ by the theorem of Adams discussed in §4.B. The simplest way to correct this deficiency in S^n would be to replace it by the free H-space that it generates, the reduced product $J(S^n)$. Re-

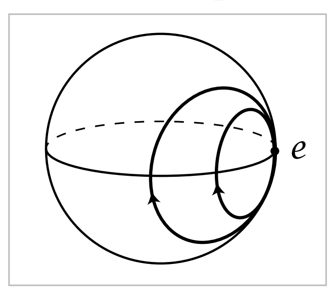

call from §3.2 that a point in $J(S^n)$ is a formal product $x_1 \cdots x_k$ of points $x_i \in S^n$, with the basepoint e acting as an identity element for the multiplication obtained by juxtaposition of formal products. We would like to define a map $\lambda : J(S^n) \to \Omega S^{n+1}$ by setting $\lambda(x_1 \cdots x_k) = \lambda(x_1) \cdots \lambda(x_k)$, the product of the loops $\lambda(x_i)$. The only difficulty is in the parametrization of this product, which needs to be adjusted so that λ is continuous. The problem is that when some x_i approaches the basepoint $e \in S^n$, one wants the loop $\lambda(x_i)$ to disappear gradually from the product $\lambda(x_1) \cdots \lambda(x_k)$, without disrupting the parametrization as simply deleting $\lambda(e)$ would do. This can be achieved by first making the time it takes to traverse each loop $\lambda(x_i)$ equal to the distance from x_i to the basepoint of S^n, then normalizing the resulting product of loops so that it takes unit time, giving a map $I \to \Sigma S^n$.

More generally, this same procedure defines a map $\lambda : J(X) \to \Omega\Sigma X$ for any connected CW complex X, where 'distance to the basepoint' is replaced by any map $d : X \to [0,1]$ with $d^{-1}(0) = e$, the basepoint of X.

Theorem 4J.1. *The map $\lambda : J(X) \to \Omega\Sigma X$ is a weak homotopy equivalence for every connected CW complex X.*

Proof: The main step will be to compute the homology of $\Omega\Sigma X$. After this is done, it will be easy to deduce that λ induces an isomorphism on homology using the calculation of the homology of $J(X)$ in Proposition 3C.8, and from this conclude that λ is a weak homotopy equivalence. It will turn out to be sufficient to consider homology with coefficients in a field F. We know that $H_*(J(X);F)$ is the tensor algebra $T\widetilde{H}_*(X;F)$ by Proposition 3C.8, so we want to show that $H_*(\Omega\Sigma X;F)$ has this same structure, a result first proved in [Bott & Samelson 1953].

Let us write the reduced suspension $Y = \Sigma X$ as the union of two reduced cones $Y_+ = C_+X$ and $Y_- = C_-X$ intersecting in the equatorial $X \subset \Sigma X$. Consider the path fibration $p : PY \to Y$ with fiber ΩY. Let $P_+Y = p^{-1}(Y_+)$ and $P_-Y = p^{-1}(Y_-)$, so P_+Y consists of paths in Y starting at the basepoint and ending in Y_+, and similarly for P_-Y. Then $P_+Y \cap P_-Y$ is $p^{-1}(X)$, the paths from the basepoint to X. Since Y_+ and Y_- are deformation retracts of open neighborhoods U_+ and U_- in Y such that $U_+ \cap U_-$ deformation retracts onto $Y_+ \cap Y_- = X$, the homotopy lifting property implies that P_+Y, P_-Y, and $p_+Y \cap P_-Y$ are deformation retracts, in the weak sense, of open neighborhoods $p^{-1}(U_+)$, $p^{-1}(U_-)$, and $p^{-1}(U_+) \cap p^{-1}(U_-)$, respectively. Therefore there is a Mayer–Vietoris sequence in homology for the decomposition of PY as $P_+Y \cup P_-Y$. Taking reduced homology and using the fact that PY is contractible, this gives an isomorphism

(i) $$\Phi : \widetilde{H}_*(P_+Y \cap P_-Y;F) \xrightarrow{\approx} \widetilde{H}_*(P_+Y;F) \oplus \widetilde{H}_*(P_-Y;F)$$

The two coordinates of Φ are induced by the inclusions, with a minus sign in one case, but Φ will still be an isomorphism if this minus sign is eliminated, so we may assume this has been done.

We claim that the isomorphism Φ can be rewritten as an isomorphism

$$\text{(ii)} \qquad \Theta : \widetilde{H}_*(\Omega Y \times X;F) \xrightarrow{\approx} \widetilde{H}_*(\Omega Y;F) \oplus \widetilde{H}_*(\Omega Y;F)$$

To see this, we observe that the fibration $P_+Y \to Y_+$ is fiber-homotopically trivial. This is true since the cone Y_+ is contractible, but we shall need an explicit fiber homotopy equivalence $P_+Y \simeq \Omega Y \times Y_+$, and this is easily constructed as follows. Define $f_+ : P_+Y \to \Omega Y \times Y_+$ by $f_+(\gamma) = (\gamma \cdot \gamma_y^+, y)$ where $y = \gamma(1)$ and γ_y^+ is the obvious path in Y_+ from $y = (x,t)$ to the basepoint along the segment $\{x\} \times I$. In the other direction, define $g_+ : \Omega Y \times Y_+ \to P_+Y$ by $g_+(\gamma, y) = \gamma \cdot \overline{\gamma}_y^+$ where the bar denotes the inverse path. Then f_+g_+ and g_+f_+ are fiber-homotopic to the respective identity maps since $\overline{\gamma}_y^+ \cdot \gamma_y^+$ and $\gamma_y^+ \cdot \overline{\gamma}_y^+$ are homotopic to the constant paths.

In similar fashion the fibration $P_-Y \to Y_-$ is fiber-homotopically trivial via maps f_- and g_-. By restricting a fiber-homotopy trivialization of either P_+Y or P_-Y to $P_+Y \cap P_-Y$, we see that the fibration $P_+Y \cap P_-Y$ is fiber-homotopy equivalent to the product $\Omega Y \times X$. Let us do this using the fiber-homotopy trivialization of P_-Y. The groups in (i) can now be replaced by those in (ii). The map Φ has coordinates induced by inclusion, and it follows that the corresponding map Θ in (ii) has coordinates induced by the two maps $\Omega Y \times X \to \Omega Y$, $(\gamma, x) \mapsto \gamma \cdot \lambda(x)$ and $(\gamma, x) \mapsto \gamma$. Namely, the first coordinate of Θ is induced by $f_+g_- \,|\, \Omega Y \times X$ followed by projection to ΩY, and the second coordinate is the same but with f_-g_- in place of f_+g_-.

Writing the two coordinates of Θ as Θ_1 and Θ_2, the fact that Θ is an isomorphism means that the restriction of Θ_1 to the kernel of Θ_2 is an isomorphism. Via the Künneth formula we can write $\widetilde{H}_*(\Omega Y \times X;F)$ as $(H_*(\Omega Y;F) \otimes \widetilde{H}_*(X;F)) \oplus \widetilde{H}_*(\Omega Y;F)$ where projection onto the latter summand is Θ_2. Hence Θ_1 gives an isomorphism from the first summand $H_*(\Omega Y;F) \otimes \widetilde{H}_*(X;F)$ onto $\widetilde{H}_*(\Omega Y;F)$. Since $\Theta_1(\gamma, x) = (\gamma \cdot \lambda(x))$, this means that the composed map

$$H_*(\Omega Y;F) \otimes \widetilde{H}_*(X;F) \xrightarrow{1\!\!1 \otimes \lambda_*} H_*(\Omega Y;F) \otimes \widetilde{H}_*(\Omega Y;F) \longrightarrow \widetilde{H}_*(\Omega Y;F)$$

with the second map Pontryagin product, is an isomorphism. Now to finish the calculation of $H_*(\Omega Y;F)$ as the tensor algebra $T\widetilde{H}_*(X;F)$, we apply the following algebraic lemma, with $A = H_*(\Omega Y;F)$, $V = \widetilde{H}_*(X;F)$, and $i = \lambda_*$.

Lemma 4J.2. *Let A be a graded algebra over a field F with $A_0 = F$ and let V be a graded vector space over F with $V_0 = 0$. Suppose we have a linear map $i : V \to A$ preserving grading, such that the multiplication map $\mu : A \otimes V \to \widetilde{A}$, $\mu(a \otimes v) = ai(v)$, is an isomorphism. Then the canonical algebra homomorphism $i : TV \to A$ extending the previous i is an isomorphism.*

For example, if V is a 1-dimensional vector space over F, as happens in the case $X = S^n$, then this says that if the map $A \to \widetilde{A}$ given by right-multiplication by an element $a = i(v)$ is an isomorphism, then A is the polynomial algebra $F[a]$. The general case can be viewed as the natural generalization of this to polynomials in any number of noncommuting variables.

Proof: Since μ is an isomorphism, each element $a \in A_n$ with $n > 0$ can be written uniquely in the form $\mu(\sum_j a_j \otimes v_j) = \sum_j a_j i(v_j)$ for $v_j \in V$ and $a_j \in A_{n(j)}$, with $n(j) < n$ since $V_0 = 0$. By induction on n, $a_j = i(\alpha_j)$ for a unique $\alpha_j \in (TV)_{n(j)}$. Thus $a = i(\sum_j \alpha_j \otimes v_j)$ so i is surjective. Since these representations are unique, i is also injective. The induction starts with the hypothesis that $A_0 = F$, the scalars in TV. $\square$

Returning now to the proof of the theorem, we observe that λ is an H–map: The two maps $J(X)\times J(X)\to\Omega\Sigma X$, $(x,y)\mapsto\lambda(xy)$ and $(x,y)\mapsto\lambda(x)\lambda(y)$, are homotopic since the loops $\lambda(xy)$ and $\lambda(x)\lambda(y)$ differ only in their parametrizations. Since λ is an H–map, the maps $X\hookrightarrow J(X)\xrightarrow{\lambda}\Omega\Sigma X$ induce the commutative diagram at the right. We have shown that the downward map on the right is an isomorphism, and the same is true of the one on the left by the calculation of $H_*(J(X);F)$ in Proposition 3C.8. Hence λ_* is an isomorphism. By Corollary 3A.7 this is also true for $\mathbb{Z}$ coefficients. When X is simply-connected, so are $J(X)$ and $\Omega\Sigma X$, so after taking a CW approximation to $\Omega\Sigma X$, Whitehead's theorem implies that λ is a weak homotopy equivalence. In the general case that X is only connected, we obtain the same conclusion from the generalization of Whitehead's theorem to abelian spaces, Proposition 4.74, since $J(X)$ and $\Omega\Sigma X$ are H–spaces, with trivial action of π_1 on all homotopy groups by Example 4A.3. $\square$

$$\begin{array}{ccc} & T\widetilde{H}_*(X;F) & \\ \swarrow & & \searrow \\ H_*(J(X);F) & \xrightarrow{\lambda_*} & H_*(\Omega\Sigma X;F) \end{array}$$

Using the natural identification $\pi_i(\Omega\Sigma X) = \pi_{i+1}(\Sigma X)$, the inclusion $X\hookrightarrow\Omega\Sigma X$ induces the suspension map $\pi_i(X)\to\pi_{i+1}(\Sigma X)$. Since this inclusion factors through $J(X)$, we can identify the relative groups $\pi_i(\Omega\Sigma X,X)$ with $\pi_i(J(X),X)$. If X is n-connected then the pair $(J(X),X)$ is $(2n+1)$-connected since we can replace X by a complex with n-skeleton a point, and then the $(2n+1)$-skeleton of $J(X)$ is contained in X. Thus we have:

Corollary 4J.3. *The suspension map $\pi_i(X)\to\pi_{i+1}(\Sigma X)$ for an n-connected CW complex X is an isomorphism if $i \le 2n$ and a surjection if $i = 2n+1$.* $\square$

In the case of a sphere we can describe what happens in the first dimension when suspension is not an isomorphism, namely the suspension $\pi_{2n-1}(S^n)\to\pi_{2n}(S^{n+1})$ which the corollary guarantees only to be a surjection. The CW structure on $J(S^n)$ consists of a single cell in each dimension a multiple of n, so from exactness of $\pi_{2n}(J(S^n),S^n)\xrightarrow{\partial}\pi_{2n-1}(S^n)\xrightarrow{\Sigma}\pi_{2n}(S^{n+1})$ we see that the kernel of the suspension $\pi_{2n-1}(S^n)\to\pi_{2n}(S^{n+1})$ is generated by the attaching map of the $2n$-cell of $J(S^n)$. This attaching map is the Whitehead product $[\mathbb{1},\mathbb{1}]$, as we noted in §4.2 when we defined Whitehead products following Example 4.52. When n is even, the Hopf invariant homomorphism $\pi_{2n-1}(S^n)\to\mathbb{Z}$ has the value ± 2 on $[\mathbb{1},\mathbb{1}]$, as we saw in §4.B. If there is no map of Hopf invariant ± 1, it follows that $[\mathbb{1},\mathbb{1}]$ generates a $\mathbb{Z}$ summand of $\pi_{2n-1}(S^n)$, and so the suspension homomorphism simply cancels this summand from $\pi_{2n-1}(S^n)$. By Adams' theorem, this is the situation for all even n except 2, 4, and 8.

When $n=2$ we have $\pi_3(S^2)\approx\mathbb{Z}$ generated by the Hopf map η with Hopf invariant 1, so $2\eta=\pm[\mathbb{1},\mathbb{1}]$, generating the kernel of the suspension, hence:

Corollary 4J.4. *$\pi_{n+1}(S^n)$ is $\mathbb{Z}_2$ for $n\ge 3$, generated by the suspension or iterated suspension of the Hopf map.* □

The situation for $n=4$ and 8 is more subtle. We do not have the tools available here to do the actual calculation, but if we consult the table near the beginning of §4.1 we see that the suspension $\pi_7(S^4)\to\pi_8(S^5)$ is a map $\mathbb{Z}\oplus\mathbb{Z}_{12}\to\mathbb{Z}_{24}$. By our preceding remarks we know this map is surjective with kernel generated by the single element $[\mathbb{1},\mathbb{1}]$. Algebraically, what must be happening is that the coordinate of $[\mathbb{1},\mathbb{1}]$ in the $\mathbb{Z}$ summand is twice a generator, while the coordinate in the $\mathbb{Z}_{12}$ summand is a generator. Thus a generator of the $\mathbb{Z}$ summand, which we may take to be the Hopf map $S^7\to S^4$, suspends to a generator of the $\mathbb{Z}_{24}$. For $n=8$ the situation is entirely similar, with the suspension $\pi_{15}(S^8)\to\pi_{16}(S^9)$ a homomorphism $\mathbb{Z}\oplus\mathbb{Z}_{120}\to\mathbb{Z}_{240}$.

We can also obtain some information about suspension somewhat beyond the edge of the stable dimension range. Since S^n is $(n-1)$-connected and $(J(S^n),S^n)$ is $(2n-1)$-connected, we have isomorphisms $\pi_i(J(S^n),S^n)\approx\pi_i(J(S^n)/S^n)$ for $i\le 3n-2$ by Proposition 4.28. The group $\pi_i(J(S^n)/S^n)$ is isomorphic to $\pi_i(S^{2n})$ in the same range $i\le 3n-2$ since $J(S^n)/S^n$ has S^{2n} as its $(3n-1)$-skeleton. Thus the terminal portion of the long exact sequence of the pair $(J(S^n),S^n)$ starting with the term $\pi_{3n-2}(S^n)$ can be written in the form

$$\pi_{3n-2}(S^n)\xrightarrow{\Sigma}\pi_{3n-1}(S^{n+1})\to\pi_{3n-2}(S^{2n})\to\pi_{3n-3}(S^n)\xrightarrow{\Sigma}\pi_{3n-2}(S^{n+1})\to\cdots$$

This is known as the **EHP sequence** since its three maps were originally called E, H, and P. (The German word for 'suspension' begins with E, the H refers to a generalization of the Hopf invariant, and the P denotes a connection with Whitehead products; see [Whitehead 1978] for more details.) Note that the terms $\pi_i(S^{2n})$ in the EHP sequence are stable homotopy groups since $i\le 3n-2$. Thus we have the curious situation that stable homotopy groups are measuring the lack of stability of the groups $\pi_i(S^n)$ in the range $2n-1\le i\le 3n-2$, the so-called *metastable* range.

Specializing to the first interesting case $n=2$, the sequence becomes

$$\begin{array}{ccccccccccc}\pi_4(S^2) & \xrightarrow{\Sigma} & \pi_5(S^3) & \longrightarrow & \pi_4(S^4) & \longrightarrow & \pi_3(S^2) & \xrightarrow{\Sigma} & \pi_4(S^3) & \longrightarrow & 0\\ \wr\wr & & & & \wr\wr & & \wr\wr & & \wr\wr & & \\ \mathbb{Z}_2 & & & & \mathbb{Z} & & \mathbb{Z} & & \mathbb{Z}_2 & & \end{array}$$

From the Hopf bundle $S^1 \to S^3 \to S^2$ we have $\pi_4(S^2) \approx \pi_4(S^3) \approx \mathbb{Z}_2$, with $\pi_4(S^2)$ generated by the composition $\eta(\Sigma\eta)$ where η is the Hopf map $S^3 \to S^2$. From exactness of the latter part of the sequence we deduce that the map $\pi_4(S^4) \to \pi_3(S^2)$ is injective, and hence that the suspension $\pi_4(S^2) \to \pi_5(S^3)$ is surjective, so $\pi_5(S^3)$ is either $\mathbb{Z}_2$ or 0. From the general suspension theorem, the suspension $\pi_5(S^3) \to \pi_6(S^4)$ is surjective as well, and the latter group is in the stable range. We show in Proposition 4L.11 that the stable group π_2^s is nonzero, and so we conclude that $\pi_{n+2}(S^n) \approx \mathbb{Z}_2$ for all $n \ge 2$, generated by the composition $(\Sigma^{n-2}\eta)(\Sigma^{n-1}\eta)$.

We will see in [SSAT] that the EHP sequence extends all the way to the left to form an infinite exact sequence when n is odd, and when n is even a weaker statement holds: The sequence extends after factoring out all odd torsion.

Replacing S^n by any $(n-1)$-connected CW complex X, our derivation of the finite EHP sequence generalizes immediately to give an exact sequence

$$\pi_{3n-2}(X) \xrightarrow{\Sigma} \pi_{3n-1}(\Sigma X) \to \pi_{3n-2}(X \wedge X) \to \pi_{3n-3}(X) \xrightarrow{\Sigma} \pi_{3n-2}(\Sigma X) \to \cdots$$

using the fact that $J_2(X)/X = X \wedge X$.

The generalization of the results of this section to $\Omega^n \Sigma^n X$ turns out to be of some importance in homotopy theory. In case we do not get to this topic in [SSAT], the reader can begin to learn about it by looking at [Carlsson & Milgram 1995].

Exercise

1. Show that $\Omega\Sigma X$ for a nonconnected CW complex X reduces to the connected case by showing that each path-component of $\Omega\Sigma X$ is homotopy equivalent to $\Omega\Sigma(\bigvee_\alpha X_\alpha)$ where the X_α's are the components of X.

4.K The Dold-Thom Theorem

In the preceding section we studied the free monoid $J(X)$ generated by a space X, and in this section we take up its commutative analog, the free abelian monoid generated by X. This is the infinite symmetric product $SP(X)$ introduced briefly in §3.C. The main result will be a theorem of [Dold & Thom 1958] asserting that $\pi_* SP(X) \approx \widetilde{H}_*(X;\mathbb{Z})$ for all connected CW complexes X. In particular this yields the surprising fact that $SP(S^n)$ is a $K(\mathbb{Z},n)$, and more generally that the functor SP takes Moore spaces $M(G,n)$ to Eilenberg–MacLane spaces $K(G,n)$. This leads to the general result that for all connected CW complexes X, $SP(X)$ has the homotopy type of a product of Eilenberg–MacLane spaces. In other words, the k-invariants of $SP(X)$ are all trivial.

The main step in the proof of the Dold–Thom theorem will be to show that the homotopy groups $\pi_*SP(X)$ define a homology theory. An easy computation of the coefficient groups $\pi_*SP(S^n)$ will then show that this must be ordinary homology with $\mathbb{Z}$ coefficients. A new idea needed for the proof of the main step is the notion of a *quasifibration*, generalizing fibrations and fiber bundles. In order to establish a few basic facts about quasifibrations we first make a small detour to prove an essentially elementary fact about relative homotopy groups.

A Mayer–Vietoris Property of Homotopy Groups

In this subsection we will be concerned largely with relative homotopy groups, and it will be impossible to avoid the awkward fact that there is no really good way to define the relative π_0. What we will do as a compromise is to take $\pi_0(X,A,x_0)$ to be the quotient set $\pi_0(X,x_0)/\pi_0(A,x_0)$. This at least allows the long exact sequence of homotopy groups for (X,A) to end with the terms

$$\pi_0(A,x_0) \to \pi_0(X,x_0) \to \pi_0(X,A,x_0) \to 0$$

An exercise for §4.1 shows that the five-lemma can be applied to the map of long exact sequences induced by a map $(X,A) \to (Y,B)$, provided the basepoint is allowed to vary. However, the long exact sequence of a triple cannot be extended through the π_0 terms with this definition, so one must proceed with some caution.

The excision theorem for homology involves a space X with subspaces A and B such that X is the union of the interiors of A and B. In this situation we call $(X;A,B)$ an **excisive triad**. By a map $f:(X;A,B) \to (Y;C,D)$ we mean $f:X \to Y$ with $f(A) \subset C$ and $f(B) \subset D$.

Proposition 4K.1. *Let $f:(X;A,B) \to (Y;C,D)$ be a map of excisive triads. If the induced maps $\pi_i(A,A\cap B) \to \pi_i(C,C\cap D)$ and $\pi_i(B,A\cap B) \to \pi_i(D,C\cap D)$ are bijections for $i<n$ and surjections for $i=n$, for all choices of basepoints, then the same is true of the induced maps $\pi_i(X,A) \to \pi_i(Y,C)$. By symmetry the conclusion holds also for the maps $\pi_i(X,B) \to \pi_i(Y,D)$.*

The corresponding statement for homology is a trivial consequence of excision which says that $H_i(X,A) \approx H_i(B,A\cap B)$ and $H_i(Y,C) \approx H_i(D,C\cap D)$, so it is not necessary to assume anything about the map $H_i(A,A\cap B) \to H_i(C,C\cap D)$. With the failure of excision for homotopy groups, however, it is not surprising that the assumption on $\pi_i(A,A\cap B) \to \pi_i(C,C\cap D)$ cannot be dropped. An example is provided by the quotient map $f:D^2 \to S^2$ collapsing ∂D^2 to the north pole of S^2, with C and D the northern and southern hemispheres of S^2, and A and B their preimages under f.

Proof: First we will establish a general fact about relative homotopy groups. Consider an inclusion $(X,A) \hookrightarrow (Y,C)$. We will show the following three conditions are equivalent for each $n \ge 1$:

(i) For all choices of basepoints the map $\pi_i(X,A)\to\pi_i(Y,C)$ induced by inclusion is surjective for $i=n$ and has trivial kernel for $i=n-1$.

(ii) Let ∂D^n be written as the union of hemispheres ∂_+D^n and ∂_-D^n intersecting in S^{n-2}. Then every map

$$(D^n\times\{0\}\cup\partial_+D^n\times I,\partial_-D^n\times\{0\}\cup S^{n-2}\times I)\longrightarrow(Y,C)$$

taking $(\partial_+D^n\times\{1\},S^{n-2}\times\{1\})$ to (X,A) extends to a map $(D^n\times I,\partial_-D^n\times I)\to(Y,C)$ taking $(D^n\times\{1\},\partial_-D^n\times\{1\})$ to (X,A).

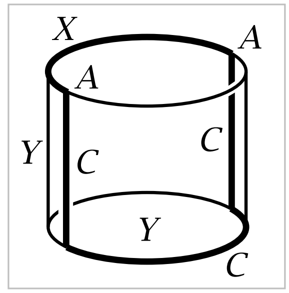

(iii) Condition (ii) with the added hypothesis that the restriction of the given map to $\partial_+D^n\times I$ is independent of the I coordinate.

It is obvious that (ii) and (iii) are equivalent since the stronger hypothesis in (iii) can always be achieved by composing with a homotopy of $D^n\times I$ that shrinks $\partial_+D^n\times I$ to $\partial_+D^n\times\{1\}$.

