Chapter 2
The Adams Spectral Sequence

The Adams spectral sequence was invented as a tool for computing stable homotopy groups of spheres, and more generally the stable homotopy groups of any space. Let us begin by explaining the underlying idea of this spectral sequence.

As a first step toward computing the set $[X,Y]$ of homotopy classes of maps $X\to Y$ one could consider induced homomorphisms on homology. This produces a map $[X,Y]\to\text{Hom}(H_*(X),H_*(Y))$. The first interesting instance of this is the notion of degree for maps $S^n\to S^n$, where it happens that the degree computes $[S^n,S^n]$ completely. For maps between spheres of different dimension we get no information this way, however, so it is natural to look for more sophisticated structure. For a start we can replace homology by cohomology since this has cup products and their stable outgrowths, Steenrod squares and powers. Changing notation by switching the roles of X and Y for convenience, we then have a map $[Y,X]\to\text{Hom}_{\mathcal{A}}(H^*(X),H^*(Y))$ where $\mathcal{A}$ is the mod p Steenrod algebra and cohomology is taken with $\mathbb{Z}_p$ coefficients. Since cohomology and Steenrod operations are stable under suspension, it makes sense to change our viewpoint and let $[Y,X]$ now denote the stable homotopy classes of maps, the direct limit under suspension of the sets of maps $\Sigma^kY\to\Sigma^kX$. This has the advantage that the map $[Y,X]\to\text{Hom}_{\mathcal{A}}(H^*(X),H^*(Y))$ is a homomorphism of abelian groups, where cohomology is now to be interpreted as reduced cohomology since we want it to be stable under suspension.

Since $\text{Hom}_{\mathcal{A}}(H^*(X),H^*(Y))$ is just a subgroup of $\text{Hom}(H^*(X),H^*(Y))$, we are not yet using the real strength of the $\mathcal{A}$-module structure. To do this, recall that $\text{Hom}_{\mathcal{A}}$ is the $n=0$ case of a whole sequence of functors $\text{Ext}^n_{\mathcal{A}}$. Since $\mathcal{A}$ has such a complicated multiplicative structure, these higher $\text{Ext}^n_{\mathcal{A}}$ groups could be nontrivial and might carry quite a bit more information than $\text{Hom}_{\mathcal{A}}$ by itself. As evidence that there may be something to this idea, consider the functor $\text{Ext}^1_{\mathcal{A}}$. This measures whether short exact sequences of $\mathcal{A}$-modules split. For a map $f:S^k\to S^\ell$ with $k>\ell$ one can form the mapping cone C_f, and then associated to the pair (C_f,S^ℓ) there is a short exact sequence of $\mathcal{A}$-modules

$$0\to H^*(S^{k+1})\to H^*(C_f)\to H^*(S^\ell)\to 0$$

Additively this splits, but whether it splits over $\mathcal{A}$ is equivalent to whether $\mathcal{A}$ acts trivially in $H^*(C_f)$ since it automatically acts trivially on the two adjacent terms in

the short exact sequence. Since $\mathcal{A}$ is generated by the squares or powers, we are therefore asking whether some Sq^i or P^i is nontrivial in $H^*(C_f)$. For $p = 2$ this is the mod 2 Hopf invariant question, and for $p > 2$ it is the mod p analog. The answer for $p = 2$ is the theorem of Adams that Sq^i can be nontrivial only for $i = 1, 2, 4, 8$. For odd p the corresponding statement is that only P^1 can be nontrivial.

Thus $\text{Ext}^1_{\mathcal{A}}$ does indeed detect some small but nontrivial part of the stable homotopy groups of spheres. One could hardly expect the higher $\text{Ext}^n_{\mathcal{A}}$ functors to give a full description of stable homotopy groups, but the Adams spectral sequence says that, rather miraculously, they give a reasonable first approximation. In the case that Y is a sphere, the Adams spectral sequence will have the form

$$E_2^{s,t} = \text{Ext}^{s,t}_{\mathcal{A}}(H^*(X), \mathbb{Z}_p) \quad \text{converging to} \quad \pi^s_*(X)/\text{non-}p\text{-torsion}$$

Here the second index t in $\text{Ext}^{s,t}_{\mathcal{A}}$ denotes merely a grading of $\text{Ext}^s_{\mathcal{A}}$ arising from the usual grading of $H^*(X)$. The fact that torsion of order prime to p is factored out should be no surprise since one would not expect $\mathbb{Z}_p$ cohomology to give any information about non-p torsion.

More generally if Y is a finite CW complex and we define $\pi^Y_k(X) = [\Sigma^k Y, X]$, the stable homotopy classes of maps, then the Adams spectral sequence is

$$E_2^{s,t} = \text{Ext}^{s,t}_{\mathcal{A}}(H^*(X), H^*(Y)) \quad \text{converging to} \quad \pi^Y_*(X)/\text{non-}p\text{-torsion}$$

Taking $Y = S^0$ gives the earlier case, which suffices for the more common applications, but the general case illuminates the formal machinery, and is really no more difficult to set up than the special case. For the space X a modest hypothesis is needed for convergence, that it is a CW complex with finitely many cells in each dimension.

The Adams spectral sequence breaks the problem of computing stable homotopy groups of spheres up into three steps. First there is the purely algebraic problem of computing $\text{Ext}^{s,t}_{\mathcal{A}}(\mathbb{Z}_p, \mathbb{Z}_p)$. Since $\mathcal{A}$ is a complicated ring, this is not easy, but at least it is pure algebra. After this has been done through some range of values for s and t there remain the two problems one usually has with a spectral sequence, computing differentials and resolving ambiguous extensions. In practice it is computing differentials that is the most difficult. As with the Serre spectral sequence for cohomology, there will be a product structure that helps considerably.

The fact that the Steenrod algebra tells a great deal about stable homotopy groups of spheres should not be quite so surprising if one recalls the calculations done in §1.3. Here the Serre spectral sequence was used repeatedly to figure out successive stages in a Postnikov tower for a sphere. The main step was computing differentials by means of computations with Steenrod squares. One can think of the Adams spectral sequence as streamlining this process. There is one spectral sequence for all the p-torsion rather than one spectral sequence for the p-torsion in each individual homotopy group, and the algebraic calculation of the E_2 page replaces much of the

calculation of differentials in the Serre spectral sequences. As we will see, the first several stable homotopy groups of spheres can be computed completely without having to do any nontrivial calculations of differentials in the Adams spectral sequence. Eventually, however, hard work is involved in computing differentials, but we will stop well short of that point in the exposition here.

A Sketch of the Construction

Our approach to constructing the Adams spectral sequence will be to try to realize the algebraic definition of the Ext functors topologically. Let us recall how $\mathrm{Ext}_R^n(M,N)$ is defined, for modules M and N over a ring R. The first step is to choose a free resolution of M, an exact sequence

$$\cdots \longrightarrow F_2 \longrightarrow F_1 \longrightarrow F_1 \longrightarrow F_0 \longrightarrow M \longrightarrow 0$$

with each F_i a free R-module. Then one applies the functor $\mathrm{Hom}_R(-,N)$ to the free resolution, dropping the term $\mathrm{Hom}_R(M,N)$, to obtain a chain complex

$$\cdots \longleftarrow \mathrm{Hom}_R(F_2,N) \longleftarrow \mathrm{Hom}_R(F_1,N) \longleftarrow \mathrm{Hom}_R(F_0,N) \longleftarrow 0$$

Finally, the homology groups of this chain complex are by definition the groups $\mathrm{Ext}_R^n(M,N)$. It is a basic lemma that these do not depend on the choice of the free resolution of M.

Now we take R to be the Steenrod algebra $\mathcal{A}$ for some prime p and M to be $H^*(X)$, the reduced cohomology of a space X with $\mathbb{Z}_p$ coefficients, and we ask whether it is possible to construct a sequence of maps

$$X \longrightarrow K_0 \longrightarrow K_1 \longrightarrow K_2 \longrightarrow \cdots$$

that induces a free resolution of $H^*(X)$ as an $\mathcal{A}$-module:

$$\cdots \longrightarrow H^*(K_2) \longrightarrow H^*(K_1) \longrightarrow H^*(K_0) \longrightarrow H^*(X) \longrightarrow 0$$

Stated in this way, this is impossible because no space can have its cohomology a free $\mathcal{A}$-module. For if $H^*(K)$ were free as an $\mathcal{A}$-module then for each basis element α we would have $Sq^i\alpha$ nonzero for all i in the case $p=2$, or $P^i\alpha$ nonzero for all i when p is odd, but this contradicts the basic property of squares and powers that $Sq^i\alpha = 0$ for $i > |\alpha|$ and $P^i\alpha = 0$ for $i > |\alpha|/2$.