To see that (iii) implies (i), let $f:(\partial_+D^n\times\{1\},S^{n-2}\times\{1\})\to(X,A)$ represent an element of $\pi_{n-1}(X,A)$. If this is in the kernel of the map to $\pi_{n-1}(Y,C)$, then we get an extension of f over $D^n\times\{0\}\cup\partial_+D^n\times I$, with the constant homotopy on $\partial_+D^n\times I$ and $(D^n\times\{0\},\partial_-D^n\times\{0\})$ mapping to (Y,C). Condition (iii) then gives an extension over $D^n\times I$, whose restriction to $D^n\times\{1\}$ shows that f is zero in $\pi_{n-1}(X,A)$, so the kernel of $\pi_{n-1}(X,A)\to\pi_{n-1}(Y,C)$ is trivial. To check surjectivity of the map $\pi_n(X,A)\to\pi_n(Y,C)$, represent an element of $\pi_n(Y,C)$ by a map $f:D^n\times\{0\}\to Y$ taking $\partial_-D^n\times\{0\}$ to C and $\partial_+D^n\times\{0\}$ to a chosen basepoint. Extend f over $\partial_+D^n\times I$ via the constant homotopy, then extend over $D^n\times I$ by applying (iii). The result is a homotopy of the given f to a map representing an element of the image of $\pi_n(X,A)\to\pi_n(Y,C)$.

Now we show that (i) implies (ii). Given a map f as in the hypothesis of (ii), the injectivity part of (i) gives an extension of f over $D^n\times\{1\}$. Choose a small disk $E^n\subset\partial_-D^n\times I$, shown shaded in the figure, intersecting $\partial_-D^n\times\{1\}$ in a hemisphere ∂_+E^n of its boundary. We may assume the extended f has a constant value $x_0\in A$ on ∂_+E^n. Viewing the extended f as representing an element of $\pi_n(Y,C,x_0)$, the surjectivity part of (i) then gives an extension of f over $D^n\times I$ taking (E^n,∂_-E^n) to (X,A) and the rest of $\partial_-D^n\times I$ to C. The argument is finished by composing this extended f with a deformation of $D^n\times I$ pushing E^n into $D^n\times\{1\}$.

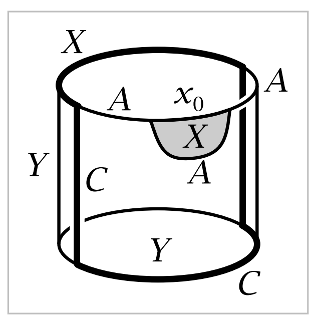

Having shown the equivalence of (i)–(iii), let us prove the proposition. We may reduce to the case that the given $f:(X;A,B)\to(Y;C,D)$ is an inclusion by using mapping cylinders. One's first guess would be to replace $(Y;C,D)$ by the triad of mapping cylinders $(M_f;M_{f|A},M_{f|B})$, where we view $f|A$ as a map $A\to C$ and $f|B$ as a map $B\to D$. However, the triad $(M_f;M_{f|A},M_{f|B})$ need not be excisive, for example if X

consists of two points A and B and Y is a single point. To remedy this problem, replace $M_{f|A}$ by its union with $f^{-1}(C)\times(1/2,1)$ in M_f, and enlarge $M_{f|B}$ similarly.

Now we prove the proposition for an inclusion $(X;A,B)\hookrightarrow(Y;C,D)$. The case $n=0$ is trivial from the definitions, so let us assume $n\ge 1$. In view of the equivalence of condition (i) with (ii) and (iii), it suffices to show that condition (ii) for the inclusions $(A,A\cap B)\hookrightarrow(C,C\cap D)$ and $(B,A\cap B)\hookrightarrow(D,C\cap D)$ implies (iii) for the inclusion $(X,A)\hookrightarrow(Y,C)$. Let a map $f:D^n\times\{0\}\cup\partial_+D^n\times I\to Y$ as in the hypothesis of (iii) be given. The argument will involve subdivision of D^n into smaller disks, and for this it is more convenient to use the cube I^n instead of D^n, so let us identify I^n with D^n in such a way that ∂_-D^n corresponds to the face $I^{n-1}\times\{1\}$, which we denote by ∂_-I^n, and ∂_+D^n corresponds to the remaining faces of I^n, which we denote by ∂_+I^n. Thus we are given f on $I^n\times\{0\}$ taking $\partial_+I^n\times\{0\}$ to X and $\partial_-I^n\times\{0\}$ to C, and on $\partial_+I^n\times I$ we have the constant homotopy.

Since $(Y;C,D)$ is an excisive triad, we can subdivide each of the I factors of $I^n\times\{0\}$ into subintervals so that f takes each of the resulting n-dimensional subcubes of $I^n\times\{0\}$ into either C or D. The extension of f we construct will have the following key property:

(∗) If K is a one of the subcubes of $I^n\times\{0\}$, or a lower-dimensional face of such a cube, then the extension of f takes $(K\times I,K\times\{1\})$ to (C,A) or (D,B) whenever f takes K to C or D, respectively.

Initially we have f defined on $\partial_+I^n\times I$ with image in X, independent of the I coordinate, and we may assume the condition $(*)$ holds here since we may assume that $A=X\cap C$ and $B=X\cap D$, these conditions holding for the mapping cylinder construction described above.

Consider the problem of extending f over $K\times I$ for K one of the subcubes. We may assume that f has already been extended to $\partial_+K\times I$ so that $(*)$ is satisfied, by induction on n and on the sequence of subintervals of the last coordinate of $I^n\times\{0\}$. To extend f over $K\times I$, let us first deal with the cases that the given f takes (K,∂_-K) to $(C,C\cap D)$ or $(D,C\cap D)$. Then by (ii) for the inclusion $(A,A\cap B)\hookrightarrow(C,C\cap D)$ or $(B,A\cap B)\hookrightarrow(D,C\cap D)$ we may extend f over $K\times I$ so that $(*)$ is still satisfied. If neither of these two cases applies, then the given f takes (K,∂_-K) just to (C,C) or (D,D), and we can apply (ii) trivially to construct the desired extension of f over $K\times I$. □

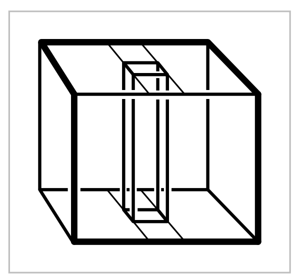

Corollary 4K.2. *Given a map $f:X\to Y$ and open covers $\{U_i\}$ of X and $\{V_i\}$ of Y with $f(U_i)\subset V_i$ for all i, then if each restriction $f:U_i\to V_i$ and more generally each $f:U_{i_1}\cap\cdots\cap U_{i_n}\to V_{i_1}\cap\cdots\cap V_{i_n}$ is a weak homotopy equivalence, so is $f:X\to Y$.*

Proof: First let us do the case of covers by two sets. By the five-lemma, the hypotheses imply that $\pi_n(U_i,U_1\cap U_2)\to\pi_n(V_i,V_1\cap V_2)$ is bijective for $i=1,2$, $n\ge 0$, and all choices of basepoints. The preceding proposition then implies that the maps $\pi_n(X, U_1) \to \pi_n(Y, V_1)$ are isomorphisms. Hence by the five-lemma again, so are the maps $\pi_n(X) \to \pi_n(Y)$.

By induction, the case of finite covers by $k > 2$ sets reduces to the case of covers by two sets, by letting one of the two sets be the union of the first $k-1$ of the given sets and the other be the k^{th} set. The case of infinite covers reduces to the finite case since for surjectivity of $\pi_n(X) \to \pi_n(Y)$, a map $S^n \to Y$ has compact image covered by finitely many V_i's, and similarly for injectivity. □

Quasifibrations

A map $p: E \to B$ with B path-connected is a **quasifibration** if the induced map $p_*: \pi_i(E, p^{-1}(b), x_0) \to \pi_i(B, b)$ is an isomorphism for all $b \in B$, $x_0 \in p^{-1}(b)$, and $i \geq 0$. We have shown in Theorem 4.41 that fiber bundles and fibrations have this property for $i > 0$, as a consequence of the homotopy lifting property, and the same reasoning applies for $i = 0$ since we assume B is path-connected.

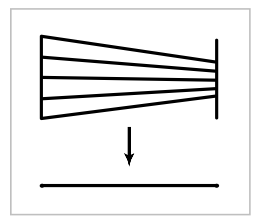

For example, consider the natural projection $M_f \to I$ of the mapping cylinder of a map $f: X \to Y$. This projection will be a quasifibration iff f is a weak homotopy equivalence, since the latter condition is equivalent to having $\pi_i(M_f, p^{-1}(b)) = 0 = \pi_i(I, b)$ for all i and all $b \in I$. Note that if f is not surjective, there are paths in I that do not lift to paths in M_f with a prescribed starting point, so p will not be a fibration in such cases.

An alternative condition for a map $p: E \to B$ to be a quasifibration is that the inclusion of each fiber $p^{-1}(b)$ into the homotopy fiber F_b of p over b is a weak homotopy equivalence. Recall that F_b is the space of all pairs (x, γ) with $x \in E$ and γ a path in B from $p(x)$ to b. The actual fiber $p^{-1}(b)$ is included in F_b as the pairs with $x \in p^{-1}(b)$ and γ the constant path at x. To see the equivalence of the two definitions, consider the commutative triangle at the right, where $F_b \to E_p \to B$ is the usual path-fibration construction applied to p. The right-hand map in the diagram is an isomorphism for all $i \geq 0$, and the upper map will be an isomorphism for all $i \geq 0$ iff the inclusion $p^{-1}(b) \hookrightarrow F_b$ is a weak equivalence since $E \simeq E_p$. Hence the two definitions are equivalent.

$$\begin{array}{ccc} \pi_i(E, p^{-1}(b)) & \longrightarrow & \pi_i(E_p, F_b) \\ & \searrow \quad \swarrow & \\ & \pi_i(B, b) & \end{array}$$

Recall from Proposition 4.61 that all fibers of a fibration over a path-connected base are homotopy equivalent. Since we are only considering quasifibrations over path-connected base spaces, this implies that all the fibers of a quasifibration have the same weak homotopy type. Quasifibrations over a base that is not path-connected are considered in the exercises, but we will not need this generality in what follows.

The following technical lemma gives various conditions for recognizing that a map is a quasifibration, which will be needed in the proof of the Dold–Thom theorem.

Lemma 4K.3. *A map $p: E \to B$ is a quasifibration if any one of the following conditions is satisfied*:

(a) *B can be decomposed as the union of open sets V_1 and V_2 such that each of the restrictions $p^{-1}(V_1) \to V_1$, $p^{-1}(V_2) \to V_2$, and $p^{-1}(V_1 \cap V_2) \to V_1 \cap V_2$ is a quasifibration.*

(b) *B is the union of an increasing sequence of subspaces $B_1 \subset B_2 \subset \cdots$ with the property that each compact set in B lies in some B_n, and such that each restriction $p^{-1}(B_n) \to B_n$ is a quasifibration.*

(c) *There is a deformation F_t of E into a subspace E', covering a deformation $\overline{F}_t$ of B into a subspace B', such that the restriction $E' \to B'$ is a quasifibration and $F_1: p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$ is a weak homotopy equivalence for each $b \in B$.*

By a 'deformation' in (c) we mean a deformation retraction in the weak sense as defined in the exercises for Chapter 0, where the homotopy is not required to be the identity on the subspace.

Proof: (a) To avoid some tedious details we will consider only the case that the fibers of p are path-connected, which will suffice for our present purposes, leaving the general case as an exercise for the reader. This hypothesis on fibers guarantees that all π_0's arising in the proof are trivial. In particular, by an exercise for §4.1 this allows us to terminate long exact sequences of homotopy groups of triples with zeros in the π_0 positions.

Let $U_1 = p^{-1}(V_1)$ and $U_2 = p^{-1}(V_2)$. The five-lemma for the long exact sequences of homotopy groups of the triples $(U_k, U_1 \cap U_2, p^{-1}(b))$ and $(V_k, V_1 \cap V_2, b)$ implies that the maps $\pi_i(U_k, U_1 \cap U_2) \to \pi_i(V_k, V_1 \cap V_2)$ are isomorphisms for $k = 1, 2$ and all i. Then Proposition 4K.1 implies that the maps $\pi_i(E, U_k) \to \pi_i(B, V_k)$ are isomorphisms for all choices of basepoints. The maps $\pi_i(U_k, p^{-1}(b)) \to \pi_i(V_k, b)$ are isomorphisms by hypothesis, so from the five-lemma we can then deduce that the maps $\pi_i(E, p^{-1}(b)) \to \pi_i(B, b)$ are isomorphisms for all $b \in V_k$, hence for all $b \in B$.

(b) Since each compact set in B lies in some B_n, each compact set in E lies in some subspace $E_n = p^{-1}(B_n)$, so $\pi_i(E, p^{-1}(b))$ is the direct limit $\varinjlim \pi_i(E_n, p^{-1}(b))$ just as $\pi_i(B, b) = \varinjlim \pi_i(B_n, b)$. It follows that the map $\pi_i(E, p^{-1}(b)) \to \pi_i(B, b)$ is an isomorphism since each of the maps $\pi_i(E_n, p^{-1}(b)) \to \pi_i(B_n, b)$ is an isomorphism by assumption. We can take the point b to be an arbitrary point in B and then discard any initial spaces B_n in the sequence that do not contain b, so we can assume b lies in B_n for all n.

(c) Consider the commutative diagram

$$\begin{array}{ccc} \pi_i(E, p^{-1}(b)) & \xrightarrow{F_{1*}} & \pi_i(E', p^{-1}(\overline{F}_1(b))) \\ \downarrow & & \downarrow \\ \pi_i(B, b) & \xrightarrow{\overline{F}_{1*}} & \pi_i(B', \overline{F}_1(b)) \end{array}$$

where b is an arbitrary point in B. The upper map in the diagram is an isomorphism by the five-lemma since the hypotheses imply that F_1 induces isomorphisms $\pi_i(E)\to\pi_i(E')$ and $\pi_i(p^{-1}(b))\to\pi_i(p^{-1}(\overline{F}_1(b)))$ for all i. The hypotheses also imply that the lower map and the right-hand map are isomorphisms. Hence the left-hand map is an isomorphism. □

Symmetric Products

Let us recall the definition from §3.C. For a space X the n-fold symmetric product $SP_n(X)$ is the quotient space of the product of n copies of X obtained by factoring out the action of the symmetric group permuting the factors. A choice of basepoint $e\in X$ gives inclusions $SP_n(X)\hookrightarrow SP_{n+1}(X)$ induced by $(x_1,\cdots,x_n)\mapsto(x_1,\cdots,x_n,e)$, and $SP(X)$ is defined to be the union of this increasing sequence of spaces, with the direct limit topology. Note that SP_n is a homotopy functor: A map $f:X\to Y$ induces $f_*:SP_n(X)\to SP_n(Y)$, and $f\simeq g$ implies $f_*\simeq g_*$. Hence $X\simeq Y$ implies $SP_n(X)\simeq SP_n(Y)$. In similar fashion SP is a homotopy functor on the category of basepointed spaces and basepoint-preserving homotopy classes of maps. It follows that $X\simeq Y$ implies $SP(X)\simeq SP(Y)$ for connected CW complexes X and Y since in this case requiring maps and homotopies to preserve basepoints does not affect the relation of homotopy equivalence.

Example 4K.4. An interesting special case is when $X=S^2$ because in this case $SP(S^2)$ can be identified with $\mathbb{C}\mathrm{P}^\infty$ in the following way. We first identify $\mathbb{C}\mathrm{P}^n$ with the nonzero polynomials of degree at most n with coefficients in $\mathbb{C}$, modulo scalar multiplication, by letting $a_0+\cdots+a_nz^n$ correspond to the line containing $(a_0,\cdots,a_n)$. The sphere S^2 we view as $\mathbb{C}\cup\{\infty\}$, and then we define $f:(S^2)^n\to\mathbb{C}\mathrm{P}^n$ by setting $f(a_1,\cdots,a_n)=(z+a_1)\cdots(z+a_n)$ with factors $z+\infty$ omitted, so in particular $f(\infty,\cdots,\infty)=1$. To check that f is continuous, suppose some a_i approaches ∞, say a_n, and all the other a_j's are finite. Then if we write

$$(z+a_1)\cdots(z+a_n)=$$
$$z^n+(a_1+\cdots+a_n)z^{n-1}+\cdots+\sum_{i_1<\cdots<i_k}a_{i_1}\cdots a_{i_k}z^{n-k}+\cdots+a_1\cdots a_n$$

we see that, dividing through by a_n and letting a_n approach ∞, this polynomial approaches $z^{n-1}+(a_1+\cdots+a_{n-1})z^{n-2}+\cdots+a_1\cdots a_{n-1}=(z+a_1)\cdots(z+a_{n-1})$. The same argument would apply if several a_i's approach ∞ simultaneously.

The value $f(a_1,\cdots,a_n)$ is unchanged under permutation of the a_i's, so there is an induced map $SP_n(S^2)\to\mathbb{C}\mathrm{P}^n$ which is a continuous bijection, hence a homeomorphism since both spaces are compact Hausdorff. Letting n go to ∞, we then get a homeomorphism $SP(S^2)\approx\mathbb{C}\mathrm{P}^\infty$.

The same argument can be used to show that $SP_n(S^1)\simeq S^1$ for all n, including $n=\infty$. Namely, the argument shows that $SP_n(\mathbb{C}-\{0\})$ can be identified with the polynomials $z^n + a_{n-1}z^{n-1} + \cdots + a_0$ with $a_0 \neq 0$, or in other words, the n-tuples $(a_0, \cdots, a_{n-1}) \in \mathbb{C}^n$ with $a_0 \neq 0$, and this subspace of $\mathbb{C}^n$ deformation retracts onto a circle.

The symmetric products of higher-dimensional spheres are more complicated, though things are not so bad for the 2-fold symmetric product:

Example 4K.5. Let us show that $SP_2(S^n)$ is homeomorphic to the mapping cone of a map $S^n\mathbb{R}\mathrm{P}^{n-1} \to S^n$ where $S^n\mathbb{R}\mathrm{P}^{n-1}$ is the n-fold unreduced suspension of $\mathbb{R}\mathrm{P}^{n-1}$. Hence $H_*(SP_2(S^n)) \approx H_*(S^n) \oplus \widetilde{H}_*(S^{n+1}\mathbb{R}\mathrm{P}^{n-1})$ from the long exact sequence of homology groups for the pair $(SP_2(S^n), S^n)$, since $SP_2(S^n)/S^n$ is $S^{n+1}\mathbb{R}\mathrm{P}^{n-1}$ with no reduced homology below dimension $n+2$.

If we view S^n as $D^n/\partial D^n$, then $SP_2(S^n)$ becomes a certain quotient of $D^n \times D^n$. Viewing $D^n \times D^n$ as the cone on its boundary $D^n \times \partial D^n \cup \partial D^n \times D^n$, the identifications that produce $SP_2(S^n)$ respect the various concentric copies of this boundary which fill up the interior of $D^n \times D^n$, so it suffices to analyze the identifications in all these copies of the boundary. The identifications on the boundary of $D^n \times D^n$ itself yield S^n. This is clear since the identification $(x,y) \sim (y,x)$ converts $D^n \times \partial D^n \cup \partial D^n \times D^n$ to $D^n \times \partial D^n$, and all points of ∂D^n are identified in S^n.

It remains to see that the identifications $(x,y) \sim (y,x)$ on each concentric copy of the boundary in the interior of $D^n \times D^n$ produce $S^n\mathbb{R}\mathrm{P}^{n-1}$. Denote by Z the quotient of $D^n \times \partial D^n \cup \partial D^n \times D^n$ under these identifications. This is the same as the quotient of $D^n \times \partial D^n$ under the identifications $(x,y) \sim (y,x)$ for $(x,y) \in \partial D^n \times \partial D^n$.

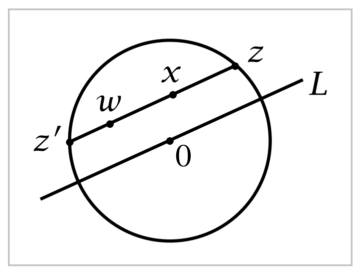

Define $f : D^n \times \mathbb{R}\mathrm{P}^{n-1} \to Z$ by $f(x,L) = (w,z)$ where x is equidistant from $z \in \partial D^n$ and $w \in D^n$ along the line through x parallel to L, as in the figure. If x is the midpoint of the segment zz' then $w = z'$ and there is no way to distinguish between w and z, but since f takes values in the quotient space Z, this is not a problem. If $x \in \partial D^n$ then $w = z = x$, independent of L. If $x \in D^n - \partial D^n$ then $w \neq z$, and conversely, given $(w,z) \in D^n \times \partial D^n$ with $w \neq z$ there is a unique (x,L) with $f(x,L) = (w,z)$, namely x is the midpoint of the segment wz and L is the line parallel to this segment. In view of these remarks, we see that Z is the quotient space of $D^n \times \mathbb{R}\mathrm{P}^{n-1}$ under the identifications $(x,L) \sim (x,L')$ if $x \in \partial D^n$. This quotient is precisely $S^n\mathbb{R}\mathrm{P}^{n-1}$.

This example illustrates that passing from a CW structure on X to a CW structure on $SP_n(X)$ or $SP(X)$ is not at all straightforward. However, if X is a simplicial complex, there is a natural way of putting Δ-complex structures on $SP_n(X)$ and $SP(X)$, as follows. A simplicial complex structure on X gives a CW structure on the product of n copies of X, with cells n-fold products of simplices. Such a product has a canonical barycentric subdivision as a simplicial complex, with vertices the points whose coordinates are barycenters of simplices of X. By induction over skeleta, this just amounts to coning off a simplicial structure on the boundary of each product cell. This simplicial structure on the product of n copies of X is in fact a Δ-complex structure since the vertices of each of its simplices have a natural ordering given by the dimensions of the cells of which they are barycenters. The action of the symmetric group permuting coordinates respects this Δ-complex structure, taking simplices homeomorphically to simplices, preserving vertex-orderings, so there is an induced Δ-complex structure on the quotient $SP_n(X)$. With the basepoint of X chosen to be a vertex, $SP_n(X)$ is a subcomplex of $SP_{n+1}(X)$ so there is a natural Δ-complex structure on the infinite symmetric product $SP(X)$ as well.

As usual with products, the CW topology on $SP_n(X)$ and $SP(X)$ is in general different from the topology arising from the original definition in terms of product topologies, but one can check that the two topologies have the same compact sets, so the distinction will not matter for our present purposes. For definiteness, we will use the CW topology in what follows, which means restricting X to be a simplicial complex. Since every CW complex is homotopy equivalent to a simplicial complex by Theorem 2C.5, and SP_n and SP are homotopy functors, there is no essential loss of generality in restricting from CW complexes to simplicial complexes.

Here is the main result of this section, the Dold–Thom theorem:

Theorem 4K.6. *The functor $X \mapsto \pi_i SP(X)$ for $i \ge 1$ coincides with the functor $X \mapsto H_i(X;\mathbb{Z})$ on the category of basepointed connected CW complexes.*

In particular this says that $SP(S^n)$ is a $K(\mathbb{Z},n)$, and more generally that for a Moore space $M(G,n)$, $SP(M(G,n))$ is a $K(G,n)$.

The fact that $SP(X)$ is a commutative, associative H–space with a strict identity element limits its weak homotopy type considerably:

Corollary 4K.7. *A path-connected, commutative, associative H–space with a strict identity element has the weak homotopy type of a product of Eilenberg–MacLane spaces.*

In particular, if X is a connected CW complex then $SP(X)$ is path-connected and has the weak homotopy type of $\prod_n K(H_n(X),n)$. Thus the functor SP essentially reduces to Eilenberg–MacLane spaces.

Proof: Let X be a path-connected, commutative, associative H–space with a strict identity element, and let $G_n = \pi_n(X)$. By Lemma 4.31 there is a map $M(G_n,n) \to X$ inducing an isomorphism on π_n when $n > 1$ and an isomorphism on H_1 when $n = 1$. We can take these maps to be basepoint-preserving, and then they combine to give a map $\bigvee_n M(G_n,n) \to X$. The very special H–space structure on X allows us to extend this to a homomorphism $f : SP(\bigvee_n M(G_n,n)) \to X$. In general, $SP(\bigvee_\alpha X_\alpha)$ can be identified with $\prod_\alpha SP(X_\alpha)$ where this is the 'weak' infinite product, the union of the finite products. This, together with the general fact that the map

$\pi_i(X) \to \pi_i SP(X) = H_i(X;\mathbb{Z})$ induced by the inclusion $X = SP_1(X) \hookrightarrow SP(X)$ is the Hurewicz homomorphism, as we will see at the end of the proof of the Dold–Thom theorem, implies that the map f induces an isomorphism on all homotopy groups. Thus we have a weak homotopy equivalence $\prod_n SP(M(G_n,n)) \to X$, and as we noted above, $SP(M(G_n,n))$ is a $K(G_n,n)$. Finally, since each factor $SP(M(G_n,n))$ has only one nontrivial homotopy group, the weak infinite product has the same weak homotopy type as the ordinary infinite product. $\square$

The main step in the proof of the theorem will be to show that for a simplicial pair (X,A) with both X and A connected, there is a long exact sequence

$$\cdots \to \pi_i SP(A) \to \pi_i SP(X) \to \pi_i SP(X/A) \to \pi_{i-1} SP(A) \to \cdots$$

This would follow if the maps $SP(A) \to SP(X) \to SP(X/A)$ formed a fiber bundle or fibration. There is some reason to think this might be true, because all the fibers of the projection $SP(X) \to SP(X/A)$ are homeomorphic to $SP(A)$. In fact, in terms of the H–space structure on $SP(X)$ as the free abelian monoid generated by X, the fibers are exactly the cosets of the submonoid $SP(A)$. The projection $SP(X) \to SP(X/A)$, however, fails to have the homotopy lifting property, even the special case of lifting paths. For if x_t, $t \in [0,1)$, is a path in $X - A$ approaching a point $x_1 = a \in A$ other than the basepoint, then regarding x_t as a path in $SP(X/A)$, any lift to $SP(X)$ would have the form $x_t\alpha_t$, $\alpha_t \in SP(A)$, ending at $x_1\alpha_1 = a\alpha_1$, a point of $SP(A)$ which is a multiple of a, so in particular there would be no lift ending at the basepoint of $SP(X)$.