The spaces whose cohomology is closest to being free over $\mathcal{A}$ are Eilenberg-MacLane spaces. The cohomology $H^*(K(\mathbb{Z}_p,n))$ is free over $\mathcal{A}$ in dimensions less than $2n$, with one basis element, the fundamental class ι in H^n. This follows from the calculations in §1.3 since below dimension $2n$ there are only linear combinations of admissible monomials, and the condition that the monomials have excess less than n is automatically satisfied in this range. Alternatively, if one defines $\mathcal{A}$ as the limit of $H^*(K(\mathbb{Z}_p,n))$ as n goes to infinity, the freeness below dimension $2n$ is automatic from the Freudenthal suspension theorem. More generally, by taking a wedge sum of

$K(\mathbb{Z}_p, n_i)$'s with $n_i \geq n$ and only finitely many n_i's below any given N we would have a space with cohomology free over $\mathcal{A}$ below dimension $2n$. Instead of the wedge sum we could just as well take the product since this would have the same cohomology as the wedge sum below dimension $2n$.

Free modules have the good property that every module is the homomorphic image of a free module, and products of Eilenberg-MacLane spaces have an analogous property: For every space X there is a product K of Eilenberg-MacLane spaces and a map $X \to K$ inducing a surjection on cohomology. Namely, choose some set of generators α_i for $H^*(X)$, either as a group or more efficiently as an $\mathcal{A}$-module, and then there are maps $f_i : X \to K(\mathbb{Z}_p, |\alpha_i|)$ sending fundamental classes to the α_i's, and the product of these maps induces a surjection on H^*.

Using this fact, we construct a diagram

$$\begin{array}{ccccccccc} X & \longrightarrow & K_0 & & \longrightarrow & & K_1 & & \longrightarrow & & K_2 & \longrightarrow & \cdots \\ & & & \searrow & & \nearrow & & \searrow & & \nearrow & & & \\ & & & & K_0/X = X_1 & & & & K_1/X_1 = X_2 & & K_2/X_2 = X_3 & & \end{array}$$

by the following inductive procedure. Start with a map $X \to K_0$ to a product of Eilenberg-MacLane spaces inducing a surjection on H^*. Then after replacing this map by an inclusion via a mapping cylinder, let $X_1 = K_0/X$ and repeat the process with X_1 in place of $X = X_0$, choosing a map $X_1 \to K_1$ to another product of Eilenberg-MacLane spaces inducing a surjection on H^*, and so on. Thus we have a diagram of short exact sequences

$$\begin{array}{ccccccccccc} 0 & \longleftarrow & H^*(X) & \longleftarrow & H^*(K_0) & \longleftarrow & H^*(K_1) & \longleftarrow & H^*(K_2) & \longleftarrow & \cdots \\ & & & & & \nwarrow \swarrow & & \nwarrow \swarrow & & \nwarrow \swarrow & \\ & & & & & H^*(X_1) & & H^*(X_2) & & H^*(X_3) & \\ & & & & & \swarrow & \nwarrow & \swarrow & \nwarrow & \swarrow & \nwarrow \\ & & & & & 0 & & 0 & & 0 & & 0 \end{array}$$

The sequence across the top is exact, so we have a resolution of $H^*(X)$ which would be a free resolution if the modules $H^*(K_i)$ were free over $\mathcal{A}$.

Since stable homotopy groups are a homology theory, when we apply them to the cofibrations $X_i \to K_i \to K_i/X_i = X_{i+1}$ we obtain a staircase diagram

$$\begin{array}{ccccccccccc} & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \longrightarrow & \pi_{t+1}X_s & \longrightarrow & \pi_{t+1}K_s & \longrightarrow & \pi_{t+1}X_{s+1} & \longrightarrow & \pi_{t+1}K_{s+1} & \longrightarrow & \pi_{t+1}X_{s+2} & \longrightarrow \\ & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\ \longrightarrow & \pi_t X_{s-1} & \longrightarrow & \pi_t K_{s-1} & \longrightarrow & \pi_t X_s & \longrightarrow & \pi_t K_s & \longrightarrow & \pi_t X_{s+1} & \longrightarrow \\ & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\ \longrightarrow & \pi_{t-1}X_{s-2} & \longrightarrow & \pi_{t-1}K_{s-2} & \longrightarrow & \pi_{t-1}X_{s-1} & \longrightarrow & \pi_{t-1}K_{s-1} & \longrightarrow & \pi_{t-1}X_s & \longrightarrow \\ & \downarrow & & & & \downarrow & & & & \downarrow & \end{array}$$

and hence a spectral sequence. Since it is stable homotopy groups we are interested in, we may assume X has been suspended often enough to be highly connected, say n-connected, and then all the spaces K_i and X_i can be taken to be n-connected as well. Then below dimension $2n$ the cohomology $H^*(K_i)$ is a free $\mathcal{A}$-module and the

stable homotopy groups of K_i coincide with its ordinary homotopy groups, hence are very simple. As we will see, these two facts allow the E^1 terms of the spectral sequence to be identified with $\mathrm{Hom}_{\mathcal{A}}$ groups and the E^2 terms with $\mathrm{Ext}_{\mathcal{A}}$ groups, at least in the range of dimensions below $2n$. The Adams spectral sequence can be obtained from the exact couple above by repeated suspension and passing to a limit as n goes to infinity. In practice this is a little awkward, and a much cleaner and more elegant way to proceed is to do the whole construction with spectra instead of spaces, so this is what we will do instead.

2.1 Spectra

The derivation of the Adams spectral sequence will be fairly easy once we have available some basic facts about spectra, so our first task will be to develop these facts. The theme here will be that spectra are much like spaces, but are better in a few key ways, behaving more like abelian groups than spaces.

A spectrum consists of a sequence of basepointed spaces X_n, $n \geq 0$, together with basepoint-preserving maps $\Sigma X_n \to X_{n+1}$. In the realm of spaces with basepoints the suspension ΣX_n should be taken to be the reduced suspension, with the basepoint cross I collapsed to a point. The two examples of spectra we will have most to do with are:

— The suspension spectrum of a space X. This has $X_n = \Sigma^n X$ with $\Sigma X_n \to X_{n+1}$ the identity map.

— An Eilenberg-MacLane spectrum for an abelian group G. Here X_n is a CW complex $K(G,n)$ and $\Sigma K(G,n) \to K(G,n+1)$ is the adjoint of a map giving a CW approximation $K(G,n) \to \Omega K(G,n+1)$. More generally we could shift dimensions and take $X_n = K(G,m+n)$ for some fixed m, with maps $\Sigma K(G,m+n) \to K(G,m+n+1)$ as before.

The idea of spectra is that they should be the objects of a category that is the natural domain for stable phenomena in homotopy theory. In particular, the homotopy groups of the suspension spectrum of a space X should be the stable homotopy groups of X. With this aim in mind, one defines $\pi_i(X)$ for an arbitrary spectrum $X = \{X_n\}$ to be the direct limit of the sequence

$$\cdots \to \pi_{i+n}(X_n) \xrightarrow{\Sigma} \pi_{i+n+1}(\Sigma X_n) \to \pi_{i+n+1}(X_{n+1}) \xrightarrow{\Sigma} \pi_{i+n+2}(\Sigma X_{n+1}) \to \cdots$$

Here the unlabeled map is induced by the map $\Sigma X_n \to X_{n+1}$ that is part of the structure of the spectrum X. For a suspension spectrum these are the identity maps, so the homotopy groups of the suspension spectrum of a space X are the stable homotopy groups of X. For the Eilenberg-MacLane spectrum with $X_n = K(G,m+n)$ the Freudenthal suspension theorem implies that the map $\Sigma K(G,m+n) \to K(G,m+n+1)$ induces an isomorphism on ordinary homotopy groups up to dimension approximately $m+n$, so the spectrum has π_i equal to G for $i = m$ and zero otherwise, just as for an Eilenberg-MacLane space.

The homology groups of a spectrum can be defined in the same way, and in this case the suspension maps Σ are isomorphisms on homology. For cohomology, however, this definition in terms of limits would involve inverse limits rather than direct limits, and inverse limits are not as nice as direct limits since they do not generally preserve exactness, so we will give a different definition of cohomology for spectra. For suspension spectra and Eilenberg-MacLane spectra the definition in terms of inverse limits turns out to give the right thing since the limits are achieved at a finite stage. But for the construction of the Adams spectra sequence we have to deal with

more general spectra than these, so we need a general definition of the cohomology a spectrum. The definition should be such that the fundamental property of CW complexes that $H^n(X;G)$ is homotopy classes of maps $X \to K(G,n)$ remains valid for spectra. Our task then is to give good definitions of CW spectra, their cohomology, and maps between them, so that this result is true.