What we will show is that the projection $SP(X) \to SP(X/A)$ has instead the weaker structure of a quasifibration, which is still good enough to deduce a long exact sequence of homotopy groups.

Proof of 4K.6: As we have said, the main step will be to associate a long exact sequence of homotopy groups to each simplicial pair (X,A) with X and A connected. This will be the long exact sequence of homotopy groups coming from the quasifibration $SP(A) \to SP(X) \to SP(X/A)$, so the major work will be in verifying the quasifibration property. Since SP is a homotopy functor, we are free to replace (X,A) by a homotopy equivalent pair, so let us replace (X,A) by (M,A) where M is the mapping cylinder of the inclusion $A \hookrightarrow X$. This new pair, which we still call (X,A), has some slight technical advantages, as we will see later in the proof.

To begin the proof that the projection $p: SP(X) \to SP(X/A)$ is a quasifibration, let $B_n = SP_n(X/A)$ and $E_n = p^{-1}(B_n)$. Thus E_n consists of those points in $SP(X)$ having at most n coordinates in $X - A$. By Lemma 4K.3(b) it suffices to show that $p: E_n \to B_n$ is a quasifibration. The proof of the latter fact will be by induction on n, starting with the trivial case $n = 0$ when B_0 is a point. The induction step will consist of showing that p is a quasifibration over a neighborhood of B_{n-1} and over

$B_n - B_{n-1}$, then applying Lemma 4K.3(a). We first tackle the problem of showing the quasifibration property over a neighborhood of B_{n-1}.

Let $f_t : X \to X$ be a homotopy of the identity map deformation retracting a neighborhood N of A onto A. Since we have replaced the original X by the mapping cylinder of the inclusion $A \hookrightarrow X$, we can take f_t simply to slide points along the segments $\{a\}\times I$ in the mapping cylinder, with $N = A\times[0, 1/2)$. Let $U \subset E_n$ consist of those points having at least one coordinate in N, or in other words, products with at least one factor in N. Thus U is a neighborhood of E_{n-1} in E_n and $p(U)$ is a neighborhood of B_{n-1} in B_n.

The homotopy f_t induces a homotopy $F_t : E_n \to E_n$ whose restriction to U is a deformation of U into E_{n-1}, where by 'deformation' we mean deformation retraction in the weak sense. Since f_t is the identity on A, F_t is the lift of a homotopy $\overline{F}_t : B_n \to B_n$ which restricts to a deformation of $\overline{U} = p(U)$ into B_{n-1}. We will deduce that the projection $U \to \overline{U}$ is a quasifibration by using Lemma 4K.3(c). To apply this to the case at hand we need to verify that $F_1 : p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$ is a weak equivalence for all b. Each point $w \in p^{-1}(b)$ is a commuting product of points in X. Let $\widehat{w}$ be the subproduct whose factors are points in $X - A$, so we have $w = \widehat{w}v$ for v a product of points in A. Since f_1 is the identity on A and F_1 is a homomorphism, we have $F_1(w) = F_1(\widehat{w})v$, which can be written $\widehat{F_1(\widehat{w})}v'v$ with v' also a product of points in A. If we fix $\widehat{w}$ and let v vary over $SP(A)$, we get all points of $p^{-1}(b)$ exactly once, or in other words, we have $p^{-1}(b)$ expressed as the coset $\widehat{w}SP(A)$. The map F_1, $\widehat{w}v \mapsto \widehat{F_1(\widehat{w})}v'v$, takes this coset to the coset $\widehat{F_1(\widehat{w})}SP(A)$ by a map that would be a homeomorphism if the factor v' were not present. Since A is connected, there is a path v'_t from v' to the basepoint, and so by replacing v' with v'_t in the product $\widehat{F_1(\widehat{w})}v'v$ we obtain a homotopy from $F_1 : p^{-1}(b) \to p^{-1}(\overline{F}_1(b))$ to a homeomorphism, so this map is a homotopy equivalence, as desired.

It remains to see that p is a quasifibration over $B_n - B_{n-1}$ and over the intersection of this set with $\overline{U}$. The argument will be the same in both cases.

Identifying $B_n - B_{n-1}$ with $SP_n(X - A)$, the projection $p : E_n - E_{n-1} \to B_n - B_{n-1}$ is the same as the operator $w \mapsto \widehat{w}$. The inclusion $SP_n(X - A) \hookrightarrow E_n - E_{n-1}$ gives a section for $p : E_n - E_{n-1} \to B_n - B_{n-1}$, so $p_* : \pi_i(E_n - E_{n-1}, p^{-1}(b)) \to \pi_i(B_n - B_{n-1}, b)$ is surjective. To see that it is also injective, represent an element of its kernel by a map $g : (D^i, \partial D^i) \to (E_n - E_{n-1}, p^{-1}(b))$. A nullhomotopy of pg gives a homotopy of g changing only its coordinates in $X - A$. This homotopy is through maps $(D^i, \partial D^i) \to (E_n - E_{n-1}, p^{-1}(b))$, and ends with a map to $p^{-1}(b)$, so the kernel of p_* is trivial. Thus the projection $E_n - E_{n-1} \to B_n - B_{n-1}$ is a quasifibration, at least if $B_n - B_{n-1}$ is path-connected. But by replacing the original X with the mapping cylinder of the inclusion $A \hookrightarrow X$, we guarantee that $X - A$ is path-connected since it deformation retracts onto X. Hence the space $B_n - B_{n-1} = SP_n(X - A)$ is also path-connected.

This argument works equally well over any open subset of $B_n - B_{n-1}$ that is path-connected, in particular over $U \cap (B_n - B_{n-1})$, so via Lemma 4K.3(a) this finishes the proof that $SP(A) \to SP(X) \to SP(X/A)$ is a quasifibration.

Since the homotopy axiom is obvious, this gives us the first two of the three axioms needed for the groups $h_i(X) = \pi_i SP(X)$ to define a reduced homology theory. There remains only the wedge sum axiom, $h_i(\bigvee_\alpha X_\alpha) \approx \bigoplus_\alpha h_i(X_\alpha)$, but this is immediate from the evident fact that $SP(\bigvee_\alpha X_\alpha) = \prod_\alpha SP(X_\alpha)$, where this is the 'weak' product, the union of the products of finitely many factors.

The homology theory $h_*(X)$ is defined on the category of connected, basepointed simplicial complexes, with basepoint-preserving maps. The coefficients of this homology theory, the groups $h_i(S^n)$, are the same as for ordinary homology with $\mathbb{Z}$ coefficients since we know this is true for $n = 2$ by the homeomorphism $SP(S^2) \approx \mathbb{C}\mathrm{P}^\infty$, and there are isomorphisms $h_i(X) \approx h_{i+1}(\Sigma X)$ in any reduced homology theory. If the homology theory $h_*(X)$ were defined on the category of all simplicial complexes, without basepoints, then Theorem 4.59 would give natural isomorphisms $h_i(X) \approx H_i(X;\mathbb{Z})$ for all X, and the proof would be complete. However, it is easy to achieve this by defining a new homology theory $h'_i(X) = h_{i+1}(\Sigma X)$, since the suspension of an arbitrary complex is connected and the suspension of an arbitrary map is basepoint-preserving, taking the basepoint to be one of the suspension points. Since $h'_i(X)$ is naturally isomorphic to $h_i(X)$ if X is connected, we are done. $\square$

It is worth noting that the map $\pi_i(X) \to \pi_i SP(X) = H_i(X;\mathbb{Z})$ induced by the inclusion $X = SP_1(X) \hookrightarrow SP(X)$ is the Hurewicz homomorphism. For by definition of the Hurewicz homomorphism and naturality this reduces to the case $X = S^i$, where the map $SP_1(S^i) \hookrightarrow SP(S^i)$ induces on π_i a homomorphism $\mathbb{Z} \to \mathbb{Z}$, which one just has to check is an isomorphism, the Hurewicz homomorphism being determined only up to sign. The suspension isomorphism gives a further reduction to the case $i = 1$, where the inclusion $SP_1(S^1) \hookrightarrow SP(S^1)$ is a homotopy equivalence, hence induces an isomorphism on π_1.

Exercises

1. Show that Corollary 4K.2 remains valid when X and Y are CW complexes and the subspaces U_i and V_i are subcomplexes rather than open sets.

2. Show that a simplicial map $f : K \to L$ is a homotopy equivalence if $f^{-1}(x)$ is contractible for all $x \in L$. [Consider the cover of L by open stars of simplices and the cover of K by the preimages of these open stars.]

3. Show that $SP_n(I) = \Delta^n$.

4. Show that $SP_2(S^1)$ is a Möbius band, and that this is consistent with the description of $SP_2(S^n)$ as a mapping cone given in Example 4K.5.

5. A map $p:E\to B$ with B not necessarily path-connected is defined to be a quasifibration if the following equivalent conditions are satisfied:

(i) For all $b\in B$ and $x_0\in p^{-1}(b)$, the map $p_*:\pi_i(E,p^{-1}(b),x_0)\to\pi_i(B,b)$ is an isomorphism for $i>0$ and $\pi_0(p^{-1}(b),x_0)\to\pi_0(E,x_0)\to\pi_0(B,b)$ is exact.

(ii) The inclusion of the fiber $p^{-1}(b)$ into the homotopy fiber F_b of p over b is a weak homotopy equivalence for all $b\in B$.

(iii) The restriction of p over each path-component of B is a quasifibration according to the definition in this section.

Show these three conditions are equivalent, and prove Lemma 4K.3 for quasifibrations over non-pathconnected base spaces.

6. Let X be a complex of spaces over a Δ-complex Γ, as defined in §4.G. Show that the natural projection $\Delta X\to\Gamma$ is a quasifibration if all the maps in X associated to edges of Γ are weak homotopy equivalences.

4.L Steenrod Squares and Powers

The main objects of study in this section are certain homomorphisms called **Steenrod squares** and **Steenrod powers**:

$$Sq^i:H^n(X;\mathbb{Z}_2)\to H^{n+i}(X;\mathbb{Z}_2)$$
$$P^i:H^n(X;\mathbb{Z}_p)\to H^{n+2i(p-1)}(X;\mathbb{Z}_p)\quad\text{for odd primes } p$$

The terms 'squares' and 'powers' arise from the fact that Sq^i and P^i are related to the maps $\alpha\mapsto\alpha^2$ and $\alpha\mapsto\alpha^p$ sending a cohomology class α to the 2-fold or p-fold cup product with itself. Unlike cup products, however, the operations Sq^i and P^i are stable, that is, invariant under suspension.

The operations Sq^i generate an algebra $\mathcal{A}_2$, called the Steenrod algebra, such that $H^*(X;\mathbb{Z}_2)$ is a module over $\mathcal{A}_2$ for every space X, and maps between spaces induce homomorphisms of $\mathcal{A}_2$-modules. Similarly, for odd primes p, $H^*(X;\mathbb{Z}_p)$ is a module over a corresponding Steenrod algebra $\mathcal{A}_p$ generated by the P^i's and Bockstein homomorphisms. Like the ring structure given by cup product, these module structures impose strong constraints on spaces and maps. For example, we will use them to show that there do not exist spaces X with $H^*(X;\mathbb{Z})$ a polynomial ring $\mathbb{Z}[\alpha]$ unless α has dimension 2 or 4, where there are the familiar examples of $\mathbb{C}\mathrm{P}^\infty$ and $\mathbb{H}\mathrm{P}^\infty$.

This rather lengthy section is divided into two main parts. The first part describes the basic properties of Steenrod squares and powers and gives a number of examples and applications. The second part is devoted to constructing the squares and powers and showing they satisfy the basic properties listed in the first part. More extensive applications will be given in [SSAT] after spectral sequences have been introduced. Most applications of Steenrod squares and powers do not depend on how these operations are actually constructed, but only on their basic properties. This is similar to the situation for ordinary homology and cohomology, where the axioms generally suffice for most applications. The construction of Steenrod squares and powers and the verification of their basic properties, or axioms, is rather interesting in its own way, but does involve a certain amount of work, particularly for the Steenrod powers, and this is why we delay the work until later in the section.

We begin with a few generalities. A **cohomology operation** is a transformation $\Theta = \Theta_X : H^m(X;G) \to H^n(X;H)$ defined for all spaces X, with a fixed choice of m, n, G, and H, and satisfying the naturality property that for all maps $f : X \to Y$ there is a commuting diagram as shown at the right.

$$\begin{array}{ccc} H^m(Y;G) & \xrightarrow{\Theta_Y} & H^n(Y;H) \\ \downarrow f^* & & \downarrow f^* \\ H^m(X;G) & \xrightarrow{\Theta_X} & H^n(X;H) \end{array}$$

For example, with coefficients in a ring R the transformation $H^m(X;R) \to H^{mp}(X;R)$, $\alpha \mapsto \alpha^p$, is a cohomology operation since $f^*(\alpha^p) = (f^*(\alpha))^p$. Taking $R = \mathbb{Z}$, this example shows that cohomology operations need not be homomorphisms. On the other hand, when $R = \mathbb{Z}_p$ with p prime, the operation $\alpha \mapsto \alpha^p$ is a homomorphism. Other examples of cohomology operations we have already encountered are the Bockstein homomorphisms defined in §3.E. As a more trivial example, a homomorphism $G \to H$ induces change-of-coefficient homomorphisms $H^m(X;G) \to H^m(X;H)$ which can be viewed as cohomology operations.

In spite of their rather general definition, cohomology operations can be described in somewhat more concrete terms:

Proposition 4L.1. *For fixed m, n, G, and H there is a bijection between the set of all cohomology operations $\Theta : H^m(X;G) \to H^n(X;H)$ and $H^n(K(G,m);H)$, defined by $\Theta \mapsto \Theta(\iota)$ where $\iota \in H^m(K(G,m);G)$ is a fundamental class.*

Proof: Via CW approximations to spaces, it suffices to restrict attention to CW complexes, so we can identify $H^m(X;G)$ with $\langle X, K(G,m)\rangle$ when $m > 0$ by Theorem 4.57, and with $[X, K(G,0)]$ when $m = 0$. If an element $\alpha \in H^m(X;G)$ corresponds to a map $\varphi : X \to K(G,m)$, so $\varphi^*(\iota) = \alpha$, then $\Theta(\alpha) = \Theta(\varphi^*(\iota)) = \varphi^*(\Theta(\iota))$ and Θ is uniquely determined by $\Theta(\iota)$. Thus $\Theta \mapsto \Theta(\iota)$ is injective. For surjectivity, given an element $\alpha \in H^n(K(G,m);H)$ corresponding to a map $\theta : K(G,m) \to K(H,n)$, then composing with θ defines a transformation $\langle X, K(G,m)\rangle \to \langle X, K(H,n)\rangle$, that is, $\Theta : H^m(X;G) \to H^n(X;H)$, with $\Theta(\iota) = \alpha$. The naturality property for Θ amounts to associativity of the compositions $X \xrightarrow{f} Y \xrightarrow{\varphi} K(G,m) \xrightarrow{\theta} K(H,n)$ and so Θ is a cohomology operation. □

A consequence of the proposition is that cohomology operations that decrease dimension are all rather trivial since $K(G,m)$ is $(m-1)$-connected. Moreover, since

$H^m(K(G,m);H) \approx \mathrm{Hom}(G,H)$, it follows that the only cohomology operations that preserve dimension are given by coefficient homomorphisms.

The Steenrod squares $Sq^i : H^n(X;\mathbb{Z}_2) \to H^{n+i}(X;\mathbb{Z}_2)$, $i \geq 0$, will satisfy the following list of properties, beginning with naturality:

(1) $Sq^i(f^*(\alpha)) = f^*(Sq^i(\alpha))$ for $f : X \to Y$.
(2) $Sq^i(\alpha + \beta) = Sq^i(\alpha) + Sq^i(\beta)$.
(3) $Sq^i(\alpha \smile \beta) = \sum_j Sq^j(\alpha) \smile Sq^{i-j}(\beta)$ (the Cartan formula).
(4) $Sq^i(\sigma(\alpha)) = \sigma(Sq^i(\alpha))$ where $\sigma : H^n(X;\mathbb{Z}_2) \to H^{n+1}(\Sigma X;\mathbb{Z}_2)$ is the suspension isomorphism given by reduced cross product with a generator of $H^1(S^1;\mathbb{Z}_2)$.
(5) $Sq^i(\alpha) = \alpha^2$ if $i = |\alpha|$, and $Sq^i(\alpha) = 0$ if $i > |\alpha|$.
(6) $Sq^0 = \mathbb{1}$, the identity.
(7) Sq^1 is the $\mathbb{Z}_2$ Bockstein homomorphism β associated with the coefficient sequence $0 \to \mathbb{Z}_2 \to \mathbb{Z}_4 \to \mathbb{Z}_2 \to 0$.

The first part of (5) says that the Steenrod squares extend the squaring operation $\alpha \mapsto \alpha^2$, which has the nice feature of being a homomorphism with $\mathbb{Z}_2$ coefficients. Property (4) says that the Sq^i's are stable operations, invariant under suspension. The actual squaring operation $\alpha \mapsto \alpha^2$ does not have this property since in a suspension ΣX all cup products of positive-dimensional classes are zero, according to an exercise for §3.2.

The fact that Steenrod squares are stable operations extending the cup product square yields the following theorem, which implies that the stable homotopy groups of spheres π_1^s, π_3^s, and π_7^s are nontrivial:

Theorem 4L.2. *If $f : S^{2n-1} \to S^n$ has Hopf invariant 1, then $[f] \in \pi_{n-1}^s$ is nonzero, so the iterated suspensions $\Sigma^k f : S^{2n+k-1} \to S^{n+k}$ are all homotopically nontrivial.*

Proof: Associated to a map $f : S^\ell \to S^m$ is the mapping cone $C_f = S^m \cup_f e^{\ell+1}$ with the cell $e^{\ell+1}$ attached via f. Assuming f is basepoint-preserving, we have the relation $C_{\Sigma f} = \Sigma C_f$ where Σ denotes reduced suspension.

If $f : S^{2n-1} \to S^n$ has Hopf invariant 1, then by (5), $Sq^n : H^n(C_f;\mathbb{Z}_2) \to H^{2n}(C_f;\mathbb{Z}_2)$ is nontrivial. By (4) the same is true for $Sq^n : H^{n+k}(\Sigma^k C_f;\mathbb{Z}_2) \to H^{2n+k}(\Sigma^k C_f;\mathbb{Z}_2)$ for all k. If $\Sigma^k f$ were homotopically trivial we would have a retraction $r : \Sigma^k C_f \to S^{n+k}$. The diagram at the right would then commute by naturality of Sq^n, but since the group in the lower left corner of the diagram is zero, this gives a contradiction. $\square$

$$\begin{array}{ccc} H^{n+k}(S^{n+k};\mathbb{Z}_2) & \xrightarrow[\approx]{r^*} & H^{n+k}(\Sigma^k C_f;\mathbb{Z}_2) \\ \downarrow{\scriptstyle Sq^n} & & \downarrow{\scriptstyle Sq^n} \\ H^{2n+k}(S^{n+k};\mathbb{Z}_2) & \xrightarrow{r^*} & H^{2n+k}(\Sigma^k C_f;\mathbb{Z}_2) \end{array}$$

The Steenrod power operations $P^i : H^n(X;\mathbb{Z}_p) \to H^{n+2i(p-1)}(X;\mathbb{Z}_p)$ for p an odd prime will satisfy analogous properties:

(1) $P^i(f^*(\alpha)) = f^*(P^i(\alpha))$ for $f : X \to Y$.

(2) $P^i(\alpha+\beta)=P^i(\alpha)+P^i(\beta)$.

(3) $P^i(\alpha\smile\beta)=\sum_j P^j(\alpha)\smile P^{i-j}(\beta)$ (the Cartan formula).

(4) $P^i(\sigma(\alpha))=\sigma(P^i(\alpha))$ where $\sigma:H^n(X;\mathbb{Z}_p)\to H^{n+1}(\Sigma X;\mathbb{Z}_p)$ is the suspension isomorphism given by reduced cross product with a generator of $H^1(S^1;\mathbb{Z}_p)$.

(5) $P^i(\alpha)=\alpha^p$ if $2i=|\alpha|$, and $P^i(\alpha)=0$ if $2i>|\alpha|$.

(6) $P^0=1\!\!1$, the identity.

The germinal property $P^i(\alpha)=\alpha^p$ in (5) can only be expected to hold for even-dimensional classes α since for odd-dimensional α the commutativity property of cup product implies that $\alpha^2=0$ with $\mathbb{Z}_p$ coefficients if p is odd, and then $\alpha^p=0$ since $\alpha^2=0$. Note that the formula $P^i(\alpha)=\alpha^p$ for $|\alpha|=2i$ implies that P^i raises dimension by $2i(p-1)$, explaining the appearance of this number.

The Bockstein homomorphism $\beta:H^n(X;\mathbb{Z}_p)\to H^{n+1}(X;\mathbb{Z}_p)$ is not included as one of the P^i's, but this is mainly a matter of notational convenience. As we shall see later when we discuss Adem relations, the operation Sq^{2i+1} is the same as the composition $Sq^1Sq^{2i}=\beta Sq^{2i}$, so the Sq^{2i}'s can be regarded as the P^i's for $p=2$.

One might ask if there are elements of π_*^s detectable by Steenrod powers in the same way that the Hopf maps are detected by Steenrod squares. The answer is yes for the operation P^1, as we show in Example 4L.6. It is a perhaps disappointing fact that no other squares or powers besides Sq^1, Sq^2, Sq^4, Sq^8, and P^1 detect elements of homotopy groups of spheres. (Sq^1 detects a map $S^n\to S^n$ of degree 2.) We will prove this for certain Sq^i's and P^i's later in this section. The general case for $p=2$ is Adams' theorem on the Hopf invariant discussed in §4.B, while the case of odd p is proved in [Adams & Atiyah 1966]; see also [VBKT].

The Cartan formulas can be expressed in a more concise form by defining *total* Steenrod square and power operations by $Sq=Sq^0+Sq^1+\cdots$ and $P=P^0+P^1+\cdots$. These act on $H^*(X;\mathbb{Z}_p)$ since by property (5), only a finite number of Sq^i's or P^i's can be nonzero on a given cohomology class. The Cartan formulas then say that $Sq(\alpha\smile\beta)=Sq(\alpha)\smile Sq(\beta)$ and $P(\alpha\smile\beta)=P(\alpha)\smile P(\beta)$, so Sq and P are ring homomorphisms.

We can use Sq and P to compute the operations Sq^i and P^i for projective spaces and lens spaces via the following general formulas:

$$(*)\quad \begin{aligned} Sq^i(\alpha^n)&=\binom{n}{i}\alpha^{n+i} \text{ for } \alpha\in H^1(X;\mathbb{Z}_2)\\ P^i(\alpha^n)&=\binom{n}{i}\alpha^{n+i(p-1)} \text{ for } \alpha\in H^2(X;\mathbb{Z}_p)\end{aligned}$$

To derive the first formula, properties (5) and (6) give $Sq(\alpha)=\alpha+\alpha^2=\alpha(1+\alpha)$, so $Sq(\alpha^n)=Sq(\alpha)^n=\alpha^n(1+\alpha)^n=\sum_i\binom{n}{i}\alpha^{n+i}$ and hence $Sq^i(\alpha^n)=\binom{n}{i}\alpha^{n+i}$. The second formula is obtained in similar fashion: $P(\alpha)=\alpha+\alpha^p=\alpha(1+\alpha^{p-1})$ so $P(\alpha^n)=\alpha^n(1+\alpha^{p-1})^n=\sum_i\binom{n}{i}\alpha^{n+i(p-1)}$.

In Lemma 3C.6 we described how binomial coefficients can be computed modulo a prime p:

$$\binom{m}{n} \equiv \prod_i \binom{m_i}{n_i} \bmod p, \qquad \text{where } m = \sum_i m_i p^i \text{ and } n = \sum_i n_i p^i \text{ are the } p\text{-adic expansions of } m \text{ and } n.$$

When $p = 2$ for example, the extreme cases of a dyadic expansion consisting of a single 1 or all 1's give

$$Sq(\alpha^{2^k}) = \alpha^{2^k} + \alpha^{2^{k+1}}$$
$$Sq(\alpha^{2^k-1}) = \alpha^{2^k-1} + \alpha^{2^k} + \alpha^{2^k+1} + \cdots + \alpha^{2^{k+1}-2}$$

for $\alpha \in H^1(X;\mathbb{Z}_2)$. More generally, the coefficients of $Sq(\alpha^n)$ can be read off from the $(n+1)^{st}$ row of the mod 2 Pascal triangle, a portion of which is shown in the figure at the right, where dots denote zeros.

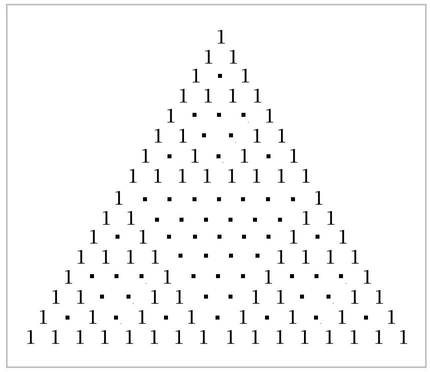

Example 4L.3: Stable Splittings. The formula $(*)$ tells us how to compute Steenrod squares for $\mathbb{R}P^\infty$, hence also for any suspension of $\mathbb{R}P^\infty$. The explicit formulas for $Sq(\alpha^{2^k})$ and $Sq(\alpha^{2^k-1})$ above show that all the powers of the generator $\alpha \in H^1(\mathbb{R}P^\infty;\mathbb{Z}_2)$ are tied together by Steenrod squares since the first formula connects α inductively to all the powers α^{2^k} and the second formula connects these powers to all the other powers. This shows that no suspension $\Sigma^k\mathbb{R}P^\infty$ has the homotopy type of a wedge sum $X \vee Y$ with both X and Y having nontrivial cohomology. In the case of $\mathbb{R}P^\infty$ itself we could have deduced this from the ring structure of $H^*(\mathbb{R}P^\infty;\mathbb{Z}_2) \approx \mathbb{Z}_2[\alpha]$, but cup products become trivial in a suspension.

The same reasoning shows that $\mathbb{C}P^\infty$ and $\mathbb{H}P^\infty$ have no nontrivial stable splittings. The $\mathbb{Z}_2$ cohomology in these cases is again $\mathbb{Z}_2[\alpha]$, though with α no longer 1-dimensional. However, we still have $Sq(\alpha) = \alpha + \alpha^2$ since these spaces have no nontrivial cohomology in the dimensions between α and α^2, so we have $Sq^{2i}(\alpha^n) = \binom{n}{i}\alpha^{n+i}$ for $\mathbb{C}P^\infty$ and $Sq^{4i}(\alpha^n) = \binom{n}{i}\alpha^{n+i}$ for $\mathbb{H}P^\infty$. Then the arguments from the real case carry over using the operations Sq^{2i} and Sq^{4i} in place of Sq^i.

Suppose we consider the same question for $K(\mathbb{Z}_3,1)$ instead of $\mathbb{R}P^\infty$. Taking cohomology with $\mathbb{Z}_3$ coefficients, the Bockstein β is nonzero on odd-dimensional classes in $H^*(K(\mathbb{Z}_3,1);\mathbb{Z}_3)$, thus tying them to the even-dimensional classes, so we only need to see which even-dimensional classes are connected by P^i's. The even-dimensional part of $H^*(K(\mathbb{Z}_3,1);\mathbb{Z}_3)$ is a polynomial algebra $\mathbb{Z}_3[\alpha]$ with $|\alpha| = 2$, so we have $P^i(\alpha^n) = \binom{n}{i}\alpha^{n+i(p-1)} = \binom{n}{i}\alpha^{n+2i}$ by our earlier formula. Since P^i raises dimension by $4i$ when $p = 3$, there is no chance that all the even-dimensional cohomology will be connected by the P^i's. In fact, we showed in Proposition 4I.3 that $\Sigma K(\mathbb{Z}_3,1) \simeq X_1 \vee X_2$ where X_1 has the cohomology of $\Sigma K(\mathbb{Z}_3,1)$ in dimensions congruent to 2 and 3 mod 4, while X_2 has the remaining cohomology. Thus the best one could hope would be that all the odd powers of α are connected by P^i's and likewise all the even powers are connected, since this would imply that neither X_1 nor X_2 splits nontrivially. This is indeed the case, as one sees by an examination of the coefficients in the formula $P^i(\alpha^n) = \binom{n}{i}\alpha^{n+2i}$. In the Pascal triangle mod 3, shown at the right, $P(\alpha^n)$ is determined by the $(n+1)^{st}$ row. For example the sixth row says that $P(\alpha^5) = \alpha^5 + 2\alpha^7 + \alpha^9 + \alpha^{11} + 2\alpha^{13} + \alpha^{15}$. A few moments' thought shows that the rows that compute $P(\alpha^n)$ for $n = 3^k m - 1$ have all nonzero entries, and these rows together with the rows right after them suffice to connect the powers of α in the desired way, so X_1 and X_2 have no stable splittings. One can also see that $\Sigma^2 X_1$ and X_2 are not homotopy equivalent, even stably, since the operations P^i act differently in the two spaces. For example P^2 is trivial on suspensions of α but not on suspensions of α^2.

```
                1
               1 1
              1 2 1
             1 · · 1
            1 1 · 1 1
           1 2 1 1 2 1
          1 · · 2 · · 1
         1 1 · 2 2 · 1 1
        1 2 1 2 1 2 1 2 1
       1 · · · · · · · · · 1
```

The situation for $K(\mathbb{Z}_p,1)$ for larger primes p is entirely similar, with $\Sigma K(\mathbb{Z}_p,1)$ splitting as a wedge sum of $p-1$ spaces. The same arguments work more generally for $K(\mathbb{Z}_{p^i},1)$, though for $i>1$ the usual Bockstein β is identically zero so one has to use instead a Bockstein involving $\mathbb{Z}_{p^i}$ coefficients. We leave the details of these arguments as exercises.

Example 4L.4: Maps of $\mathbb{H}\mathrm{P}^\infty$. We can use the operations P^i together with a bit of number theory to demonstrate an interesting distinction between $\mathbb{H}\mathrm{P}^\infty$ and $\mathbb{C}\mathrm{P}^\infty$, namely, we will show that if a map $f:\mathbb{H}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ has $f^*(\gamma) = d\gamma$ for γ a generator of $H^4(\mathbb{H}\mathrm{P}^\infty;\mathbb{Z})$, then the integer d, which we call the *degree* of f, must be a square. By contrast, since $\mathbb{C}\mathrm{P}^\infty$ is a $K(\mathbb{Z},2)$, there are maps $\mathbb{C}\mathrm{P}^\infty \to \mathbb{C}\mathrm{P}^\infty$ carrying a generator $\alpha \in H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z})$ onto any given multiple of itself. Explicitly, the map $z \mapsto z^d$, $z \in \mathbb{C}$, induces a map f of $\mathbb{C}\mathrm{P}^\infty$ with $f^*(\alpha) = d\alpha$, but commutativity of $\mathbb{C}$ is needed for this construction so it does not extend to the quaternionic case.

We shall deduce the action of Steenrod powers on $H^*(\mathbb{H}\mathrm{P}^\infty;\mathbb{Z}_p)$ from their action on $H^*(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}_p)$, given by the earlier formula $(*)$ which says that $P^i(\alpha^n) = \binom{n}{i}\alpha^{n+i(p-1)}$ for α a generator of $H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}_p)$. There is a natural quotient map $\mathbb{C}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ arising from the definition of both spaces as quotients of S^∞. This map takes the 4-cell of $\mathbb{C}\mathrm{P}^\infty$ homeomorphically onto the 4-cell of $\mathbb{H}\mathrm{P}^\infty$, so the induced map on cohomology sends a generator $\gamma \in H^4(\mathbb{H}\mathrm{P}^\infty;\mathbb{Z}_p)$ to α^2, hence γ^n to α^{2n}. Thus the formula $P^i(\alpha^{2n}) = \binom{2n}{i}\alpha^{2n+i(p-1)}$ implies that $P^i(\gamma^n) = \binom{2n}{i}\gamma^{n+i(p-1)/2}$. For example, $P^1(\gamma) = 2\gamma^{(p+1)/2}$.

Now let $f:\mathbb{H}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ be any map. Applying the formula $P^1(\gamma) = 2\gamma^{(p+1)/2}$ in two ways, we get

$$P^1 f^*(\gamma) = f^* P^1(\gamma) = f^*(2\gamma^{(p+1)/2}) = 2d^{(p+1)/2}\gamma^{(p+1)/2}$$
$$\text{and}\quad P^1 f^*(\gamma) = P^1(d\gamma) = 2d\gamma^{(p+1)/2}$$

Hence the degree d satisfies $d^{(p+1)/2} \equiv d \bmod p$ for all odd primes p. Thus either $d \equiv 0 \bmod p$ or $d^{(p-1)/2} \equiv 1 \bmod p$. In both cases d is a square mod p since the congruence $d^{(p-1)/2} \equiv 1 \bmod p$ is equivalent to d being a nonzero square mod p, the multiplicative group of nonzero elements of the field $\mathbb{Z}_p$ being cyclic of order $p-1$.

The argument is completed by appealing to the number theory fact that an integer which is a square mod p for all sufficiently large primes p must be a square. This can be deduced from quadratic reciprocity and Dirichlet's theorem on primes in arithmetic progressions as follows. Suppose on the contrary that the result is false for the integer d. Consider primes p not dividing d. Since the product of two squares in $\mathbb{Z}_p$ is again a square, we may assume that d is a product of distinct primes $q_1, \cdots, q_n$, where one of these primes is allowed to be -1 if d is negative. In terms of the Legendre symbol $(\frac{d}{p})$ which is defined to be $+1$ if d is a square mod p and -1 otherwise, we have

$$\Big(\frac{d}{p}\Big) = \Big(\frac{q_1}{p}\Big)\cdots\Big(\frac{q_n}{p}\Big)$$

The left side is $+1$ for all large p by hypothesis, so it will suffice to see that p can be chosen to give each term on the right an arbitrary preassigned value. The values of $(\frac{-1}{p})$ and $(\frac{2}{p})$ depend only on p mod 8, and the four combinations of values are realized by the four residues $1,3,5,7$ mod 8. Having specified the value of p mod 8, the quadratic reciprocity law then says that for odd primes q, specifying $(\frac{q}{p})$ is equivalent to specifying $(\frac{p}{q})$. Thus we need only choose p in the appropriate residue classes mod 8 and mod q_i for each odd q_i. By the Chinese remainder theorem, this means specifying p modulo 8 times a product of odd primes. Dirichlet's theorem guarantees that in fact infinitely many primes p exist satisfying this congruence condition.

It is known that the integers realizable as degrees of maps $\mathbb{H}\mathrm{P}^\infty \to \mathbb{H}\mathrm{P}^\infty$ are exactly the odd squares and zero. The construction of maps of odd square degree will be given in [SSAT] using localization techniques, following [Sullivan 1974]. Ruling out nonzero even squares can be done using K–theory; see [Feder & Gitler 1978], which also treats maps $\mathbb{H}\mathrm{P}^n \to \mathbb{H}\mathrm{P}^n$.

The preceding calculations can also be used to show that every map $\mathbb{H}\mathrm{P}^n \to \mathbb{H}\mathrm{P}^n$ must have a fixed point if $n > 1$. For, taking $p = 3$, the element $P^1(\gamma)$ lies in $H^8(\mathbb{H}\mathrm{P}^n;\mathbb{Z}_3)$ which is nonzero if $n > 1$, so, when the earlier argument is specialized to the case $p = 3$, the congruence $d^{(p+1)/2} \equiv d \bmod p$ becomes $d^2 = d$ in $\mathbb{Z}_3$, which is satisfied only by 0 and 1 in $\mathbb{Z}_3$. In particular, d is not equal to -1. The Lefschetz number $\lambda(f) = 1 + d + \cdots + d^n = (d^{n+1} - 1)/(d-1)$ is therefore nonzero since the only integer roots of unity are ± 1. The Lefschetz fixed point theorem then gives the result.

Example 4L.5: Vector Fields on Spheres. Let us now apply Steenrod squares to determine the maximum number of orthonormal tangent vector fields on a sphere in all cases except when the dimension of the sphere is congruent to -1 mod 16. The first step is to rephrase the question in terms of Stiefel manifolds. Recall from the end of §3.D and Example 4.53 the space $V_{n,k}$ of orthonormal k-frames in $\mathbb{R}^n$. Projection of a k-frame onto its first vector gives a map $p: V_{n,k} \to S^{n-1}$, and a section for this projection, that is, a map $f: S^{n-1} \to V_{n,k}$ such that $pf = \mathbb{1}$, is exactly a set of $k-1$ orthonormal tangent vector fields $v_1, \cdots, v_{k-1}$ on S^{n-1} since f assigns to each $x \in S^{n-1}$ an orthonormal k-frame $(x, v_1(x), \cdots, v_{k-1}(x))$.

We described a cell structure on $V_{n,k}$ at the end of §3.D, and we claim that the $(n-1)$-skeleton of this cell structure is $\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1}$ if $2k-1 \le n$. The cells of $V_{n,k}$ were products $e^{i_1} \times \cdots \times e^{i_m}$ with $n > i_1 > \cdots > i_m \ge n-k$, so the products with a single factor account for all of the $(2n-2k)$-skeleton, hence they account for all of the $(n-1)$-skeleton if $n-1 \le 2n-2k$, that is, if $2k-1 \le n$. The cells that are products with a single factor are the homeomorphic images of cells of $\mathbb{R}\mathrm{P}^{n-1}$ under a map $\mathbb{R}\mathrm{P}^{n-1} \to SO(n) \to SO(n)/SO(n-k) = V_{n,k}$. This map collapses $\mathbb{R}\mathrm{P}^{n-k-1}$ to a point, so we get the desired conclusion that $\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1}$ is the $(n-1)$-skeleton of $V_{n,k}$ if $2k-1 \le n$.

Now suppose we have $f: S^{n-1} \to V_{n,k}$ with $pf = \mathbb{1}$. In particular, f^* is surjective on $H^{n-1}(-;\mathbb{Z}_2)$. If we deform f to a cellular map, with image in the $(n-1)$-skeleton, then by the preceding paragraph this will give a map $g: S^{n-1} \to \mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1}$ if $2k-1 \le n$, and this map will still induce a surjection on $H^{n-1}(-;\mathbb{Z}_2)$, hence an isomorphism. If the number k happens to be such that $\binom{n-k}{k-1} \equiv 1 \bmod 2$, then by the earlier formula $(*)$ the operation

$$Sq^{k-1}: H^{n-k}(\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1};\mathbb{Z}_2) \to H^{n-1}(\mathbb{R}\mathrm{P}^{n-1}/\mathbb{R}\mathrm{P}^{n-k-1};\mathbb{Z}_2)$$

will be nonzero, contradicting the existence of the map g since obviously the operation $Sq^{k-1}: H^{n-k}(S^{n-1};\mathbb{Z}_2) \to H^{n-1}(S^{n-1};\mathbb{Z}_2)$ is zero.

In order to guarantee that $\binom{n-k}{k-1} \equiv 1 \bmod 2$, write $n = 2^r(2s+1)$ and choose $k = 2^r + 1$. Assume for the moment that $s \ge 1$. Then $\binom{n-k}{k-1} = \binom{2^{r+1}s-1}{2^r}$, and in view of the rule for computing binomial coefficients in $\mathbb{Z}_2$, this is nonzero since the dyadic expansion of $2^{r+1}s-1$ ends with a string of 1's including a 1 in the single digit where the expansion of 2^r is nonzero. Note that the earlier condition $2k-1 \le n$ is satisfied since it becomes $2^{r+1}+1 \le 2^{r+1}s + 2^r$ and we assume $s \ge 1$.

Summarizing, we have shown that for $n = 2^r(2s+1)$, the sphere S^{n-1} cannot have 2^r orthonormal tangent vector fields if $s \ge 1$. This is also trivially true for $s = 0$ since S^{n-1} cannot have n orthonormal tangent vector fields.

It is easy to see that this result is best possible when $r \le 3$ by explicitly constructing $2^r - 1$ orthonormal tangent vector fields on S^{n-1} when $n = 2^r m$. When $r = 1$, view S^{n-1} as the unit sphere in $\mathbb{C}^m$, and then $x \mapsto ix$ defines a tangent vector field since the unit complex numbers 1 and i are orthogonal and multiplication by a unit complex number is an isometry of $\mathbb{C}$, so x and ix are orthogonal in each coordinate of $\mathbb{C}^m$, hence are orthogonal. When $r = 2$ the same construction works with $\mathbb{H}$ in place of $\mathbb{C}$, using the maps $x \mapsto ix$, $x \mapsto jx$, and $x \mapsto kx$ to define three orthonormal tangent vector fields on the unit sphere in $\mathbb{H}^m$. When $r = 3$ we can follow the same procedure with the octonions, constructing seven orthonormal tangent vector fields to the unit sphere in $\mathbb{O}^m$ via an orthonormal basis $1, i, j, k, \cdots$ for $\mathbb{O}$.

The upper bound of $2^r - 1$ for the number of orthonormal vector fields on S^{n-1} is not best possible in the remaining case $n \equiv 0 \bmod 16$. The optimal upper bound is obtained instead using K–theory; see [VBKT] or [Husemoller 1966]. The construction of the requisite number of vector fields is again algebraic, this time using Clifford algebras.

Example 4L.6: A Map of mod p Hopf Invariant One. Let us describe a construction for a map $f: S^{2p} \to S^3$ such that in the mapping cone $C_f = S^3 \cup_f e^{2p+1}$, the first Steenrod power $P^1: H^3(C_f; \mathbb{Z}_p) \to H^{2p+1}(C_f; \mathbb{Z}_p)$ is nonzero, hence f is nonzero in π^s_{2p-3}. The construction starts with the fact that a generator of $H^2(K(\mathbb{Z}_p, 1); \mathbb{Z}_p)$ has nontrivial p^{th} power, so the operation $P^1: H^2(K(\mathbb{Z}_p, 1); \mathbb{Z}_p) \to H^{2p}(K(\mathbb{Z}_p, 1); \mathbb{Z}_p)$ is nontrivial by property (5). This remains true after we suspend to $\Sigma K(\mathbb{Z}_p, 1)$, and we showed in Proposition 4I.3 that $\Sigma K(\mathbb{Z}_p, 1)$ has the homotopy type of a wedge sum of CW complexes X_i, $1 \le i \le p-1$, with $\widetilde{H}_*(X_i; \mathbb{Z})$ consisting only of a $\mathbb{Z}_p$ in each dimension congruent to $2i$ mod $2p-2$. We are interested here in the space $X = X_1$, which has nontrivial $\mathbb{Z}_p$ cohomology in dimensions $2, 3, 2p, 2p+1, \cdots$. Since X is, up to homotopy, a wedge summand of $\Sigma K(\mathbb{Z}_p, 1)$, the operation $P^1: H^3(X; \mathbb{Z}_p) \to H^{2p+1}(X; \mathbb{Z}_p)$ is nontrivial. Since X is simply-connected, the construction in §4.C shows that we may take X to have $(2p+1)$-skeleton of the form $S^2 \cup e^3 \cup e^{2p} \cup e^{2p+1}$. In fact, using the notion of homology decomposition in §4.H, we can take this skeleton to be the reduced mapping cone C_g of a map of Moore spaces $g: M(\mathbb{Z}_p, 2p-1) \to M(\mathbb{Z}_p, 2)$. It follows that the quotient C_g/S^2 is the reduced mapping cone of the composition $h: M(\mathbb{Z}_p, 2p-1) \xrightarrow{g} M(\mathbb{Z}_p, 2) \to M(\mathbb{Z}_p, 2)/S^2 = S^3$. The restriction $h|S^{2p-1}$ represents an element of $\pi_{2p-1}(S^3)$ that is either trivial or has order p since this restriction extends over the $2p$-cell of $M(\mathbb{Z}_p, 2p-1)$ which is attached by a map $S^{2p-1} \to S^{2p-1}$ of degree p. In fact, $h|S^{2p-1}$ is nullhomotopic since, as we will see in [SSAT] using the Serre spectral sequence, $\pi_i(S^3)$ contains no elements of order p if $i \le 2p-1$. This implies that the space $C_h = C_g/S^2$ is homotopy equivalent to a CW complex Y obtained from $S^3 \vee S^{2p}$ by attaching a cell e^{2p+1}. The quotient Y/S^{2p} then has the form $S^3 \cup e^{2p+1}$, so it is the mapping cone of a map $f: S^{2p} \to S^3$. By construction there is a map $C_g \to C_f$ inducing an isomorphism on $\mathbb{Z}_p$ cohomology in dimensions 3 and $2p+1$, so the operation P^1 is nontrivial in $H^*(C_f; \mathbb{Z}_p)$ since this was true for C_g, the $(2p+1)$-skeleton of X.

Example 4L.7: Moore Spaces. Let us use the operation Sq^2 to show that for $n \ge 2$, the identity map of $M(\mathbb{Z}_2, n)$ has order 4 in the group of basepoint-preserving homotopy classes of maps $M(\mathbb{Z}_2, n) \to M(\mathbb{Z}_2, n)$, with addition defined via the suspension structure on $M(\mathbb{Z}_2, n) = \Sigma M(\mathbb{Z}_2, n-1)$. According to Proposition 4H.2, this group is the middle term of a short exact sequence, the remaining terms of which contain only elements of order 2. Hence if the identity map of $M(\mathbb{Z}_2,n)$ has order 4, this short exact sequence cannot split.

In view of the short exact sequence just referred to, it will suffice to show that twice the identity map of $M(\mathbb{Z}_2,n)$ is not nullhomotopic. If twice the identity were nullhomotopic, then the mapping cone C of this map would have the homotopy type of $M(\mathbb{Z}_2,n)\vee\Sigma M(\mathbb{Z}_2,n)$. This would force $Sq^2\colon H^n(C;\mathbb{Z}_2)\to H^{n+2}(C;\mathbb{Z}_2)$ to be trivial since the source and target groups would come from different wedge summands. However, we will now show that this Sq^2 operation is nontrivial. Twice the identity map of $M(\mathbb{Z}_2,n)$ can be regarded as the smash product of the degree 2 map $S^1\to S^1$, $z\mapsto z^2$, with the identity map of $M(\mathbb{Z}_2,n-1)$. If we smash the cofibration sequence $S^1\to S^1\to\mathbb{R}P^2$ for this degree 2 map with $M(\mathbb{Z}_2,n-1)$ we get the cofiber sequence $M(\mathbb{Z}_2,n)\to M(\mathbb{Z}_2,n)\to C$, in view of the identity $(X/A)\wedge Y=(X\wedge Y)/(A\wedge Y)$. This means we can view C as $\mathbb{R}P^2\wedge M(\mathbb{Z}_2,n-1)$. The Cartan formula translated to cross products gives $Sq^2(\alpha\times\beta)=Sq^0\alpha\times Sq^2\beta+Sq^1\alpha\times Sq^1\beta+Sq^2\alpha\times Sq^0\beta$. This holds for smash products as well as ordinary products, by naturality. Taking α to be a generator of $H^1(\mathbb{R}P^2;\mathbb{Z}_2)$ and β a generator of $H^{n-1}(M(\mathbb{Z}_2,n-1);\mathbb{Z}_2)$, we have $Sq^2\alpha=0=Sq^2\beta$, but $Sq^1\alpha$ and $Sq^1\beta$ are nonzero since Sq^1 is the Bockstein. By the Künneth formula, $Sq^1\alpha\times Sq^1\beta$ then generates $H^{n+2}(\mathbb{R}P^2\wedge M(\mathbb{Z}_2,n-1);\mathbb{Z}_2)$ and we are done.

Adem Relations and the Steenrod Algebra

When Steenrod squares or powers are composed, the compositions satisfy certain relations, unfortunately rather complicated, known as **Adem relations**:

$$
\begin{aligned}
Sq^aSq^b &= \textstyle\sum_j \binom{b-j-1}{a-2j} Sq^{a+b-j}Sq^j \qquad \text{if } a<2b\\
P^aP^b &= \textstyle\sum_j (-1)^{a+j}\binom{(p-1)(b-j)-1}{a-pj} P^{a+b-j}P^j \qquad \text{if } a<pb\\
P^a\beta P^b &= \textstyle\sum_j (-1)^{a+j}\binom{(p-1)(b-j)}{a-pj}\beta P^{a+b-j}P^j\\
&\quad -\textstyle\sum_j (-1)^{a+j}\binom{(p-1)(b-j)-1}{a-pj-1}P^{a+b-j}\beta P^j \qquad \text{if } a\le pb
\end{aligned}
$$

By convention, the binomial coefficient $\binom{m}{n}$ is taken to be zero if m or n is negative or if $m<n$. Also $\binom{m}{0}=1$ for $m\ge 0$.