CW Spectra

For a spectrum X whose spaces X_n are CW complexes it is always possible to find an equivalent spectrum of CW complexes for which the structure maps $\Sigma X_n \to X_{n+1}$ are inclusions of subcomplexes, since one can first deform the structure maps to be cellular and then replace each X_n by the union of the reduced mapping cylinders of the maps

$$\Sigma^n X_0 \longrightarrow \Sigma^{n-1} X_1 \longrightarrow \cdots \longrightarrow \Sigma X_{n-1} \longrightarrow X_n$$

This leads us to define a **CW spectrum** to be a spectrum X consisting of CW complexes X_n with the maps $\Sigma X_n \hookrightarrow X_{n+1}$ inclusions of subcomplexes. The basepoints are assumed to be 0-cells. For example, the suspension spectrum associated to a CW complex is certainly a CW spectrum. An Eilenberg-MacLane CW spectrum with X_n a $K(G, m+n)$ can be constructed inductively by letting X_{n+1} be obtained from ΣX_n by attaching cells to kill π_i for $i > m+n+1$. By the Freudenthal theorem the attached cells can be taken to have dimension greater than $2m+2n$, approximately.

In a CW spectrum X each nonbasepoint cell e_α^i of X_n becomes a cell e_α^{i+1} of X_{n+1}. Regarding these two cells as being equivalent, one can define the cells of X to be the equivalence classes of nonbasepoint cells of all the X_n's. Thus a cell of X consists of cells e_α^{k+n} of X_n for all $n \geq n_\alpha$ for some n_α. The dimension of this cell is said to be k. The terminology is chosen so that for the suspension spectrum of a CW complex the definitions agree with the usual ones for CW complexes.

The cells of a spectrum can have negative dimension. A somewhat artificial example is the CW spectrum X with X_n the infinite wedge sum $S^1 \vee S^2 \vee \cdots$ for each n and with $\Sigma X_n \hookrightarrow X_{n+1}$ the evident inclusion. In this case there is one cell in every dimension, both positive and negative. There are other less artificial examples that arise in some contexts, but for the Adams spectral sequence we will only be concerned with CW spectra whose cells have dimensions that are bounded below. Such spectra are called **connective**. For a connective spectrum the connectivity of the spaces X_n goes to infinity as n goes to infinity.

The homology and cohomology groups of a CW spectrum X can be defined in terms of cellular chains and cochains. If one considers cellular chains relative to the basepoint, then the inclusions $\Sigma X_n \hookrightarrow X_{n+1}$ induce inclusions $C_*(X_n;G) \hookrightarrow C_*(X_{n+1};G)$ with a dimension shift to account for the suspension. The union $C_*(X;G)$ of this increasing sequence of chain complexes is then a chain complex having one G summand for each cell of X, just as for CW complexes. We define $H_i(X;G)$ to be the

i^{th} homology group of this chain complex $C_*(X;G)$. Since homology commutes with direct limits, this is the same as the direct limit of the homology groups $H_{i+n}(X_n;G)$. Note that this can be nonzero for negative values of i, as in the earlier example having $X_n = \bigvee_k S^k$ for each n, which has $H_i(X;\mathbb{Z}) = \mathbb{Z}$ for all $i \in \mathbb{Z}$.

For cohomology we define $C^*(X;G)$ to be simply the dual cochain complex, so $C^i(X;G)$ is $\mathrm{Hom}(C_i(X;\mathbb{Z}),G)$, the functions assigning an element of G to each cell of X, and $H^*(X;G)$ is defined to be the homology of this cochain complex. This assures that the universal coefficient theorem remains valid, for example.