For example, taking $a=1$ in the Adem relation for the Steenrod squares we have $Sq^1Sq^b=(b-1)Sq^{b+1}$, so $Sq^1Sq^{2i}=Sq^{2i+1}$ and $Sq^1Sq^{2i+1}=0$. The relations $Sq^1Sq^{2i}=Sq^{2i+1}$ and $Sq^1=\beta$ explain the earlier comment that Sq^{2i} is the analog of P^i for $p=2$.

The **Steenrod algebra** $\mathcal{A}_2$ is defined to be the algebra over $\mathbb{Z}_2$ that is the quotient of the algebra of polynomials in the noncommuting variables $Sq^1,Sq^2,\cdots$ by the two-sided ideal generated by the Adem relations, that is, by the polynomials given by the differences between the left and right sides of the Adem relations. In similar fashion, $\mathcal{A}_p$ for odd p is defined to be the algebra over $\mathbb{Z}_p$ formed by polynomials in the noncommuting variables $\beta,P^1,P^2,\cdots$ modulo the Adem relations and the relation

$\beta^2 = 0$. Thus for every space X, $H^*(X;\mathbb{Z}_p)$ is a module over $\mathcal{A}_p$, for all primes p. The Steenrod algebra is a graded algebra, the elements of degree k being those that map $H^n(X;\mathbb{Z}_p)$ to $H^{n+k}(X;\mathbb{Z}_p)$ for all n.

The next proposition implies that $\mathcal{A}_2$ is generated as an algebra by the elements Sq^{2^k}, while $\mathcal{A}_p$ for p odd is generated by β and the elements P^{p^k}.

Proposition 4L.8. *There is a relation $Sq^i = \sum_{0<j<i} a_j Sq^{i-j}Sq^j$ with coefficients $a_j \in \mathbb{Z}_2$ whenever i is not a power of 2. Similarly, if i is not a power of p there is a relation $P^i = \sum_{0<j<i} a_j P^{i-j}P^j$ with $a_j \in \mathbb{Z}_p$.*

Proof: The argument is the same for $p = 2$ and p odd, so we describe the latter case. The idea is to write i as the sum $a + b$ of integers $a > 0$ and $b > 0$ with $a < pb$, such that the coefficient of the $j = 0$ term in the Adem relation for $P^a P^b$ is nonzero. Then one can solve this relation for $P^{a+b} = P^i$.

Let the p-adic representation of i be $i = i_0 + i_1 p + \cdots + i_k p^k$ with $i_k \neq 0$. Let $b = p^k$ and $a = i - p^k$, so $b > 0$ and $a > 0$ if i is not a power of p. The claim is that $\binom{(p-1)b-1}{a}$ is nonzero in $\mathbb{Z}_p$. The p-adic expansion of $(p-1)b - 1 = (p^{k+1} - 1) - p^k$ is $(p-1) + (p-1)p + \cdots + (p-2)p^k$, and the p-adic expansion of a is $i_0 + i_1 p + \cdots + (i_k - 1)p^k$. Hence $\binom{(p-1)b-1}{a} \equiv \binom{p-1}{i_0} \cdots \binom{p-2}{i_k-1}$ and in each factor of the latter product the numerator is nonzero in $\mathbb{Z}_p$ so the product is nonzero in $\mathbb{Z}_p$. When $p = 2$ the last factor is omitted, and the product is still nonzero in $\mathbb{Z}_2$. $\square$

This proposition says that most of the Sq^i's and P^i's are decomposable, where an element a of a graded algebra such as $\mathcal{A}_p$ is **decomposable** if it can be expressed in the form $\sum_i a_i b_i$ with each a_i and b_i having lower degree than a. The operation Sq^{2^k} is indecomposable since for α a generator of $H^1(\mathbb{R}\mathrm{P}^\infty;\mathbb{Z}_2)$ we saw that $Sq^{2^k}(\alpha^{2^k}) = \alpha^{2^{k+1}}$ but $Sq^i(\alpha^{2^k}) = 0$ for $0 < i < 2^k$. Similarly P^{p^k} is indecomposable since if $\alpha \in H^2(\mathbb{C}\mathrm{P}^\infty;\mathbb{Z}_p)$ is a generator then $P^{p^k}(\alpha^{p^k}) = \alpha^{p^{k+1}}$ but $P^i(\alpha^{p^k}) = 0$ for $0 < i < p^k$ and also $\beta(\alpha^{p^k}) = 0$.

Here is an application of the preceding proposition:

Theorem 4L.9. *Suppose $H^*(X;\mathbb{Z}_p)$ is the polynomial algebra $\mathbb{Z}_p[\alpha]$ on a generator α of dimension n, possibly truncated by the relation $\alpha^m = 0$ for $m > p$. Then if $p = 2$, n must be a power of 2, and if p is an odd prime, n must be of the form $p^k \ell$ where ℓ is an even divisor of $2(p-1)$.*

As we mentioned in §3.2, there is a stronger theorem that n must be 1, 2, 4, or 8 when $p = 2$, and n must be an even divisor of $2(p-1)$ when p is an odd prime. We also gave examples showing the necessity of the hypothesis $m > p$ in the case of a truncated polynomial algebra.

Proof: In the case $p = 2$, $Sq^n(\alpha) = \alpha^2 \neq 0$. If n is not a power of 2 then Sq^n decomposes into compositions $Sq^{n-j}Sq^j$ with $0 < j < n$. Such compositions must be zero since they pass through the group $H^{n+j}(X;\mathbb{Z}_2)$ which is zero for $0 < j < n$.

For odd p, the fact that α^2 is nonzero implies that n is even, say $n = 2k$. Then $P^k(\alpha) = \alpha^p \neq 0$. Since P^k can be expressed in terms of P^{p^i}'s, some P^{p^i} must be nonzero in $H^*(X;\mathbb{Z}_p)$. This implies that $2p^i(p-1)$, the amount by which P^{p^i} raises dimension, must be a multiple of n since $H^*(X;\mathbb{Z}_p)$ is concentrated in dimensions that are multiples of n. Since n divides $2p^i(p-1)$, it must be a power of p times a divisor of $2(p-1)$, and this divisor must be even since n is even and p is odd. □

Corollary 4L.10. *If $H^*(X;\mathbb{Z})$ is a polynomial algebra $\mathbb{Z}[\alpha]$, possibly truncated by $\alpha^m = 0$ with $m > 3$, then $|\alpha| = 2$ or 4.*

Proof: Passing from $\mathbb{Z}$ to $\mathbb{Z}_2$ coefficients, the theorem implies that $|\alpha|$ is a power of 2, and taking $\mathbb{Z}_3$ coefficients we see that $|\alpha|$ is a power of 3 times a divisor of $2(3-1) = 4$. □

In particular, the octonionic projective plane $\mathbb{O}\mathrm{P}^2$, constructed in Example 4.47 by attaching a 16-cell to S^8 via the Hopf map $S^{15} \to S^8$, does not generalize to an octonionic projective n-space $\mathbb{O}\mathrm{P}^n$ with $n \geq 3$.

In a similar vein, decomposability implies that if an element of π_*^s is detected by a Sq^i or P^i then i must be a power of 2 for Sq^i and a power of p for P^i. For if Sq^i is decomposable, then the map $Sq^i : H^n(C_f : \mathbb{Z}_2) \to H^{n+i}(C_f;\mathbb{Z}_2)$ must be trivial since it is a sum of compositions that pass through trivial cohomology groups, and similarly for P^i.

Interestingly enough, the Adem relations can also be used in a positive way to detect elements of π_*^s, as the proof of the following result will show.

Proposition 4L.11. *If $\eta \in \pi_1^s$ is represented by the Hopf map $S^3 \to S^2$, then η^2 is nonzero in π_2^s. Similarly, the other two Hopf maps represent elements $\nu \in \pi_3^s$ and $\sigma \in \pi_7^s$ whose squares are nontrivial in π_6^s and π_{14}^s.*

Proof: Let $\eta : S^{n+1} \to S^n$ be a suspension of the Hopf map, with mapping cone C_η obtained from S^n by attaching a cell e^{n+2} via η. If we assume the composition $(\Sigma\eta)\eta$ is nullhomotopic, then we can define a map $f : S^{n+3} \to C_\eta$ in the following way. Decompose S^{n+3} as the union of two cones CS^{n+2}. On one of these cones let f be a nullhomotopy of $(\Sigma\eta)\eta$. On the other cone let f be the composition $CS^{n+2} \to CS^{n+1} \to C_\eta$ where the first map is obtained by coning $\Sigma\eta$ and the second map is a characteristic map for the cell e^{n+2}.

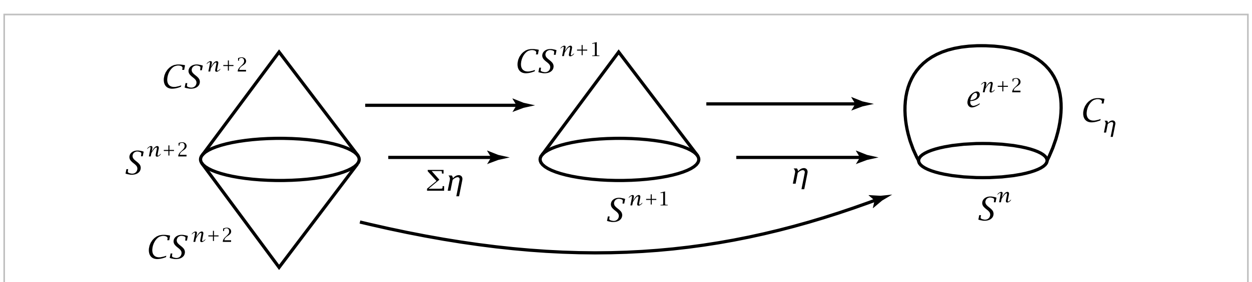

We use the map f to attach a cell e^{n+4} to C_η, forming a space X. This has C_η as its $(n+2)$-skeleton, so $Sq^2 : H^n(X;\mathbb{Z}_2) \to H^{n+2}(X;\mathbb{Z}_2)$ is an isomorphism. The map $Sq^2 : H^{n+2}(X;\mathbb{Z}_2) \to H^{n+4}(X;\mathbb{Z}_2)$ is also an isomorphism since the quotient map $X \to X/S^n$ induces an isomorphism on cohomology groups above dimension n and X/S^n is homotopy equivalent to the mapping cone of $\Sigma^2\eta$. Thus the composition $Sq^2Sq^2 : H^n(X;\mathbb{Z}_2) \to H^{n+4}(X;\mathbb{Z}_2)$ is an isomorphism. But this is impossible in view of the Adem relation $Sq^2Sq^2 = Sq^3Sq^1$, since Sq^1 is trivial on $H^n(X;\mathbb{Z}_2)$.

The same argument shows that ν^2 and σ^2 are nontrivial using the relations $Sq^4Sq^4 = Sq^7Sq^1 + Sq^6Sq^2$ and $Sq^8Sq^8 = Sq^{15}Sq^1 + Sq^{14}Sq^2 + Sq^{12}Sq^4$. ☐

This line of reasoning does not work for odd primes and the element $\alpha \in \pi^s_{2p-3}$ detected by P^1 since the Adem relation for P^1P^1 is $P^1P^1 = 2P^2$, which is not helpful. And in fact $\alpha^2 = 0$ by the commutativity property of the product in π^s_*.

When dealing with $\mathcal{A}_2$ it is often convenient to abbreviate notation by writing a monomial $Sq^{i_1}Sq^{i_2}\cdots$ as Sq^I where I is the finite sequence of nonnegative integers $i_1, i_2, \cdots$. Call Sq^I **admissible** if no Adem relation can be applied to it, that is, if $i_j \geq 2i_{j+1}$ for all j. The Adem relations imply that every monomial Sq^I can be written as a sum of admissible monomials. For if Sq^I is not admissible, it contains a pair Sq^aSq^b to which an Adem relation can be applied, yielding a sum of terms Sq^J for which $J > I$ with respect to the lexicographic ordering on finite sequences of integers. These Sq^J's have the same degree $i_1 + \cdots + i_k$ as Sq^I, and since the number of monomials Sq^I of a fixed degree is finite, successive applications of the Adem relations eventually reduce any Sq^I to a sum of admissible monomials.

For odd p, elements of $\mathcal{A}_p$ are linear combinations of monomials $\beta^{\varepsilon_1}P^{i_1}\beta^{\varepsilon_2}P^{i_2}\cdots$ with each $\varepsilon_j = 0$ or 1. Such a monomial is **admissible** if $i_j \geq \varepsilon_{j+1} + pi_{j+1}$ for all j, which again means that no Adem relation can be applied to the monomial. As with $\mathcal{A}_2$, the Adem relations suffice to reduce every monomial to a linear combination of admissible monomials, by the same argument as before but now using the lexicographic ordering on tuples $(\varepsilon_1 + pi_1, \varepsilon_2 + pi_2, \cdots)$.

Define the **excess** of the admissible monomial Sq^I to be $\sum_j (i_j - 2i_{j+1})$, the amount by which Sq^I exceeds being admissible. For odd p one might expect the excess of an admissible monomial $\beta^{\varepsilon_1}P^{i_1}\beta^{\varepsilon_2}P^{i_2}\cdots$ to be defined as $\sum_j (i_j - pi_{j+1} - \varepsilon_{j+1})$, but instead it is defined to be $\sum_j (2i_j - 2pi_{j+1} - \varepsilon_{j+1})$, for reasons which will become clear below.

As we explained at the beginning of this section, cohomology operations correspond to elements in the cohomology of Eilenberg–MacLane spaces. Here is a rather important theorem which will be proved in [SSAT] since the proof makes heavy use of spectral sequences:

Theorem. *For each prime* p, $H^*(K(\mathbb{Z}_p,n);\mathbb{Z}_p)$ *is the free commutative algebra on the generators* $\Theta(\iota_n)$ *where* $\iota_n \in H^n(K(\mathbb{Z}_p,n);\mathbb{Z}_p)$ *is a generator and* Θ *ranges over all admissible monomials of excess less than* n.

Here 'free commutative algebra' means 'polynomial algebra' when $p = 2$ and 'polynomial algebra on even-dimensional generators tensor exterior algebra on odd-dimensional generators' when p is odd. We will say something about the rationale behind the 'excess less than n' condition in a moment.

Specializing the theorem to the first two cases $n = 1,2$, we have the following cohomology algebras:

$$\begin{aligned}
K(\mathbb{Z}_2,1):&\quad \mathbb{Z}_2[\iota]\\
K(\mathbb{Z}_p,1):&\quad \Lambda_{\mathbb{Z}_p}[\iota]\otimes\mathbb{Z}_p[\beta\iota]\\
K(\mathbb{Z}_2,2):&\quad \mathbb{Z}_2[\iota, Sq^1\iota, Sq^2Sq^1\iota, Sq^4Sq^2Sq^1\iota,\cdots]\\
K(\mathbb{Z}_p,2):&\quad \mathbb{Z}_p[\iota, \beta P^1\beta\iota, \beta P^pP^1\beta\iota, \beta P^{p^2}P^pP^1\beta\iota,\cdots]\\
&\quad\quad \otimes\Lambda_{\mathbb{Z}_p}[\beta\iota, P^1\beta\iota, P^pP^1\beta\iota, P^{p^2}P^pP^1\beta\iota,\cdots]
\end{aligned}$$

The theorem implies that the admissible monomials in $\mathcal{A}_p$ are linearly independent, hence form a basis for $\mathcal{A}_p$ as a vector space over $\mathbb{Z}_p$. For if some linear combination of admissible monomials were zero, then it would be zero when applied to the class ι_n, but if we choose n larger than the excess of each monomial in the linear combination, this would contradict the freeness of the algebra $H^*(K(\mathbb{Z}_p,n);\mathbb{Z}_p)$. Even though the multiplicative structure of the Steenrod algebra is rather complicated, the Adem relations provide a way of performing calculations algorithmically by systematically reducing all products to sums of admissible monomials. A proof of the linear independence of admissible monomials using more elementary techniques can be found in [Steenrod & Epstein 1962].

Another consequence of the theorem is that all cohomology operations with $\mathbb{Z}_p$ coefficients are polynomials in the Sq^i's when $p = 2$ and polynomials in the P^i's and β when p is odd, in view of Proposition 4L.1. We can also conclude that $\mathcal{A}_p$ consists precisely of all the $\mathbb{Z}_p$ cohomology operations that are stable, commuting with suspension. For consider the map $\Sigma K(\mathbb{Z}_p,n)\to K(\mathbb{Z}_p,n+1)$ that pulls ι_{n+1} back to the suspension of ι_n. This map induces an isomorphism on homotopy groups π_i for $i \le 2n$ and a surjection for $i = 2n+1$ by Corollary 4.24, hence the same is true for homology and cohomology. Letting n go to infinity, the limit $\varprojlim \widetilde{H}^*(K(\mathbb{Z}_p,n);\mathbb{Z}_p)$ then exists in a strong sense. On the one hand, this limit is exactly the stable operations by Proposition 4L.1 and the definition of a stable operation. On the other hand, the preceding theorem implies that this limit is $\mathcal{A}_p$ since it says that all elements of $H^*(K(\mathbb{Z}_p,n);\mathbb{Z}_p)$ below dimension $2n$ are uniquely expressible as sums of admissible monomials applied to ι_n.

Now let us explain why the condition 'excess less than n' in the theorem is natural. For a monomial $Sq^I = Sq^{i_1}Sq^{i_2}\cdots$ the definition of the excess $e(I)$ can be rewritten as an equation $i_1 = e(I) + i_2 + i_3 + \cdots$. Thus if $e(I) > n$, we have $i_1 > |Sq^{i_2}Sq^{i_3}\cdots(\iota_n)|$, hence $Sq^I(\iota_n) = 0$. And if $e(I) = n$ then $Sq^I(\iota_n) = (Sq^{i_2}Sq^{i_3}\cdots(\iota_n))^2$ and either $Sq^{i_2}Sq^{i_3}\cdots$ has excess less than n or it has excess equal to n and we can repeat the process to write $Sq^{i_2}Sq^{i_3}\cdots(\iota_n) = (Sq^{i_3}\cdots(\iota_n))^2$, and so on, until we obtain an equation $Sq^I(\iota_n) = \big(Sq^J(\iota_n)\big)^{2^k}$ with $e(J) < n$, so that $Sq^I(\iota_n)$ is already in the algebra generated by the elements $Sq^J(\iota_n)$ with $e(J) < n$. The situation for odd p is similar. For an admissible monomial $P^I = \beta^{\varepsilon_1}P^{i_1}\beta^{\varepsilon_2}P^{i_2}\cdots$ the definition of excess gives $2i_1 = e(I) + \varepsilon_2 + 2(p-1)i_2 + \cdots$, so if $e(I) > n$ we must have $P^I(\iota_n) = 0$, and if $e(I) = n$ then either $P^I(\iota_n)$ is a power $(P^J(\iota_n))^{p^k}$ with $e(J) < n$, or, if P^I begins with β, then $P^I(\iota_n) = \beta((P^J(\iota_n))^{p^k}) = 0$ by the formula $\beta(x^m) = mx^{m-1}\beta(x)$, which is valid when $|x|$ is even, as we may assume is the case here, otherwise $(P^J(\iota_n))^{p^k} = 0$ by commutativity of cup product.

There is another set of relations among Steenrod squares equivalent to the Adem relations and somewhat easier to remember:

$$\textstyle\sum_j \binom{k}{j} Sq^{2n-k+j-1}Sq^{n-j} = 0$$

When $k = 0$ this is simply the relation $Sq^{2n-1}Sq^n = 0$, and the cases $k > 0$ are obtained from this via Pascal's triangle. For example, from $Sq^7Sq^4 = 0$ we obtain the following table of relations:

$$\begin{aligned}
Sq^7Sq^4 &= 0\\
Sq^6Sq^4 + Sq^7Sq^3 &= 0\\
Sq^5Sq^4 + Sq^7Sq^2 &= 0\\
Sq^4Sq^4 + Sq^5Sq^3 + Sq^6Sq^2 + Sq^7Sq^1 &= 0\\
Sq^3Sq^4 + Sq^7Sq^0 &= 0\\
Sq^2Sq^4 + Sq^3Sq^3 + Sq^6Sq^0 &= 0\\
Sq^1Sq^4 + Sq^3Sq^2 + Sq^5Sq^0 &= 0\\
Sq^0Sq^4 + Sq^1Sq^3 + Sq^2Sq^2 + Sq^3Sq^1 + Sq^4Sq^0 &= 0
\end{aligned}$$

These relations are not in simplest possible form. For example, $Sq^5Sq^3 = 0$ in the fourth row and $Sq^3Sq^2 = 0$ in the seventh row, instances of $Sq^{2n-1}Sq^n = 0$. For Steenrod powers there are similar relations $\sum_j \binom{k}{j} P^{pn-k+j-1}P^{n-j} = 0$ derived from the basic relation $P^{pn-1}P^n = 0$. We leave it to the interested reader to show that these relations follow from the Adem relations.

Constructing the Squares and Powers

Now we turn to the construction of the Steenrod squares and powers, and the proof of their basic properties including the Adem relations. As will be seen, this all hinges on the fact that cohomology is maps into Eilenberg–MacLane spaces. The case $p = 2$ is in some ways simpler than the case p odd, so in the first part of the development we will specialize p to 2 whenever there is a significant difference between the two cases.

Before giving the construction in detail, let us describe the idea in the case $p = 2$. The cup product square α^2 of an element $\alpha \in H^n(X;\mathbb{Z}_2)$ can be viewed as a composition $X \to X \times X \to K(\mathbb{Z}_2, 2n)$, with the first map the diagonal map and the second map representing the cross product $\alpha \times \alpha$. Since we have $\mathbb{Z}_2$ coefficients, cup product and cross product are strictly commutative, so if $T : X \times X \to X \times X$ is the map $T(x_1, x_2) = (x_2, x_1)$ transposing the two factors, then $T^*(\alpha \times \alpha) = \alpha \times \alpha$. Thinking of $\alpha \times \alpha$ as a map $X \times X \to K(\mathbb{Z}_2, 2n)$, this says there is a homotopy f_t from $\alpha \times \alpha$ to $(\alpha \times \alpha)T$. If we follow the homotopy f_t by the homotopy $f_t T$, we obtain a homotopy from $\alpha \times \alpha$ to $(\alpha \times \alpha)T$ and then to $(\alpha \times \alpha)T^2 = \alpha \times \alpha$, in other words a loop of maps $X \times X \to K(\mathbb{Z}_2, 2n)$. We can view this loop as a map $S^1 \times X \times X \to K(\mathbb{Z}_2, 2n)$. As we will see, if the homotopy f_t is chosen appropriately, the loop of maps will be nullhomotopic, extending to a map $D^2 \times X \times X \to K(\mathbb{Z}_2, 2n)$. Regarding D^2 as the upper hemisphere of S^2, this gives half of a map $S^2 \times X \times X \to K(\mathbb{Z}_2, 2n)$, and once again we obtain the other half by composition with T. This process can in fact be repeated infinitely often to yield a map $S^\infty \times X \times X \to K(\mathbb{Z}_2, 2n)$ with the property that each pair of points (s, x_1, x_2) and $(-s, x_2, x_1)$ is sent to the same point in $K(\mathbb{Z}_2, 2n)$. This means that when we compose with the diagonal map $S^\infty \times X \to S^\infty \times X \times X$, $(s, x) \mapsto (s, x, x)$, there is an induced quotient map $\mathbb{R}P^\infty \times X \to K(\mathbb{Z}_2, 2n)$ extending $\alpha^2 : X \to K(\mathbb{Z}_2, 2n)$. This extended map represents a class in $H^{2n}(\mathbb{R}P^\infty \times X;\mathbb{Z}_2)$. By the Künneth formula and the fact that $H^*(\mathbb{R}P^\infty;\mathbb{Z}_2)$ is the polynomial ring $\mathbb{Z}_2[\omega]$, this cohomology class in $H^{2n}(\mathbb{R}P^\infty \times X;\mathbb{Z}_2)$ can be written in the form $\sum_i \omega^{n-i} \times a_i$ with $a_i \in H^{n+i}(X;\mathbb{Z}_2)$. Then we define $Sq^i(\alpha) = a_i$.

The construction of the map $S^\infty \times X \times X \to K(\mathbb{Z}_2, 2n)$ will proceed cell by cell, so it will be convenient to eliminate any unnecessary cells. This is done by replacing $X \times X$ by the smash product $X \wedge X$ and factoring out a cross-sectional slice S^∞ in $S^\infty \times X \wedge X$. A further simplification will be to use naturality to reduce to the case $X = K(\mathbb{Z}_2, n)$.