A CW spectrum is said to be **finite** if it has just finitely many cells, and it is of **finite type** if it has only finitely many cells in each dimension. If X is of finite type then for each i there is an n such that X_n contains all the i-cells of X. It follows that $H_i(X;G) = H_i(X_n;G)$ for all sufficiently large n, and the same is true for cohomology. The corresponding statement for homotopy groups is not always true, as the example with $X_n = \bigvee_k S^k$ for each n shows. In this case the groups $\pi_{i+n}(X_n)$ never stabilize since, for example, there are elements of $\pi_{2p}(S^3)$ of order p that are stably nontrivial, for all primes p. But for a connective CW spectrum of finite type the groups $\pi_{i+n}(X_n)$ do eventually stabilize by the Freudenthal theorem.

Maps between CW Spectra

Now we come to the slightly delicate question of how to define a map between CW spectra. A reasonable goal would be that a map $f: X \to Y$ of CW spectra should induce maps $f_*: \pi_i(X) \to \pi_i(Y)$, and likewise for homology and cohomology. Certainly a sequence of basepoint-preserving maps $f_n: X_n \to Y_n$ forming commutative diagrams as at the right would induce maps on homotopy groups, and also on homology and cohomology groups if the individual f_n's were cellular. Let us call such a map f a **strict map**, since it is not the most general sort of map that works. For example, it would suffice to have the maps f_n defined only for all sufficiently large n. This would be enough to yield an induced map on π_i, thinking of $\pi_i(X)$ as $\varinjlim \pi_{i+n}(X_n)$ and $\pi_i(Y)$ as $\varinjlim \pi_{i+n}(Y_n)$. If the maps f_n were cellular there would also be an induced chain map $C_*(X) \to C_*(Y)$ and hence induced maps on H_* and H^*.

$$\begin{array}{ccc} \Sigma X_n & \longrightarrow & X_{n+1} \\ \downarrow{\scriptstyle \Sigma f_n} & & \downarrow{\scriptstyle f_{n+1}} \\ \Sigma Y_n & \longrightarrow & Y_{n+1} \end{array}$$

It turns out that a weaker condition will suffice: For each cell e^i_α of an X_n, the map f_{n+k} is defined on $\Sigma^k e^i_\alpha$ for all sufficiently large k. Here each f_n should be defined on a subcomplex $X'_n \subset X_n$ such that $\Sigma X'_n \subset X'_{n+1}$. Such a sequence of subcomplexes is called a **subspectrum** of X. The condition that for each n and each cell e^i_α of X_n the cell $\Sigma^k e^i_\alpha$ belongs to X'_{n+k} for all sufficiently large k is what is meant by saying that X' is a **cofinal** subspectrum of X. Thus we define a **map of CW spectra** $f: X \to Y$ to be a strict map $X' \to Y$ for some cofinal subspectrum X' of X. If the maps $f_n: X'_n \to Y_n$ defining f are cellular it is clear that there is an induced chain map $f_*: C_*(X) \to C_*(Y)$ and hence induced maps on homology and cohomology. A map of CW spectra $f: X \to Y$ also induces maps $f_*: \pi_i(X) \to \pi_i(Y)$ since each map $S^{i+n} \to X_n$

has compact image contained in a finite union of cells, whose k-fold suspensions lie in X'_{n+k} for sufficiently large k, and similarly for homotopies $S^{i+n}\times I\to X_n$.

Two maps of CW spectra $X\to Y$ are regarded as the same if they take the same values on a common cofinal subspectrum. Since the intersection of two cofinal subspectra is a cofinal subspectrum, this amounts to saying that replacing the cofinal subspectrum on which a spectrum map is defined by a smaller cofinal subspectrum is regarded as giving the same map.

It needs to be checked that the composition of two spectrum maps $X\xrightarrow{f} Y\xrightarrow{g} Z$ is defined. If f and g are given by strict maps on subspectra X' and Y', let X'' be the subspectrum of X' consisting of the cells of the complexes X'_n mapped by f to Y'_n. Then X'' is also cofinal in X' and hence in X since f takes each cell e^i_α of X'_n to a union of finitely many cells of Y_n, suspending to cells of Y'_{n+k} for some k since Y' is cofinal in Y, and then f_{n+k} takes $\Sigma^k e^i_\alpha$ to Y'_{n+k} so $\Sigma^k e^i_\alpha$ is in X''_{n+k}. Thus X'' is cofinal in X and the composition gf is a strict map $X''\to Z$.

The inclusion of a subspectrum X' into a spectrum X is of course a map of spectra, in fact a strict map. If X' is cofinal in X then the identity maps $X'_n\to X'_n$ define a map $X\to X'$ which is an inverse to the inclusion $X'\hookrightarrow X$, in the sense that the compositions of these two maps, in either order, are the identity. This means that a spectrum is always equivalent to any cofinal subspectrum.

For example, for any spectrum X the subspectrum X' with X'_n defined to be $\Sigma X_{n-1}\subset X_n$ is cofinal and hence equivalent to X. This means that every spectrum X is equivalent to the suspension of another spectrum. Namely, if we define the suspension ΣY of a spectrum Y by setting $(\Sigma Y)_n = \Sigma Y_n$, then a given spectrum X is equivalent to ΣY for Y the spectrum with $Y_n = X_{n-1}$. It is reasonable to denote this spectrum Y by $\Sigma^{-1}X$, so that $X = \Sigma(\Sigma^{-1}X)$. More generally we could define $\Sigma^k X$ for any $k\in\mathbb{Z}$ by setting $(\Sigma^k X)_n = X_{n+k}$, where X_{n+k} is taken to be the basepoint if $n+k<0$. (Alternatively, we could define spectra in terms of sequences X_n for $n\in\mathbb{Z}$, and then use the fact that such a spectrum is equivalent to the cofinal subspectrum obtained by replacing X_n for $n<0$ with the basepoint.)

A homotopy of maps between spectra is defined as one would expect, as a map $X\times I\to Y$, where $X\times I$ is the spectrum with $(X\times I)_n = X_n\times I$, this being the reduced product, with basepoint cross I collapsed to a point, so that $\Sigma(X_n\times I) = \Sigma X_n\times I$. The set of homotopy classes of maps $X\to Y$ is denoted $[X,Y]$. When X is S^i, by which we mean the suspension spectrum of the sphere S^i, we have $[S^i,Y] = \pi_i(Y)$ since spectrum maps $S^i\to Y$ are space maps $S^{i+n}\to Y_n$ for some n, and spectrum homotopies $S^i\times I\to Y$ are space homotopies $S^{i+n}\times I\to Y_n$ for some n.

One way in which spectra are better than spaces is that $[X,Y]$ is always a group, in fact an abelian group, since as noted above, every CW spectrum X is equivalent to a suspension spectrum, hence also to a double suspension spectrum, allowing an abelian sum operation to be defined just as in ordinary homotopy theory. The sus-

pension map $[X,Y]\to[\Sigma X,\Sigma Y]$ is a homomorphism, and in fact an isomorphism, as one can see in the following way. To show surjectivity, start with a map $f:\Sigma X\to\Sigma Y$, which we may assume is a strict map. For clarity write this as $f:X\wedge S^1\to Y\wedge S^1$, consisting of map $f_n:X_n\wedge S^1\to Y_n\wedge S^1$. Passing to cofinal subspectra, we can replace this by its restriction $\Sigma f_{n-1}:\Sigma(X_{n-1}\wedge S^1)\to\Sigma(Y_{n-1}\wedge S^1)$. The parentheses here are redundant and can be omitted. This map is independent of the suspension coordinate Σ, and we want it to be independent of the last coordinate S^1. This can be achieved by a homotopy rotating the sphere ΣS^1 by 90 degrees. So Σf_{n-1} is homotopic to a map $h_n\wedge 1\!\!1$, as desired, proving surjectivity. Injectivity is similar using $X\times I$ in place of X.

The homotopy extension property is valid for CW spectra as well as for CW complexes. Given a map $f:X\to Y$ and a homotopy $F:A\times I\to Y$ of $f|A$ for a subspectrum A of X, we may assume these are given by strict maps, after passing to cofinal subspectra. Assuming inductively that F has already been extended over $X_n\times I$, we suspend to get a map $(\Sigma X_n\times I\to\Sigma Y_n\hookrightarrow Y_{n+1}$, then extend the union of this map with the given $A_{n+1}\times I\to Y_{n+1}$ over $X_{n+1}\times I$.

The cellular approximation theorem for CW spectra can be proved in the same way. To deform a map $f:X\to Y$ to be cellular, staying fixed on a subcomplex A where it is already cellular, we may assume we are dealing with strict maps, and that f is already cellular on X_n, hence also its suspension $\Sigma X_n\to Y_{n+1}$. Then we deform f to be cellular on X_{n+1}, staying fixed where it is already cellular, and extend this deformation to all of X to finish the induction step.