Now we begin the actual construction. For a space X with basepoint x_0, let $X^{\wedge p}$ denote the smash product $X \wedge \cdots \wedge X$ of p copies of X. There is a map $T : X^{\wedge p} \to X^{\wedge p}$, $T(x_1, \cdots, x_p) = (x_2, \cdots, x_p, x_1)$, permuting the factors cyclically. Note that when $p = 2$ this is just the transposition $(x_1, x_2) \mapsto (x_2, x_1)$. The map T generates an action of $\mathbb{Z}_p$ on $X^{\wedge p}$. There is also the standard action of $\mathbb{Z}_p$ on S^∞ viewed as the union of the unit spheres S^{2n-1} in $\mathbb{C}^n$, a generator of $\mathbb{Z}_p$ rotating each $\mathbb{C}$ factor through an angle $2\pi/p$, with quotient space an infinite-dimensional lens space L^∞, or $\mathbb{R}P^\infty$ when $p = 2$. On the product $S^\infty \times X^{\wedge p}$ there is then the diagonal action $g(s, x) = (g(s), g(x))$ for $g \in \mathbb{Z}_p$. Let ΓX denote the orbit space $(S^\infty \times X^{\wedge p})/\mathbb{Z}_p$ of this diagonal action. This is the same as the Borel construction $S^\infty \times_{\mathbb{Z}_p} X^{\wedge p}$ described in §3.G. The projection $S^\infty \times X^{\wedge p} \to S^\infty$ induces a projection $\pi : \Gamma X \to L^\infty$ with $\pi^{-1}(z) = X^{\wedge p}$ for all $z \in L^\infty$ since the action of $\mathbb{Z}_p$ on S^∞ is free. This projection $\Gamma X \to L^\infty$ is in fact a fiber bundle, though we shall not need this fact and so we leave the proof as an exercise. The $\mathbb{Z}_p$ action on $X^{\wedge p}$ fixes the basepoint $x_0 \in X^{\wedge p}$, so the inclusion $S^\infty \times \{x_0\} \hookrightarrow S^\infty \times X^{\wedge p}$

induces an inclusion $L^\infty \hookrightarrow \Gamma X$. The composition $L^\infty \hookrightarrow \Gamma X \to L^\infty$ is the identity, so in fiber bundle terminology this subspace $L^\infty \subset \Gamma X$ is a section of the bundle. Let ΛX denote the quotient $\Gamma X/L^\infty$ obtained by collapsing the section L^∞ to a point. Note that the fibers $X^{\wedge p}$ in ΓX are still embedded in the quotient ΛX since each fiber meets the section L^∞ in a single point.

If we replace S^∞ by S^1 in these definitions, we get subspaces $\Gamma^1 X \subset \Gamma X$ and $\Lambda^1 X \subset \Lambda X$. All these spaces have natural CW structures if X is a CW complex having x_0 as a 0-cell. To see this, let L^∞ be given its standard CW structure with one cell in each dimension. This lifts to a CW structure on S^∞ with p cells in each dimension, and then T freely permutes the product cells of $S^\infty \times X^{\wedge p}$ so there is an induced quotient CW structure on ΓX. The section $L^\infty \subset \Gamma X$ is a subcomplex, so the quotient ΛX inherits a CW structure from ΓX. In particular, note that if the n-skeleton of X is S^n with its usual CW structure, then the pn-skeleton of ΛX is S^{pn} with its usual CW structure.

We remark also that Γ, Γ^1, Λ, and Λ^1 are functors: A map $f:(X,x_0)\to(Y,y_0)$ induces maps $\Gamma f:\Gamma X \to \Gamma Y$, etc., in the evident way.

For brevity we write $H^*(-;\mathbb{Z}_p)$ simply as $H^*(-)$. For $n>0$ let K_n denote a CW complex $K(\mathbb{Z}_p,n)$ with $(n-1)$-skeleton a point and n-skeleton S^n. Let $\iota \in H^n(K_n)$ be the canonical fundamental class described in the discussion following Theorem 4.57. It will be notationally convenient to regard an element $\alpha \in H^n(X)$ also as a map $\alpha: X \to K_n$ such that $\alpha^*(\iota) = \alpha$. Here we are assuming X is a CW complex.

From §3.2 we have a reduced p-fold cross product $\widetilde{H}^*(X)^{\otimes p} \to \widetilde{H}^*(X^{\wedge p})$ where $\widetilde{H}^*(X)^{\otimes p}$ denotes the p-fold tensor product of $\widetilde{H}^*(X)$ with itself. This cross product map $\widetilde{H}^*(X)^{\otimes p} \to \widetilde{H}^*(X^{\wedge p})$ is an isomorphism since we are using $\mathbb{Z}_p$ coefficients. With this isomorphism in mind, we will use the notation $\alpha_1 \otimes \cdots \otimes \alpha_p$ rather than $\alpha_1 \times \cdots \times \alpha_p$ for p-fold cross products in $\widetilde{H}^*(X^{\wedge p})$. In particular, for each element $\alpha \in H^n(X)$, $n>0$, we have its p-fold cross product $\alpha^{\otimes p} \in \widetilde{H}^{pn}(X^{\wedge p})$. Our first task will be to construct an element $\lambda(\alpha) \in H^{pn}(\Lambda X)$ restricting to $\alpha^{\otimes p}$ in each fiber $X^{\wedge p} \subset \Lambda X$. By naturality it will suffice to construct $\lambda(\iota) \in H^{pn}(\Lambda K_n)$.

The key point in the construction of $\lambda(\iota)$ is the fact that $T^*(\iota^{\otimes p}) = \iota^{\otimes p}$. In terms of maps $K_n^{\wedge p} \to K_{pn}$, this says the composition $\iota^{\otimes p}\, T$ is homotopic to $\iota^{\otimes p}$, preserving basepoints. Such a homotopy can be constructed as follows. The pn-skeleton of $K_n^{\wedge p}$ is $(S^n)^{\wedge p} = S^{pn}$, with T permuting the factors cyclically. Thinking of S^n as $(S^1)^{\wedge n}$, the permutation T is a product of $(p-1)n^2$ transpositions of adjacent factors, so T has degree $(-1)^{(p-1)n^2}$ on S^{pn}. If p is odd, this degree is $+1$, so the restriction of T to this skeleton is homotopic to the identity, hence $\iota^{\otimes p}\, T$ is homotopic to $\iota^{\otimes p}$ on this skeleton. This conclusion also holds when $p=2$, signs being irrelevant in this case since we are dealing with maps $S^{2n} \to K_{2n}$ and $\pi_{2n}(K_{2n}) = \mathbb{Z}_2$. Having a homotopy $\iota^{\otimes p}\, T \simeq \iota^{\otimes p}$ on the pn-skeleton, there are no obstructions to extending the homotopy over all higher-dimensional cells $e^i \times (0,1)$ since $\pi_i(K_{pn}) = 0$ for $i > pn$.

The homotopy $\iota^{\otimes p}T \simeq \iota^{\otimes p}: K_n^{\wedge p} \to K_{pn}$ defines a map $\Gamma^1 K_n \to K_{pn}$ since $\Gamma^1 X$ is the quotient of $I \times X^{\wedge p}$ under the identifications $(0,x) \sim (1,T(x))$. The homotopy is basepoint-preserving, so the map $\Gamma^1 K_n \to K_{pn}$ passes down to a quotient map $\lambda_1 : \Lambda^1 K_n \to K_{pn}$. Since K_n is obtained from S^n by attaching cells of dimension greater than n, ΛK_n is obtained from $\Lambda^1 K_n$ by attaching cells of dimension greater than $pn+1$. There are then no obstructions to extending λ_1 to a map $\lambda : \Lambda K_n \to K_{pn}$ since $\pi_i(K_{pn}) = 0$ for $i > pn$.

The map λ gives the desired element $\lambda(\iota) \in H^{pn}(\Lambda K_n)$ since the restriction of λ to each fiber $K_n^{\wedge p}$ is homotopic to $\iota^{\otimes p}$. Note that this property determines λ uniquely up to homotopy since the restriction map $H^{pn}(\Lambda K_n) \to H^{pn}(K_n^{\wedge p})$ is injective, the pn-skeleton of ΛK_n being contained in $K_n^{\wedge p}$. We shall have occasion to use this argument again in the proof, so we refer to it as 'the uniqueness argument'.

For any $\alpha \in H^n(X)$ let $\lambda(\alpha)$ be the composition $\Lambda X \xrightarrow{\Lambda\alpha} \Lambda K_n \xrightarrow{\lambda} K_{pn}$. This restricts to $\alpha^{\otimes p}$ in each fiber $X^{\wedge p}$ since $\Lambda\alpha$ restricts to $\alpha^{\otimes p}$ in each fiber.

Now we are ready to define some cohomology operations. There is an inclusion $L^\infty \times X \hookrightarrow \Gamma X$ as the quotient of the diagonal embedding $S^\infty \times X \hookrightarrow S^\infty \times X^{\wedge p}$, $(s,x) \mapsto (s,x,\cdots,x)$. Composing with the quotient map $\Gamma X \to \Lambda X$, we get a map $\nabla : L^\infty \times X \to \Lambda X$ inducing $\nabla^* : H^*(\Lambda X) \to H^*(L^\infty \times X) \approx H^*(L^\infty) \otimes H^*(X)$. For each $\alpha \in H^n(X)$ the element $\nabla^*(\lambda(\alpha)) \in H^{pn}(L^\infty \times X)$ may be written in the form

$$\nabla^*(\lambda(\alpha)) = \sum_i \omega_{(p-1)n-i} \otimes \theta_i(\alpha)$$

where ω_j is a generator of $H^j(L^\infty)$ and $\theta_i(\alpha) \in H^{n+i}(X)$. Thus θ_i increases dimension by i. When $p = 2$ there is no ambiguity about ω_j. For odd p we choose ω_1 to be the class dual to the 1-cell of L^∞ in its standard cell structure, then we take ω_2 to be the Bockstein $\beta\omega_1$ and we set $\omega_{2j} = \omega_2^j$ and $\omega_{2j+1} = \omega_1\omega_2^j$.

It is clear that θ_i is a cohomology operation since $\theta_i(\alpha) = \alpha^*(\theta_i(\iota))$. Note that $\theta_i = 0$ for $i < 0$ since $H^{n+i}(K_n) = 0$ for $i < 0$ except for $i = -n$, and in this special case $\theta_i = 0$ since $\nabla : L^\infty \times X \to \Lambda X$ sends $L^\infty \times \{x_0\}$ to a point.

For $p = 2$ we set $Sq^i(\alpha) = \theta_i(\alpha)$. For odd p we will show that $\theta_i = 0$ unless $i = 2k(p-1)$ or $2k(p-1)+1$. The operation P^k will be defined to be a certain constant times $\theta_{2k(p-1)}$, and $\theta_{2k(p-1)+1}$ will be a constant times βP^k, for β the mod p Bockstein.

Theorem 4L.12. *The operations Sq^i satisfy the properties* (1)–(7).

Proof: We have already observed that the θ_i's are cohomology operations, so property (1) holds. The basic property that $\lambda(\alpha)$ restricts to $\alpha^{\otimes p}$ in each fiber implies that $\theta_{(p-1)n}(\alpha) = \alpha^p$ since $\omega_0 = 1$. This gives the first half of property (5) for Sq^i. The second half follows from the fact that $\theta_i = 0$ for $i > (p-1)n$ since the factor $\omega_{(p-1)n-i}$ vanishes in this case.

Next we turn to the Cartan formula. For any prime p we will show that $\lambda(\alpha \smile \beta) = (-1)^{p(p-1)mn/2}\lambda(\alpha) \smile \lambda(\beta)$ for $m = |\alpha|$ and $n = |\beta|$. This implies (3) when $p = 2$

since if we let $\omega = \omega_1$, hence $\omega_j = \omega^j$, then

$$\begin{aligned}
\textstyle\sum_i Sq^i(\alpha \smile \beta) \otimes \omega^{n+m-i} &= \nabla^*(\lambda(\alpha \smile \beta)) = \nabla^*(\lambda(\alpha) \smile \lambda(\beta)) \\
&= \nabla^*(\lambda(\alpha)) \smile \nabla^*(\lambda(\beta)) \\
&= \textstyle\sum_j Sq^j(\alpha) \otimes \omega^{n-j} \smile \sum_k Sq^k(\beta) \otimes \omega^{m-k} \\
&= \textstyle\sum_i \Big(\sum_{j+k=i} Sq^j(\alpha) \smile Sq^k(\beta)\Big) \otimes \omega^{n+m-i}
\end{aligned}$$

To show that $\lambda(\alpha \smile \beta) = (-1)^{p(p-1)mn/2}\lambda(\alpha) \smile \lambda(\beta)$ we use the following diagram:

$$\begin{array}{ccccccc}
\Lambda X & \xrightarrow{\Lambda(\Delta)} & \Lambda(X \wedge X) & \xrightarrow{\Lambda(\alpha\wedge\beta)} & \Lambda(K_m \wedge K_n) & \xrightarrow{\Lambda(\iota_m \otimes \iota_n)} & \Lambda K_{m+n} \\
& \searrow^{\Delta} & \downarrow & & \downarrow & \searrow^{\lambda(\iota_m \otimes \iota_n)} & \downarrow \lambda \\
& & \Lambda X \wedge \Lambda X & \xrightarrow{\Lambda\alpha \wedge \Lambda\beta} & \Lambda K_m \wedge \Lambda K_n & \xrightarrow{\lambda(\iota_m)\otimes\lambda(\iota_n)} & K_{pm+pn}
\end{array}$$

Here Δ is a generic symbol for diagonal maps $x \mapsto (x,x)$. These relate cross product to cup product via $\Delta^*(\varphi \otimes \psi) = \varphi \smile \psi$. The two unlabeled vertical maps are induced by $(s, x_1, y_1, \cdots, x_p, y_p) \mapsto (s, x_1, \cdots, x_p, s, y_1, \cdots, y_p)$. The composition $\Lambda X \to K_{pm+pn}$ going across the top of the diagram is $\lambda(\alpha \smile \beta)$ since the composition $\Lambda X \to \Lambda K_{m+n}$ is $\Lambda(\alpha \smile \beta)$. The composition $\Lambda X \wedge \Lambda X \to K_{pm+pn}$ is $\lambda(\alpha) \otimes \lambda(\beta)$ so the composition $\Lambda X \to K_{pm+pn}$ across the bottom of the diagram is $\lambda(\alpha) \smile \lambda(\beta)$. The triangle on the left, the square, and the upper triangle on the right obviously commute from the definitions. It remains to see that the third triangle commutes up to the sign $(-1)^{p(p-1)mn/2}$. Since $(K_m \wedge K_n)^{\wedge p}$ includes the $(pm+pn)$-skeleton of $\Lambda(K_m \wedge K_n)$, restriction to this fiber is injective on H^{pm+pn}. On this fiber the two routes around the triangle give $(\iota_m \otimes \iota_n)^{\otimes p}$ and $\iota_m^{\otimes p} \otimes \iota_n^{\otimes p}$. These differ by a permutation that is the product of $(p-1) + (p-2) + \cdots + 1 = p(p-1)/2$ transpositions of adjacent factors. Since ι_m and ι_n have dimensions m and n, this permutation introduces a sign $(-1)^{p(p-1)mn/2}$ by the commutativity property of cup product. This finishes the verification of the Cartan formula when $p = 2$.

Before proceeding further we need to make an explicit calculation to show that Sq^0 is the identity on $H^1(S^1)$. Viewing S^1 as the one-point compactification of $\mathbb{R}$, with the point at infinity as the basepoint, the 2-sphere $S^1 \wedge S^1$ becomes the one-point compactification of $\mathbb{R}^2$. The map $T : S^1 \wedge S^1 \to S^1 \wedge S^1$ then corresponds to reflecting $\mathbb{R}^2$ across the line $x = y$, so after a rotation of coordinates this becomes reflection of S^2 across the equator. Hence $\Gamma^1 S^1$ is obtained from the shell $I \times S^2$ by identifying its inner and outer boundary spheres via a reflection across the equator. The diagonal $\mathbb{R}P^1 \times S^1 \subset \Gamma^1 S^1$ is a torus, obtained from the equatorial annulus $I \times S^1 \subset I \times S^2$ by identifying the two ends via the identity map since the equator is fixed by the reflection. This $\mathbb{R}P^1 \times S^1$ represents the same element of $H_2(\Gamma^1 S^1; \mathbb{Z}_2)$ as the fiber sphere $S^1 \wedge S^1$ since the upper half of the shell is a 3-cell whose mod 2 boundary in $\Gamma^1 S^1$ is the union of these two surfaces.

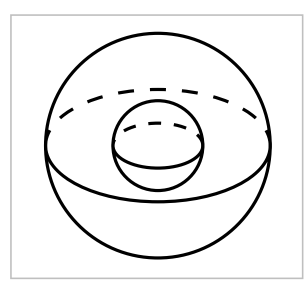

For a generator $\alpha \in H^1(S^1)$, consider the element $\nabla^*(\lambda(\alpha))$ in $H^2(\mathbb{R}\mathrm{P}^\infty \times S^1) \approx \mathrm{Hom}(H_2(\mathbb{R}\mathrm{P}^\infty \times S^1;\mathbb{Z}_2),\mathbb{Z}_2)$. A basis for $H_2(\mathbb{R}\mathrm{P}^\infty \times S^1;\mathbb{Z}_2)$ is represented by $\mathbb{R}\mathrm{P}^2 \times \{x_0\}$ and $\mathbb{R}\mathrm{P}^1 \times S^1$. A cocycle representing $\nabla^*(\lambda(\alpha))$ takes the value 0 on $\mathbb{R}\mathrm{P}^2 \times \{x_0\}$ since $\mathbb{R}\mathrm{P}^\infty \times \{x_0\}$ collapses to a point in ΛS^1 and $\lambda(\alpha)$ lies in $H^2(\Lambda S^1)$. On $\mathbb{R}\mathrm{P}^1 \times S^1$, $\nabla^*(\lambda(\alpha))$ takes the value 1 since when $\lambda(\alpha)$ is pulled back to ΓS^1 it takes the same value on the homologous cycles $\mathbb{R}\mathrm{P}^1 \times S^1$ and $S^1 \wedge S^1$, namely 1 by the defining property of $\lambda(\alpha)$ since $\alpha \otimes \alpha \in H^2(S^1 \wedge S^1)$ is a generator. Thus $\nabla^*(\lambda(\alpha)) = \omega_1 \otimes \alpha$ and hence $Sq^0(\alpha) = \alpha$ by the definition of Sq^0.

We use this calculation to prove that Sq^i commutes with the suspension σ, where σ is defined by $\sigma(\alpha) = \varepsilon \otimes \alpha \in H^*(S^1 \wedge X)$ for ε a generator of $H^1(S^1)$ and $\alpha \in H^*(X)$. We have just seen that $Sq^0(\varepsilon) = \varepsilon$. By (5), $Sq^1(\varepsilon) = \varepsilon^2 = 0$ and $Sq^i(\varepsilon) = 0$ for $i > 1$. The Cartan formula then gives $Sq^i(\sigma(\alpha)) = Sq^i(\varepsilon \otimes \alpha) = \sum_j Sq^j(\varepsilon) \otimes Sq^{i-j}(\alpha) = \varepsilon \otimes Sq^i(\alpha) = \sigma(Sq^i(\alpha))$.

From this it follows that Sq^0 is the identity on $H^n(S^n)$ for all $n > 0$. Since S^n is the n-skeleton of K_n, this implies that Sq^0 is the identity on the fundamental class ι_n, hence Sq^0 is the identity on all positive-dimensional classes.

Property (7) is proved similarly: Sq^1 coincides with the Bockstein β on the generator $\omega \in H^1(\mathbb{R}\mathrm{P}^2)$ since both equal ω^2. Hence $Sq^1 = \beta$ on the iterated suspensions of ω, and the n-fold suspension of $\mathbb{R}\mathrm{P}^2$ is the $(n+2)$-skeleton of K_{n+1}.

Finally we have the additivity property (2). This holds in fact for any cohomology operation that commutes with suspension. For such operations, it suffices to prove additivity in spaces that are suspensions. Consider a composition

$$\Sigma X \xrightarrow{c} \Sigma X \vee \Sigma X \xrightarrow{\alpha \vee \beta} K_n \xrightarrow{\theta} K_m$$

where c is the map that collapses an equatorial copy of X in ΣX to a point. The composition of the first two maps is $\alpha + \beta$, as in Lemma 4.60. Composing with the third map then gives $\theta(\alpha + \beta)$. On the other hand, if we first compose the second and third maps we get $\theta(\alpha) \vee \theta(\beta)$, and then composing with the first map gives $\theta(\alpha) + \theta(\beta)$. The two ways of composing are equal, so $\theta(\alpha + \beta) = \theta(\alpha) + \theta(\beta)$. □

Theorem 4L.13. *The Adem relations hold for Steenrod squares.*

Proof: The idea is to imitate the construction of ΛX using $\mathbb{Z}_p \times \mathbb{Z}_p$ in place of $\mathbb{Z}_p$. The Adem relations will come from the symmetry of $\mathbb{Z}_p \times \mathbb{Z}_p$ interchanging the factors.

The group $\mathbb{Z}_p \times \mathbb{Z}_p$ acts on $S^\infty \times S^\infty$ via $(g,h)(s,t) = (g(s),h(t))$, with quotient $L^\infty \times L^\infty$. There is also an action of $\mathbb{Z}_p \times \mathbb{Z}_p$ on $X^{\wedge p^2}$, obtained by writing points of $X^{\wedge p^2}$ as p^2-tuples (x_{ij}) with subscripts i and j varying from 1 to p, and then letting the first $\mathbb{Z}_p$ act on the first subscript and the second $\mathbb{Z}_p$ act on the second. Factoring out the diagonal action of $\mathbb{Z}_p \times \mathbb{Z}_p$ on $S^\infty \times S^\infty \times X^{\wedge p^2}$ gives a quotient space $\Gamma_2 X$. This projects to $L^\infty \times L^\infty$ with a section, and collapsing the section gives $\Lambda_2 X$. The fibers of the projection $\Lambda_2 X \to L^\infty \times L^\infty$ are $X^{\wedge p^2}$ since the action of $\mathbb{Z}_p \times \mathbb{Z}_p$ on $S^\infty \times S^\infty$ is free. We could also obtain $\Lambda_2 X$ from the product $S^\infty \times S^\infty \times X^{p^2}$ by first collapsing the subspace of points having at least one X coordinate equal to the basepoint x_0 and then factoring out the $\mathbb{Z}_p \times \mathbb{Z}_p$ action.

It will be useful to compare $\Lambda_2 X$ with $\Lambda(\Lambda X)$. The latter space is the quotient of $S^\infty \times (S^\infty \times X^p)^p$ in which one first identifies all points having at least one X coordinate equal to x_0 and then one factors out by an action of the wreath product $\mathbb{Z}_p \wr \mathbb{Z}_p$, the group of order p^{p+1} defined by a split exact sequence $0 \to \mathbb{Z}_p^p \to \mathbb{Z}_p \wr \mathbb{Z}_p \to \mathbb{Z}_p \to 0$ with conjugation by the quotient group $\mathbb{Z}_p$ given by cyclic permutations of the p $\mathbb{Z}_p$ factors of $\mathbb{Z}_p^p$. In the coordinates $(s, t_1, x_{11}, \cdots, x_{1p}, \cdots, t_p, x_{p1}, \cdots, x_{pp})$ the i^{th} factor $\mathbb{Z}_p$ of $\mathbb{Z}_p^p$ acts in the block $(t_i, x_{i1}, \cdots, x_{ip})$, and the quotient $\mathbb{Z}_p$ acts by cyclic permutation of the index i and by rotation in the s coordinate. There is a natural map $\Lambda_2 X \to \Lambda(\Lambda X)$ induced by $(s, t, x_{11}, \cdots, x_{1p}, \cdots, x_{p1}, \cdots, x_{pp}) \mapsto (s, t, x_{11}, \cdots, x_{1p}, \cdots, t, x_{p1}, \cdots, x_{pp})$. In $\Lambda_2 X$ one is factoring out by the action of $\mathbb{Z}_p \times \mathbb{Z}_p$. This is the subgroup of $\mathbb{Z}_p \wr \mathbb{Z}_p$ obtained by restricting the action of the quotient $\mathbb{Z}_p$ on $\mathbb{Z}_p^p$ to the diagonal subgroup $\mathbb{Z}_p \subset \mathbb{Z}_p^p$, where this action becomes trivial so that one has the direct product $\mathbb{Z}_p \times \mathbb{Z}_p$.

Since it suffices to prove that the Adem relations hold on the class $\iota \in H^n(K_n)$, we take $X = K_n$. There is a map $\lambda_2 : \Lambda_2 K_n \to K_{p^2 n}$ restricting to $\iota^{\otimes p^2}$ in each fiber. This is constructed by the same method used to construct λ. One starts with a map representing $\iota^{\otimes p^2}$ in a fiber, then extends this over the part of $\Lambda_2 K_n$ projecting to the 1-skeleton of $L^\infty \times L^\infty$, and finally one extends inductively over higher-dimensional cells of $\Lambda_2 K_n$ using the fact that $K_{p^2 n}$ is an Eilenberg–MacLane space. The map λ_2 fits into the diagram at the right, where ∇_2 is induced by the map $(s, t, x) \mapsto (s, t, x, \cdots, x)$ and the unlabeled map is the one defined above. It is clear that the square commutes.

$$\begin{array}{ccccc} L^\infty \times L^\infty \times K_n & \xrightarrow{\nabla_2} & \Lambda_2 K_n & \xrightarrow{\lambda_2} & K_{p^2 n} \\ \downarrow {\scriptstyle 1\!\!1 \times \nabla} & & \downarrow & \nearrow {\scriptstyle \lambda(\lambda)} & \\ L^\infty \times \Lambda K_n & \xrightarrow{\nabla} & \Lambda(\Lambda K_n) & & \end{array}$$

Commutativity of the triangle up to homotopy follows from the fact that λ_2 is uniquely determined, up to homotopy, by its restrictions to fibers.

The element $\nabla_2^* \lambda_2^*(\iota)$ may be written in the form $\sum_{r,s} \omega_r \otimes \omega_s \otimes \varphi_{rs}$, and we claim that the elements φ_{rs} satisfy the symmetry relation $\varphi_{rs} = (-1)^{rs + p(p-1)n/2} \varphi_{sr}$. To verify this we use the commutative diagram at the right where the map τ on the left switches the two L^∞ factors and the τ on the right is induced by switching the two S^∞ factors of $S^\infty \times S^\infty \times K_n^{\wedge p^2}$ and permuting the K_n factors of the smash product by interchanging the two subscripts in p^2-tuples (x_{ij}).

$$\begin{array}{ccc} L^\infty \times L^\infty \times K_n & \xrightarrow{\nabla_2} & \Lambda_2 K_n \\ \downarrow \tau & & \downarrow \tau \\ L^\infty \times L^\infty \times K_n & \xrightarrow{\nabla_2} & \Lambda_2 K_n \end{array}$$

This permutation is a product of $p(p-1)/2$ transpositions, one for each pair (i, j) with $1 \le i < j \le p$, so in a fiber the class $\iota^{\otimes p^2}$ is sent to $(-1)^{p(p-1)n/2} \iota^{\otimes p^2}$. By the uniqueness property of λ_2 this means that $\tau^* \lambda_2^*(\iota) = (-1)^{p(p-1)n/2} \lambda_2^*(\iota)$. Commutativity of the square then gives

$$(-1)^{p(p-1)n/2} \nabla_2^* \lambda_2^*(\iota) = \nabla_2^* \tau^* \lambda_2^*(\iota) = \tau^* \nabla_2^* \lambda_2^*(\iota) = \textstyle\sum_{r,s} (-1)^{rs} \omega_s \otimes \omega_r \otimes \varphi_{rs}$$

where the last equality follows from the commutativity property of cross products. The symmetry relation $\varphi_{rs} = (-1)^{rs+p(p-1)n/2}\varphi_{sr}$ follows by interchanging the indices r and s in the last summation.

If we compute $\nabla_2^*\lambda_2^*(\iota)$ using the lower route across the earlier diagram containing the map λ_2, we obtain

$$\begin{aligned}\nabla_2^*\lambda_2^*(\iota) &= \sum\nolimits_i \omega_{(p-1)pn-i}\otimes\theta_i\Big(\sum\nolimits_j \omega_{(p-1)n-j}\otimes\theta_j(\iota)\Big)\\ &= \sum\nolimits_{i,j}\omega_{(p-1)pn-i}\otimes\theta_i(\omega_{(p-1)n-j}\otimes\theta_j(\iota))\end{aligned}$$

Now we specialize to $p = 2$, so $\theta_i = Sq^i$ for all i. The Cartan formula converts the last summation above into $\sum_{i,j,k}\omega^{2n-i}\otimes Sq^k(\omega^{n-j})\otimes Sq^{i-k}Sq^j(\iota)$. Plugging in the value for $Sq^k(\omega^{n-j})$ computed in the discussion preceding Example 4L.3, we obtain $\sum_{i,j,k}\binom{n-j}{k}\omega^{2n-i}\otimes\omega^{n-j+k}\otimes Sq^{i-k}Sq^j(\iota)$. To write this summation more symmetrically with respect to the two ω terms, let $n-j+k = 2n-\ell$. Then we get

$$\sum\nolimits_{i,j,\ell}\binom{n-j}{n+j-\ell}\omega^{2n-i}\otimes\omega^{2n-\ell}\otimes Sq^{i+\ell-n-j}Sq^j(\iota)$$

In view of the symmetry property of φ_{rs}, which becomes $\varphi_{rs} = \varphi_{sr}$ for $p = 2$, switching i and ℓ in this formula leaves it unchanged. Hence we get the relation

$$(*)\qquad \sum\nolimits_j\binom{n-j}{n+j-\ell}Sq^{i+\ell-n-j}Sq^j(\iota) = \sum\nolimits_j\binom{n-j}{n+j-i}Sq^{i+\ell-n-j}Sq^j(\iota)$$

This holds for all n, i, and ℓ, and the idea is to choose these numbers so that the left side of this equation has only one nonzero term. Given integers r and s, let $n = 2^r-1+s$ and $\ell = n+s$, so that $\binom{n-j}{n+j-\ell} = \binom{2^r-1-(j-s)}{j-s}$. If r is sufficiently large, this will be 0 unless $j = s$. This is because the dyadic expansion of 2^r-1 consists entirely of 1's, so the expansion of $2^r-1-(j-s)$ will have 0's in the positions where the expansion of $j-s$ has 1's, hence these positions contribute factors of $\binom{0}{1} = 0$ to $\binom{2^r-1-(j-s)}{j-s}$. Thus with n and ℓ chosen as above, the relation $(*)$ becomes

$$Sq^iSq^s(\iota) = \sum\nolimits_j\binom{2^r-1+s-j}{2^r-1+s+j-i}Sq^{i+s-j}Sq^j(\iota) = \sum\nolimits_j\binom{2^r+s-j-1}{i-2j}Sq^{i+s-j}Sq^j(\iota)$$

where the latter equality comes from the general relation $\binom{x}{y} = \binom{x}{x-y}$.

The final step is to show that $\binom{2^r+s-j-1}{i-2j} = \binom{s-j-1}{i-2j}$ if $i < 2s$. Both of these binomial coefficients are zero if $i < 2j$. If $i \ge 2j$ then we have $2j \le i < 2s$, so $j < s$, hence $s-j-1 \ge 0$. The term 2^r then makes no difference in $\binom{2^r+s-j-1}{i-2j}$ if r is large since this 2^r contributes only a single 1 to the dyadic expansion of $2^r+s-j-1$, far to the left of all the nonzero entries in the dyadic expansions of $s-j-1$ and $i-2j$.

This gives the Adem relations for the classes ι of dimension $n = 2^r-1+s$ with r large. This implies the relations hold for all classes of these dimensions, by naturality. Since we can suspend repeatedly to make any class have dimension of this form, the Adem relations must hold for all cohomology classes. ◻

Steenrod Powers

Our remaining task is to verify the axioms and Adem relations for the Steenrod powers for an odd prime p. Unfortunately this is quite a bit more complicated than the $p = 2$ case, largely because one has to be very careful in computing the many coefficients in $\mathbb{Z}_p$ that arise. Even for the innocent-looking axiom $P^0 = 1\!\!1$ it will take three pages to calculate the normalization constants needed to make the axiom hold. One could wish that the whole process was a lot cleaner.

Lemma 4L.14. *$\theta_i = 0$ unless $i = 2k(p-1)$ or $2k(p-1)+1$.*

Proof: The group of automorphisms of $\mathbb{Z}_p$ is the multiplicative group $\mathbb{Z}_p^*$ of nonzero elements of $\mathbb{Z}_p$. Since p is prime, $\mathbb{Z}_p$ is a field and $\mathbb{Z}_p^*$ is cyclic of order $p-1$. Let r be a generator of $\mathbb{Z}_p^*$. Define a map $\varphi : S^\infty \times X^{\wedge p} \to S^\infty \times X^{\wedge p}$ permuting the factors X_j of $X^{\wedge p}$ by $\varphi(s, X_j) = (s^r, X_{rj})$ where subscripts are taken mod p and s^r means raise each coordinate of s, regarded as a unit vector in $\mathbb{C}^\infty$, to the r^{th} power and renormalize the resulting vector to have unit length. Then if γ is a generator of the $\mathbb{Z}_p$ action on $S^\infty \times X^{\wedge p}$, we have $\varphi(\gamma(s, X_j)) = \varphi(e^{2\pi i/p}s, X_{j-1}) = (e^{2r\pi i/p}s^r, X_{rj-r}) = \gamma^r(\varphi(s, X_j))$. This says that φ takes orbits to orbits, so φ induces maps $\varphi : \Gamma X \to \Gamma X$ and $\varphi : \Lambda X \to \Lambda X$. Restricting to the first coordinate, there is also an induced map $\varphi : L^\infty \to L^\infty$. Taking $X = K_n$, these maps fit into the diagram at the right. The square obviously commutes. The triangle commutes up to homotopy and a sign of $(-1)^n$ since it suffices to verify this on the pn-skeleton $(S^n)^{\wedge p}$, and here the map φ is an odd permutation of the S^n factors since it is a cyclic permutation of order $p-1$, which is even, and a transposition of two S^n factors has degree 1 if n is even and degree -1 if n is odd.

$$\begin{array}{ccccc} L^\infty \times K_n & \xrightarrow{\ \nabla\ } & \Lambda K_n & \xrightarrow{\ \lambda\ } & K_{pn} \\ {\scriptstyle \varphi \times 1\!\!1}\downarrow & & {\scriptstyle \varphi}\downarrow & \nearrow {\scriptstyle \lambda} & \\ L^\infty \times K_n & \xrightarrow{\ \nabla\ } & \Lambda K_n & & \end{array}$$

Suppose first that n is even. Then commutativity of the diagram means that $\sum_i \omega_{(p-1)n-i} \otimes \theta_i(\iota)$ is invariant under $\varphi^* \otimes 1\!\!1$, hence $\varphi^*(\omega_{(p-1)n-i}) = \omega_{(p-1)n-i}$ if $\theta_i(\iota)$ is nonzero. The map φ induces multiplication by r in $\pi_1(L^\infty)$, hence also in $H_1(L^\infty)$ and $H^1(L^\infty;\mathbb{Z}_p)$, sending ω_1 to $r\omega_1$. Since ω_2 was chosen to be the Bockstein of ω_1, it is also multiplied by r. We chose r to have order $p-1$ in $\mathbb{Z}_p^*$, so $\varphi^*(\omega_\ell) = \omega_\ell$ only when the total number of ω_1 and ω_2 factors in ω_ℓ is a multiple of $p-1$. For $\omega_\ell = \omega_2^k$ this means $\ell = (p-1)n - i = 2k(p-1)$, while for $\omega_\ell = \omega_1\omega_2^{k-1}$ it means ℓ is 1 less than this, $2k(p-1)-1$. Solving these equations for i gives $i = (n-2k)(p-1)$ or $i = (n-2k)(p-1)+1$. Since n is even this says that i is congruent to 0 or 1 mod $2(p-1)$, which is what the lemma asserts.

When n is odd the condition $\varphi^*(\omega_\ell) = \omega_\ell$ becomes $\varphi^*(\omega_\ell) = -\omega_\ell$. In the cyclic group $\mathbb{Z}_p^*$ the element -1 is the only element of order 2, and this element is $(p-1)/2$ times a generator, so the total number of ω_1 and ω_2 factors in ω_ℓ must be $(2k+1)(p-1)/2$ for some integer k. This implies that $\ell = (p-1)n - i =$

$2(2k+1)(p-1)/2$ or 1 less than this, hence $i=(n-2k-1)(p-1)$ or 1 greater than this. As n is odd, this again says that i is congruent to 0 or 1 mod $2(p-1)$. □

Since $\theta_0:H^n(X)\to H^n(X)$ is a cohomology operation that preserves dimension, it must be defined by a coefficient homomorphism $\mathbb{Z}_p\to\mathbb{Z}_p$, multiplication by some $a_n\in\mathbb{Z}_p$. We claim that these a_n's satisfy

$$a_{m+n}=(-1)^{p(p-1)mn/2}a_m a_n \qquad\text{and}\qquad a_n=(-1)^{p(p-1)n(n-1)/4}a_1^n$$

To see this, recall the formula $\lambda(\alpha\smile\beta)=(-1)^{p(p-1)mn/2}\lambda(\alpha)\lambda(\beta)$ for $|\alpha|=m$ and $|\beta|=n$. From the definition of the θ_i's it then follows that $\theta_0(\alpha\smile\beta)=(-1)^{p(p-1)mn/2}\theta_0(\alpha)\theta_0(\beta)$, which gives the first part of the claim. The second part follows from this by induction on n.

Lemma 4L.15. *$a_1=\pm m!$ for $m=(p-1)/2$, so $p=2m+1$.*

Proof: It suffices to compute $\theta_0(\alpha)$ where α is any nonzero 1-dimensional class, so the simplest thing is to choose α to be a generator of $H^1(S^1)$, say a generator coming from a generator of $H^1(S^1;\mathbb{Z})$. This determines α up to a sign. Since $H^i(S^1)=0$ for $i>1$, we have $\theta_i(\alpha)=0$ for $i>0$, so the defining formula for $\theta_0(\alpha)$ has the form $\nabla^*(\lambda(\alpha))=\omega_{p-1}\otimes\theta_0(\alpha)=a_1\omega_{p-1}\otimes\alpha$ in $H^p(L^\infty\times S^1)$. To compute a_1 there is no harm in replacing L^∞ by a finite-dimensional lens space, say L^p, the p-skeleton of L^∞. Thus we may restrict the bundle $\Lambda S^1\to L^\infty$ to a bundle $\Lambda^p S^1\to L^p$ with the same fibers $(S^1)^{\wedge p}=S^p$. We regard S^1 as the one-point compactification of $\mathbb{R}$ with basepoint the added point at infinity, and then $(S^1)^{\wedge p}$ becomes the one-point compactification of $\mathbb{R}^p$ with $\mathbb{Z}_p$ acting by permuting the coordinates of $\mathbb{R}^p$ cyclically, preserving the origin and the point at infinity. This action defines the bundle $\Gamma^p S^1\to L^p$ with fibers S^p, containing a zero section and a section at infinity, and $\Lambda^p S^1$ is obtained by collapsing the section at infinity. We can also describe $\Lambda^p S^1$ as the one-point compactification of the complement of the section at infinity in $\Gamma^p S^1$, since the base space L^p is compact. The complement of the section at infinity is a bundle $E\to L^p$ with fibers $\mathbb{R}^p$. In general, the one-point compactification of a fiber bundle E over a compact base space with fibers $\mathbb{R}^n$ is called the **Thom space** $T(E)$ of the bundle, and a class in $H^n(T(E))$ that restricts to a generator of H^n of the one-point compactification of each fiber $\mathbb{R}^n$ is called a **Thom class.** In our situation, $\lambda(\alpha)$ is such a Thom class.

Our first task is to construct subbundles $E_0,E_1,\cdots,E_m$ of E, where E_0 has fiber $\mathbb{R}$ and the other E_j's have fiber $\mathbb{R}^2$, so $p=2m+1$. The bundle E comes from the linear transformation $T:\mathbb{R}^p\to\mathbb{R}^p$ permuting the coordinates cyclically. We claim there is a decomposition $\mathbb{R}^p=V_0\oplus V_1\oplus\cdots\oplus V_m$ with V_0 1-dimensional and the other V_j's 2-dimensional, such that $T(V_j)=V_j$ for all j, with $T\,|\,V_0$ the identity and $T\,|\,V_j$ a rotation by the angle $2\pi j/p$ for $j>0$. Thus T defines an action of $\mathbb{Z}_p$ on V_j and we can define E_j just as E was defined, as the quotient $(S^p\times V_j)/\mathbb{Z}_p$ with respect to the diagonal action.

An easy way to get the decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ is to regard $\mathbb{R}^p$ as a module over the principal ideal domain $\mathbb{R}[t]$ by setting $tv = T(v)$ for $v \in \mathbb{R}^p$. Then $\mathbb{R}^p$ is isomorphic as a module to the module $\mathbb{R}[t]/(t^p - 1)$ since T permutes coordinates cyclically; this amounts to identifying the standard basis vectors $v_1, \cdots, v_p$ in $\mathbb{R}^p$ with $1, t, \cdots, t^{p-1}$. The polynomial $t^p - 1$ factors over $\mathbb{C}$ into the linear factors $t - e^{2\pi i j/p}$ for $j = 0, \cdots, p-1$. Combining complex conjugate factors, this gives a factorization over $\mathbb{R}$, $t^p - 1 = (t-1)\prod_{1 \le j \le m}(t^2 - 2(\cos\varphi_j)t + 1)$, where $\varphi_j = 2\pi j/p$. These are distinct monic irreducible factors, so the module $\mathbb{R}[t]/(t^p - 1)$ splits as $\mathbb{R}[t]/(t-1)\bigoplus_{1\le j \le m}\mathbb{R}[t]/(t^2 - 2(\cos\varphi_j)t + 1)$ by the basic structure theory of modules over a principal ideal domain. This translates into a decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ with $T(V_j) \subset V_j$. Here V_0 corresponds to $\mathbb{R}[t]/(t-1) \approx \mathbb{R}$ with t acting as the identity, and V_j for $j > 0$ corresponds to $\mathbb{R}[t]/(t^2 - 2(\cos\varphi_j)t + 1)$. The latter module is isomorphic to $\mathbb{R}^2$ with t acting as rotation by the angle φ_j since the characteristic polynomial of this rotation is readily computed to be $t^2 - 2(\cos\varphi_j)t + 1$, hence this rotation satisfies $t^2 - 2(\cos\varphi_j)t + 1 = 0$ so there is a module homomorphism $\mathbb{R}[t]/(t^2 - 2(\cos\varphi_j)t + 1) \to \mathbb{R}^2$ which is obviously an isomorphism.

From the decomposition $\mathbb{R}^p = V_0 \oplus V_1 \oplus \cdots \oplus V_m$ and the action of T on each factor we can see that the only vectors fixed by T are those in the line V_0. The vectors $(x, \cdots, x)$ are fixed by T, so V_0 must be this diagonal line.

Next we compute Thom classes for the bundles E_j. This is easy for E_0 which is the product $L^p \times \mathbb{R}$, so the projection $E_0 \to \mathbb{R}$ one-point compactifies to a map $T(E_0) \to S^1$ and we can pull back the chosen generator $\alpha \in H^1(S^1)$ to a Thom class for E_0. The other E_j's have 2-dimensional fibers, which we now view as $\mathbb{C}$ rather than $\mathbb{R}^2$. Just as E_j is the quotient of $S^p \times \mathbb{C}$ via the identifications $(v, z) \sim (e^{2\pi i/p}v, e^{2\pi i j/p}z)$, we can define a bundle $\overline{E}_j \to \mathbb{C}\mathrm{P}^m$ with fiber $\mathbb{C}$ by the identifications $(v,z) \sim (\lambda v, \lambda^j z)$ for $\lambda \in S^1 \subset \mathbb{C}$. We then have the left half of the commutative diagram shown at the right, where the quotient map $\widetilde{q}$ restricts to a homeomorphism on each fiber. The maps $\widetilde{f}$ and f are induced by the map $S^p \times \mathbb{C} \to S^p \times \mathbb{C}$ sending (v,z) to (v^j, z) where v^j means raise each coordinate of v to the j^{th} power and then rescale to get a vector of unit length. The map $\widetilde{f}$ is well-defined since equivalent pairs $(v,z) \sim (\lambda v, \lambda^j z)$ in $\overline{E}_j$ are carried to pairs (v^j, z) and $(\lambda^j v^j, \lambda^j z)$ that are equivalent in $\overline{E}_1$.

$$\begin{array}{ccccc} E_j & \xrightarrow{\widetilde{q}} & \overline{E}_j & \xrightarrow{\widetilde{f}} & \overline{E}_1 \\ \downarrow & & \downarrow & & \downarrow \\ L^p & \xrightarrow{q} & \mathbb{C}\mathrm{P}^m & \xrightarrow{f} & \mathbb{C}\mathrm{P}^m \end{array}$$

Since both $\widetilde{q}$ and $\widetilde{f}$ restrict to homeomorphisms in each fiber, they extend to maps of Thom spaces that pull a Thom class for $\overline{E}_1$ back to Thom classes for $\overline{E}_j$ and E_j. To construct a Thom class for $\overline{E}_1$, observe that the Thom space $T(\overline{E}_1)$ is homeomorphic to $\mathbb{C}\mathrm{P}^{m+1}$, namely, view the sphere $S^p = S^{2m+1}$ as the unit sphere in $\mathbb{C}^{m+1}$, and then the inclusion $S^p \times \mathbb{C} \hookrightarrow \mathbb{C}^{m+1} \times \mathbb{C} = \mathbb{C}^{m+2}$ induces a map $g : \overline{E}_1 \to \mathbb{C}\mathrm{P}^{m+1}$ since the equivalence relation defining $\overline{E}_1$ is $(v,z) \sim (\lambda v, \lambda z)$ for $\lambda \in S^1$. It is evident that g is a homeomorphism onto the complement of the point $[0,\cdots,0,1]$ in $\mathbb{C}P^{m+1}$, so sending the point at infinity in $T(\overline{E}_1)$ to $[0,\cdots,0,1]$ gives an extension of g to a homeomorphism $T(\overline{E}_1)\approx\mathbb{C}P^{m+1}$. Under this homeomorphism the one-point compactifications of the fibers of $\overline{E}_1$ correspond to the 2-spheres S^2_v consisting of $[0,\cdots,0,1]$ and the points $[v,z]\in\mathbb{C}P^{m+1}$ with fixed $v\in S^p$ and varying $z\in\mathbb{C}$. Each S^2_v is a $\mathbb{C}P^1$ in $\mathbb{C}P^{m+1}$ equivalent to the standard $\mathbb{C}P^1$ under a homeomorphism of $\mathbb{C}P^{m+1}$ coming from a linear isomorphism of $\mathbb{C}^{m+2}$, so a generator γ of $H^2(\mathbb{C}P^{m+1})$ is a Thom class, restricting to a generator of $H^2(S^2_v)$ for each v. We choose γ to be the $\mathbb{Z}_p$ reduction of a generator of $H^2(\mathbb{C}P^{m+1};\mathbb{Z})$, so γ is determined up to a sign.

A slightly different view of Thom classes will be useful. For the bundle $E\to L^p$, for example, we have isomorphisms

$$\begin{aligned}\widetilde{H}^*(T(E)) &\approx H^*(T(E),\infty) \quad \text{where } \infty \text{ is the compactification point}\\ &\approx H^*(T(E),T(E)-L^p) \quad \text{where } L^p \text{ is embedded in } T(E) \text{ as the zero}\\ &\qquad\qquad \text{section, so } T(E)-L^p \text{ deformation retracts onto } \infty\\ &\approx H^*(E,E-L^p) \quad \text{by excision}\end{aligned}$$

Thus we can view a Thom class as lying in $H^*(E,E-L^p)$, and similarly for the bundles E_j.

We have projections $\pi_j:E\to E_j$ via the projections $V_0\oplus V_1\oplus\cdots\oplus V_m\to V_j$ in fibers. If $\tau_j\in H^*(E_j,E_j-L^p)$ denotes the Thom class constructed above, then we have the pullback $\pi_j^*(\tau_j)\in H^*(E,E-\pi_j^{-1}(L^p))$, and the cup product $\prod_j\pi_j^*(\tau_j)$ in $H^*(E,E-L^p)$ is a Thom class for E, as one sees by applying the calculation at the end of Example 3.11 in each fiber. Under the isomorphism $H^*(E,E-L^p)\approx\widetilde{H}^*(T(E))$, the class $\prod_j\pi_j^*(\tau_j)$ corresponds to $\pm\lambda(\alpha)$ since both classes restrict to $\pm\alpha^{\otimes p}$ in each fiber $S^p\subset T(E)$ and $\lambda(\alpha)$ is uniquely determined by its restriction to fibers.

Now we can finish the proof of the lemma. The class $\nabla^*(\lambda(\alpha))$ is obtained by restricting $\lambda(\alpha)\in H^p(T(E))$ to the diagonal $T(E_0)$, then pulling back to $L^p\times S^1$ via the quotient map $L^p\times S^1\to T(E_0)$ which collapses the section at infinity to a point. Restricting $\prod_j\pi_j^*(\tau_j)$ to $H^p(E_0,E_0-L^p)\approx H^p(T(E_0))$ gives $\tau_0\smile e_1\smile\cdots\smile e_m$ where $e_j\in H^2(E_0)$ is the image of τ_j under $H^2(E_j,E_j-L^p)\to H^2(E_j)\approx H^2(L^p)\approx H^2(E_0)$, these last two isomorphisms coming from including L^p in E_j and E_0 as the zero section, to which they deformation retract. To compute e_j, we use the diagram

$$\begin{array}{ccccc}
H^2(E_j,E_j-L^p) & \xleftarrow{\;\widetilde{q}^*\;} & H^2(\overline{E}_j,\overline{E}_j-\mathbb{C}P^m) & \xleftarrow{\;\widetilde{f}^*\;} & H^2(\overline{E}_1,\overline{E}_1-\mathbb{C}P^m)\\
\downarrow & & \downarrow & & \downarrow\\
H^2(E_j) & \xleftarrow{\;\widetilde{q}^*\;} & H^2(\overline{E}_j) & \xleftarrow{\;\widetilde{f}^*\;} & H^2(\overline{E}_1)\\
\downarrow\approx & & \downarrow\approx & & \downarrow\approx\\
H^2(L^p) & \xleftarrow{\;q^*\;} & H^2(\mathbb{C}P^m) & \xleftarrow{\;f^*\;} & H^2(\mathbb{C}P^m)
\end{array}$$

The Thom class for $\overline{E}_1$ lies in the upper right group. Following this class across the top of the diagram and then down to the lower left corner gives the element e_j. To compute e_j we take the alternate route through the lower right corner of the diagram. The image of the Thom class for $\overline{E}_1$ in the lower right $H^2(\mathbb{C}\mathrm{P}^m)$ is the generator γ since $T(\overline{E}_1) = \mathbb{C}\mathrm{P}^{m+1}$. The map f^* is multiplication by j since f has degree j on $\mathbb{C}\mathrm{P}^1 \subset \mathbb{C}\mathrm{P}^m$. And $q^*(\gamma) = \pm\omega_2$ since q restricts to a homeomorphism on the 2-cell of L^p in the CW structure defined in Example 2.43. Thus $e_j = \pm j\omega_2$, and so $\tau_0 \smile e_1 \smile \cdots \smile e_m = \pm m!\tau_0 \smile \omega_2^m = \pm m!\tau_0 \smile \omega_{p-1}$. Since τ_0 was the pullback of α via the projection $T(E_0) \to S^1$, when we pull τ_0 back to $L^p \times S^1$ via ∇ we get $1 \otimes \alpha$, so $\tau_0 \smile e_1 \smile \cdots \smile e_m$ pulls back to $\pm m!\omega_{p-1} \otimes \alpha$. Hence $a_1 = \pm m!$. $\square$

The lemma implies in particular that a_n is not zero in $\mathbb{Z}_p$, so a_n has a multiplicative inverse a_n^{-1} in $\mathbb{Z}_p$. We then define

$$P^i(\alpha) = (-1)^i a_n^{-1}\theta_{2i(p-1)}(\alpha) \qquad \text{for } \alpha \in H^n(X)$$

The factor a_n^{-1} guarantees that P^0 is the identity. The factor $(-1)^i$ is inserted in order to make $P^i(\alpha) = \alpha^p$ if $|\alpha| = 2i$, as we show next. We know that $\theta_{2i(p-1)}(\alpha) = \alpha^p$, so what must be shown is that $(-1)^i a_{2i}^{-1} = 1$, or equivalently, $a_{2i} = (-1)^i$.

To do this we need a number theory fact: $((p-1)/2)!^2 \equiv (-1)^{(p+1)/2} \bmod p$. To derive this, note first that the product of all the elements $\pm 1, \pm 2, \cdots, \pm(p-1)/2$ of $\mathbb{Z}_p^*$ is $((p-1)/2)!^2(-1)^{(p-1)/2}$. On the other hand, this group is cyclic of even order, so the product of all its elements is the unique element of order 2, which is -1, since all the other nontrivial elements cancel their inverses in this product. Thus $((p-1)/2)!^2(-1)^{(p-1)/2} \equiv -1$ and hence $((p-1)/2)!^2 \equiv (-1)^{(p+1)/2} \bmod p$.

Using the formulas $a_n = (-1)^{p(p-1)n(n-1)/4}a_1^n$ and $a_1 = \pm((p-1)/2)!$ we then have

$$\begin{aligned} a_{2i} &= (-1)^{p(p-1)2i(2i-1)/4}((p-1)/2)!^{2i} \\ &= (-1)^{p[(p-1)/2]i(2i-1)}(-1)^{i(p+1)/2} \\ &= (-1)^{i(p-1)/2}(-1)^{i(p+1)/2} \quad \text{since } p \text{ and } 2i-1 \text{ are odd} \\ &= (-1)^{ip} = (-1)^i \quad \text{since } p \text{ is odd.} \end{aligned}$$

Theorem 4L.16. *The operations P^i satisfy the properties* (1)–(6) *and the Adem relations.*

Proof: Naturality and the fact that $P^i(\alpha) = 0$ if $2i > |\alpha|$ are inherited from the θ_i's. Property (6) and the other half of (5) have just been shown above. For the Cartan formula we have, for $\alpha \in H^m$ and $\beta \in H^n$, $\lambda(\alpha \smile \beta) = (-1)^{p(p-1)mn/2}\lambda(\alpha)\lambda(\beta)$ and hence

$$\begin{aligned} &\textstyle\sum_i \omega_{(p-1)(m+n)-i} \otimes \theta_i(\alpha \smile \beta) = \\ &\qquad (-1)^{p(p-1)mn/2}\Big(\textstyle\sum_j \omega_{(p-1)m-j} \otimes \theta_j(\alpha)\Big)\Big(\textstyle\sum_k \omega_{(p-1)n-k} \otimes \theta_k(\beta)\Big) \end{aligned}$$

Recall that $\omega_{2r} = \omega_2^r$ and $\omega_{2r+1} = \omega_1\omega_2^r$, with $\omega_1^2 = 0$. Therefore terms with i even on the left side of the equation can only come from terms with j and k even on the right side. This leads to the second equality in the following sequence:

$$\begin{aligned}
P^i(\alpha\smile\beta) &= (-1)^i a_{m+n}^{-1}\theta_{2i(p-1)}(\alpha\smile\beta)\\
&= (-1)^i a_{m+n}^{-1}(-1)^{p(p-1)mn/2}\textstyle\sum_j \theta_{2(i-j)(p-1)}(\alpha)\theta_{2j(p-1)}(\beta)\\
&= \textstyle\sum_j (-1)^{i-j}a_m^{-1}\theta_{2(i-j)(p-1)}(\alpha)(-1)^j a_n^{-1}\theta_{2j(p-1)}(\beta)\\
&= \textstyle\sum_j P^{i-j}(\alpha)P^j(\beta)
\end{aligned}$$

Property (4), the invariance of P^i under suspension, follows from the Cartan formula just as in the case $p=2$, using the fact that P^0 is the only P^i that can be nonzero on 1-dimensional classes, by (5). The additivity property follows just as before.

It remains to prove the Adem relations for Steenrod powers. We will need a Bockstein calculation:

Lemma 4L.17. $\beta\theta_{2k} = -\theta_{2k+1}$.

Proof: Let us first reduce the problem to showing that $\beta\nabla^*(\lambda(\iota)) = 0$. If we compute $\beta\nabla^*(\lambda(\iota))$ using the product formula for β, we get

$$\beta\Big(\textstyle\sum_i \omega_{(p-1)n-i}\otimes\theta_i(\iota)\Big) = \textstyle\sum_i\Big(\beta\omega_{(p-1)n-i}\otimes\theta_i(\iota) + (-1)^i\omega_{(p-1)n-i}\otimes\beta\theta_i(\iota)\Big)$$

Since $\beta\omega_{2j-1} = \omega_{2j}$ and $\beta\omega_{2j} = 0$, the terms with $i = 2k$ and $i = 2k+1$ give $\sum_k \omega_{(p-1)n-2k}\otimes\beta\theta_{2k}(\iota)$ and $\sum_k \omega_{(p-1)n-2k}\otimes\theta_{2k+1}(\iota) - \sum_k \omega_{(p-1)n-2k-1}\otimes\beta\theta_{2k+1}(\iota)$, respectively. Thus the coefficient of $\omega_{(p-1)n-2k}$ in $\beta\nabla^*(\lambda(\iota))$ is $\beta\theta_{2k}(\iota)+\theta_{2k+1}(\iota)$, so if we assume that $\beta\nabla^*(\lambda(\iota)) = 0$, this coefficient must vanish since we are in the tensor product $H^*(L^\infty)\otimes H^*(K_n)$. So we get $\beta\theta_{2k}(\iota) = -\theta_{2k+1}(\iota)$ and hence $\beta\theta_{2k}(\alpha) = -\theta_{2k+1}(\alpha)$ for all α. Note that $\beta\theta_{2k+1} = 0$ from the coefficient of $\omega_{(p-1)n-2k-1}$. This also follows from the formula $\beta\theta_{2k} = -\theta_{2k+1}$ since $\beta^2 = 0$.

In order to show that $\beta\nabla^*(\lambda(\iota)) = 0$ we first compute $\beta\lambda(\iota)$. We may assume K_n has a single n-cell and a single $(n+1)$-cell, attached by a map of degree p. Let φ and ψ be the cellular cochains assigning the value 1 to the n-cell and the $(n+1)$-cell, respectively, so $\delta\varphi = p\psi$. In $K_n^{\wedge p}$ we then have

$$(*)\qquad \delta(\varphi^{\otimes p}) = \textstyle\sum_i(-1)^{in}\varphi^{\otimes i}\otimes\delta\varphi\otimes\varphi^{\otimes(p-i-1)} = p\sum_i(-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$$

where the tensor notation means cellular cross product, so for example $\varphi^{\otimes p}$ is the cellular cochain dual to the np-cell $e^n\times\cdots\times e^n$ of $K_n^{\wedge p}$. The formula $(*)$ holds also in ΛK_n since the latter space has only one $(np+1)$-cell not in $K_n^{\wedge p}$, with cellular boundary zero. Namely, this cell is the product of the 1-cell of L^∞ and the np-cell of $K_n^{\wedge p}$ with one end of this product attached to the np-cell by the identity map and the other end by the cyclic permutation T, which has degree $+1$ since p is odd, so these two terms in the boundary of this cell cancel, and there are no other terms since the rest of the attachment of this cell is at the basepoint.

Bockstein homomorphisms can be computed using cellular cochain complexes, so the formula $(*)$ says that $\sum_i(-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$ represents $\beta\lambda(\iota)$. Via the quotient map $\Gamma K_n \to \Lambda K_n$, the class $\lambda(\iota)$ pulls back to a class $\gamma(\iota)$ with $\beta\gamma(\iota)$ also represented by $\sum_i (-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$. To see what happens when we pull $\beta\gamma(\iota)$ back to $\beta\nabla^*(\lambda(\iota))$ via the inclusion $L^\infty\times K_n \hookrightarrow \Gamma K_n$, consider the commutative diagram at the right. In the left-hand square the maps π^* are induced by the covering space projections $\pi: S^\infty\times K_n^{\wedge p}\to\Gamma K_n$ and $\pi: S^\infty\times K_n\to L^\infty\times K_n$ arising from the free $\mathbb{Z}_p$ actions. The vertical maps are induced by the diagonal inclusion $S^\infty\times K \hookrightarrow S^\infty\times K_n^{\wedge p}$. The maps τ are the transfer homomorphisms defined in §3.G. Recall the definition: If $\pi:\widetilde{X}\to X$ is a p-sheeted covering space, a chain map $C_*(X)\to C_*(\widetilde{X})$ is defined by sending a singular simplex $\sigma:\Delta^k\to X$ to the sum of its p lifts to $\widetilde{X}$, and τ is the induced map on cohomology. The key property of τ is that $\tau\pi^*: H^*(X)\to H^*(X)$ is multiplication by p, for any choice of coefficient group, since when we project the p lifts of $\sigma:\Delta^k\to X$ back to X we get $p\sigma$. When X is a CW complex and $\widetilde{X}$ is given the lifted CW structure, then τ can also be defined in cellular cohomology by the same procedure.

$$\begin{array}{ccccc}
H^*(\Gamma K_n) & \xrightarrow{\ \pi^*\ } & H^*(S^\infty\times K_n^{\wedge p}) & \xrightarrow{\ \tau\ } & H^*(\Gamma K_n) \\
\downarrow & & \downarrow & & \downarrow \\
H^*(L^\infty\times K_n) & \xrightarrow{\ \pi^*\ } & H^*(S^\infty\times K_n) & \xrightarrow{\ \tau\ } & H^*(L^\infty\times K_n)
\end{array}$$

Let us compute the value of the upper τ in the diagram on $1\otimes\psi\otimes\varphi^{\otimes(p-1)}$ where '1' is the cellular cocycle assigning the value 1 to each 0-cell of S^∞. By the definition of τ we have $\tau(1\otimes\psi\otimes\varphi^{\otimes(p-1)}) = \sum_i T^i(\psi\otimes\varphi^{\otimes(p-1)})$ where $T:K_n^{\wedge p}\to K_n^{\wedge p}$ permutes the factors cyclically. It does not matter whether T moves coordinates one unit leftwards or one unit rightwards since we are summing over all the powers of T, so let us say T moves coordinates rightward. Then $T(\psi\otimes\varphi^{\otimes(p-1)}) = \varphi\otimes\psi\otimes\varphi^{\otimes(p-2)}$, with the last φ moved into the first position. This move is achieved by transposing this φ with each of the preceding $p-2$ φ's and with ψ. Transposing two φ's introduces a sign $(-1)^{n^2}$, and transposing φ with ψ introduces a sign $(-1)^{n(n+1)} = +1$, by the commutativity property of cross product. Thus the total sign introduced by T is $(-1)^{n^2(p-2)}$, which equals $(-1)^n$ since p is odd. Each successive iterate of T also introduces a sign of $(-1)^n$, so T^i introduces a sign $(-1)^{in}$ for $0\le i\le p-1$. Thus

$$\tau(1\otimes\psi\otimes\varphi^{\otimes(p-1)}) = \sum_i T^i(\psi\otimes\varphi^{\otimes(p-1)}) = \sum_i (-1)^{in}\varphi^{\otimes i}\otimes\psi\otimes\varphi^{\otimes(p-i-1)}$$

As observed earlier, this last cocycle represents the class $\beta\gamma(\iota)$.

Since $\beta\gamma(\iota)$ is in the image of the upper τ in the diagram, the image of $\beta\gamma(\iota)$ in $H^*(L^\infty\times K_n)$, which is $\nabla^*(\beta\lambda(\iota))$, is in the image of the lower τ since the right-hand square commutes. The map π^* in the lower row is obviously onto since S^∞ is contractible, so $\nabla^*(\beta\lambda(\iota))$ is in the image of the composition $\tau\pi^*$ across the bottom of the diagram. But this composition is multiplication by p, which is zero for $\mathbb{Z}_p$ coefficients, so $\beta\nabla^*(\lambda(\iota)) = \nabla^*(\beta\lambda(\iota)) = 0$. ☐

The derivation of the Adem relations now follows the pattern for the case $p=2$. We had the formula $\nabla_2^*\lambda_2^*(\iota) = \sum_{i,j}\omega_{(p-1)pn-i}\otimes\theta_i(\omega_{(p-1)n-j}\otimes\theta_j(\iota))$. Since we are letting $p = 2m+1$, this can be rewritten as $\sum_{i,j}\omega_{2mpn-i}\otimes\theta_i(\omega_{2mn-j}\otimes\theta_j(\iota))$. The only nonzero θ_i's are $\theta_{2i(p-1)} = (-1)^i a_n P^i$ and $\theta_{2i(p-1)+1} = -\beta\theta_{2i(p-1)}$ so we have

$$\begin{aligned}\textstyle\sum_{i,j}\omega_{2mpn-i}\otimes\theta_i(\omega_{2mn-j}\otimes\theta_j(\iota)) = {}&\\ &\textstyle\sum_{i,j}(-1)^{i+j}a_{2mn}a_n\omega_{2m(pn-2i)}\otimes P^i(\omega_{2m(n-2j)}\otimes P^j(\iota))\\ &\textstyle-\sum_{i,j}(-1)^{i+j}a_{2mn}a_n\omega_{2m(pn-2i)}\otimes P^i(\omega_{2m(n-2j)-1}\otimes\beta P^j(\iota))\\ &\textstyle-\sum_{i,j}(-1)^{i+j}a_{2mn}a_n\omega_{2m(pn-2i)-1}\otimes\beta P^i(\omega_{2m(n-2j)}\otimes P^j(\iota))\\ &\textstyle+\sum_{i,j}(-1)^{i+j}a_{2mn}a_n\omega_{2m(pn-2i)-1}\otimes\beta P^i(\omega_{2m(n-2j)-1}\otimes\beta P^j(\iota))\end{aligned}$$

Since m and n will be fixed throughout the discussion, we may factor out the nonzero constant $a_{2mn}a_n$. Then applying the Cartan formula to expand the P^i terms, using also the formulas $P^k(\omega_{2r}) = \binom{r}{k}\omega_{2r+2k(p-1)}$ and $P^k(\omega_{2r+1}) = \binom{r}{k}\omega_{2r+2k(p-1)+1}$ derived earlier in the section, we obtain

$$\begin{aligned}&\textstyle\sum_{i,j,k}(-1)^{i+j}\binom{m(n-2j)}{k}\omega_{2m(pn-2i)}\otimes\omega_{2m(n-2j+2k)}\otimes P^{i-k}P^j(\iota)\\ -&\textstyle\sum_{i,j,k}(-1)^{i+j}\binom{m(n-2j)-1}{k}\omega_{2m(pn-2i)}\otimes\omega_{2m(n-2j+2k)-1}\otimes P^{i-k}\beta P^j(\iota)\\ -&\textstyle\sum_{i,j,k}(-1)^{i+j}\binom{m(n-2j)}{k}\omega_{2m(pn-2i)-1}\otimes\omega_{2m(n-2j+2k)}\otimes\beta P^{i-k}P^j(\iota)\\ +&\textstyle\sum_{i,j,k}(-1)^{i+j}\binom{m(n-2j)-1}{k}\omega_{2m(pn-2i)-1}\otimes\omega_{2m(n-2j+2k)}\otimes P^{i-k}\beta P^j(\iota)\\ -&\textstyle\sum_{i,j,k}(-1)^{i+j}\binom{m(n-2j)-1}{k}\omega_{2m(pn-2i)-1}\otimes\omega_{2m(n-2j+2k)-1}\otimes\beta P^{i-k}\beta P^j(\iota)\end{aligned}$$

Letting $\ell = mn+j-k$, so that $n-2j+2k = pn-2\ell$, the first of these five summations becomes

$$\textstyle\sum_{i,j,\ell}(-1)^{i+j}\binom{m(n-2j)}{mn+j-\ell}\omega_{2m(pn-2i)}\otimes\omega_{2m(pn-2\ell)}\otimes P^{i+\ell-mn-j}P^j(\iota)$$

and similarly for the other four summations.

Now we bring in the symmetry property $\varphi_{rs} = (-1)^{rs+mnp}\varphi_{sr}$, where, as before, $\nabla_2^*\lambda_2^*(\iota) = \sum_{r,s}\omega_r\otimes\omega_s\otimes\varphi_{rs}$. Of the five summations, only the first has both ω terms with even subscripts, namely $r = 2m(pn-2i)$ and $s = 2m(pn-2\ell)$, so the coefficient of $\omega_r\otimes\omega_s$ in this summation must be symmetric with respect to switching i and ℓ, up to a sign which will be $+$ if we choose n to be even, as we will do. This gives the relation

$$(1)\qquad \textstyle\sum_j(-1)^{i+j}\binom{m(n-2j)}{mn+j-\ell}P^{i+\ell-mn-j}P^j(\iota) = \sum_j(-1)^{\ell+j}\binom{m(n-2j)}{mn+j-i}P^{i+\ell-mn-j}P^j(\iota)$$

Similarly, the second, third, and fourth summations involve ω's with subscripts of opposite parity, yielding the relation

$$\begin{aligned}(2)\qquad &\textstyle\sum_j(-1)^{i+j}\binom{m(n-2j)-1}{mn+j-\ell}P^{i+\ell-mn-j}\beta P^j(\iota) =\\ &\textstyle\sum_j(-1)^{\ell+j}\binom{m(n-2j)}{mn+j-i}\beta P^{i+\ell-mn-j}P^j(\iota) - \sum_j(-1)^{\ell+j}\binom{m(n-2j)-1}{mn+j-i}P^{i+\ell-mn-j}\beta P^j(\iota)\end{aligned}$$

The relations (1) and (2) will yield the two Adem relations, so we will not need to consider the relation arising from the fifth summation.

To get the first Adem relation from (1) we choose n and ℓ so that the left side of (1) has only one term, namely we take $n = 2(1+p+\cdots+p^{r-1})+2s$ and $\ell = mn+s$ for given integers r and s. Then

$$\binom{m(n-2j)}{mn+j-\ell} = \binom{p^r-1-(p-1)(j-s)}{j-s}$$

and if r is large, this binomial coefficient is 1 if $j = s$ and 0 otherwise since if the rightmost nonzero digit in the p-adic expansion of the 'denominator' $j-s$ is x, then the corresponding digit of the 'numerator' $(p-1)[(1+p+\cdots+p^{r-1})-(j-s)]$ is obtained by reducing $(p-1)(1-x)$ mod p, giving $x-1$, and $\binom{x-1}{x} = 0$. Then (1) becomes

$$(-1)^{i+s}P^iP^s(\iota) = \sum_j (-1)^{\ell+j}\binom{m(n-2j)}{mn+j-i}P^{i+s-j}P^j(\iota)$$

$$\text{or} \qquad P^iP^s(\iota) = \sum_j (-1)^{i+j}\binom{m(n-2j)}{mn+j-i}P^{i+s-j}P^j(\iota) \qquad \text{since} \quad \ell \equiv s \bmod 2$$

$$= \sum_j (-1)^{i+j}\binom{m(n-2j)}{i-pj}P^{i+s-j}P^j(\iota) \qquad \text{since} \quad \binom{x}{y} = \binom{x}{x-y}$$

$$= \sum_j (-1)^{i+j}\binom{p^r+(p-1)(s-j)-1}{i-pj}P^{i+s-j}P^j(\iota)$$

If r is large and $i < ps$, the term p^r in the binomial coefficient can be omitted since we may assume $i \ge pj$, hence $j < s$, so $-1+(p-1)(s-j) \ge 0$ and the p^r has no effect on the binomial coefficient if r is large. This shows the first Adem relation holds for the class ι, and the general case follows as in the case $p = 2$.

To get the second Adem relation we choose $n = 2p^r + 2s$ and $\ell = mn + s$. Reasoning as before, the left side of (2) then reduces to $(-1)^{i+s}P^i\beta P^s(\iota)$ and (2) becomes

$$P^i\beta P^s(\iota) = \sum_j (-1)^{i+j}\binom{(p-1)(p^r+s-j)}{i-pj}\beta P^{i+s-j}P^j(\iota)$$
$$- \sum_j (-1)^{i+j}\binom{(p-1)(p^r+s-j)-1}{i-pj-1}P^{i+s-j}\beta P^j(\iota)$$

This time the term p^r can be omitted if r is large and $i \le ps$. □

Exercises

1. Determine all cohomology operations $H^1(X;\mathbb{Z})\to H^n(X;\mathbb{Z})$, $H^2(X;\mathbb{Z})\to H^n(X;\mathbb{Z})$, and $H^1(X;\mathbb{Z}_p)\to H^n(X;\mathbb{Z}_p)$ for p prime.

2. Use cohomology operations to show that the spaces $(S^1\times\mathbb{C}P^\infty)/(S^1\times\{x_0\})$ and $S^3\times\mathbb{C}P^\infty$ are not homotopy equivalent.

3. Since there is a fiber bundle $S^2\to\mathbb{C}P^5\to\mathbb{H}P^2$ by Exercise 35 in §4.2, one might ask whether there is an analogous bundle $S^4\to\mathbb{H}P^5\to\mathbb{O}P^2$. Use Steenrod powers for the prime 3 to show that such a bundle cannot exist. [The Gysin sequence can be used to determine the map on cohomology induced by the bundle projection $\mathbb{H}P^5\to\mathbb{O}P^2$.]

4. Show there is no fiber bundle $S^7\to S^{23}\to\mathbb{O}P^2$. [Compute the cohomology ring of the mapping cone of the projection $S^{23}\to\mathbb{O}P^2$ via Poincaré duality or the Thom isomorphism.]

5. Show that the subalgebra of $\mathcal{A}_2$ generated by Sq^i for $i \le 2$ has dimension 8 as a vector space over $\mathbb{Z}_2$, with multiplicative structure encoded in the following diagram, where diagonal lines indicate left-multiplication by Sq^1 and horizontal lines indicate left-multiplication by Sq^2.