Whitehead's theorem also translates to spectra:

Proposition 2.1. *A map between CW spectra that induces isomorphisms on all homotopy groups is a homotopy equivalence.*

Proof: Without loss we may assume the map is cellular. We will use the same scheme as in the standard proof for CW complexes, showing that if $f:X\to Y$ induces isomorphisms on homotopy groups, then the mapping cylinder M_f deformation retracts onto X as well as Y. First we need to define the mapping cylinder of a cellular map $f:X\to Y$ of CW spectra. This is the CW spectrum M_f obtained by first passing to a strict map $f:X'\to Y$ for a cofinal subspectrum X' of X, then taking the usual reduced mapping cylinders of the maps $f_n:X'_n\to Y_n$. These form a CW spectrum since the mapping cylinder of Σf_n is the suspension of the mapping cylinder of f_n. Replacing X' by a cofinal subspectrum replaces the spectrum M_f by a cofinal subspectrum, so M_f is independent of the choice of X', up to equivalence. The usual deformation retractions of M_{f_n} onto Y_n give a deformation retraction of the spectrum M_f onto the subspectrum Y.

If f induces isomorphisms on homotopy groups, the relative groups $\pi_*(M_f,X)$ are zero, so the proof of the proposition will be completed by applying the following

result to the identity map of (M_f, X): □

Lemma 2.2. *If (Y, B) is a pair of CW spectra with $\pi_*(Y, B) = 0$ and (X, A) is an arbitrary pair of CW spectra, then every map $(X, A) \to (Y, B)$ is homotopic, staying fixed on A, to a map with image in B.*

Proof: The corresponding result for CW complexes is proved by the usual method of induction over skeleta, but if we filter a CW spectrum by its skeleta there may be no place to start the induction unless the spectrum is connective. To deal with nonconnective spectra we will instead use a different filtration. In a CW complex the closure of each cell is compact, hence is contained in a finite subcomplex. There is in fact a unique smallest such subcomplex, the intersection of all the finite subcomplexes containing the given cell. Define the *width* of the cell to be the number of cells in this minimal subcomplex. In the basepointed situation we do not count the basepoint 0-cell, so cells that attach only to the basepoint have width 1. Reduced suspension preserves width, so we have a notion of width for cells of a CW spectrum. The key fact is that cells of width k attach only to cells of width strictly less than k, if $k > 1$. Thus a CW spectrum X is filtered by its subspectra $X(k)$ consisting of cells of width at most k.

Using this filtration by width we can now prove the lemma. Suppose inductively that for a given map $f : (X, A) \to (Y, B)$, which we may assume is a strict map, we have a cofinite subspectrum $X'(k)$ of $X(k)$ for which we have constructed a homotopy of $f|X'(k)$ to a map to B, staying fixed on $A \cap X'(k)$. Choose a cofinite subspectrum $X'(k+1)$ of $X(k+1)$ with $X'(k+1) \cap X(k) = X'(k)$. This is possible since each cell of width $k+1$ will have some sufficiently high suspension that attaches only to cells in $X'(k)$. Extend the homotopy of $f|X'(k)$ to a homotopy of $f|X'(k+1)$. The restriction of the homotoped f to each cell of width $k+1$ then defines an element of $\pi_*(Y, B)$. Since $\pi_*(Y, B) = 0$, this restriction will be nullhomotopic after some number of suspensions. Thus after replacing $X'(k+1)$ by a cofinite subspectrum that still contains $X'(k)$, there will be a homotopy of the restriction of f to the new $X'(k+1)$ to a map to B. We may assume this homotopy is fixed on cells of A, so this finishes the induction step. In the end we have a cofinal subspectrum X' of X, the union of the $X'(k)$'s, with a homotopy of f on X' to a map to B, fixing A. □

Proposition 2.3. *If a CW spectrum X is n-connected in the sense that $\pi_i(X) = 0$ for $i \le n$, then X is homotopy equivalent to a CW spectrum with no cells of dimension $\le n$.*

In particular this says that a CW spectrum that is n-connected for some n is homotopy equivalent to a connective CW spectrum, so one could broaden the definition of a connective spectrum to mean one whose homotopy groups vanish below some dimension.

Another consequence of this proposition is the Hurewicz theorem for CW spectra: If a CW spectrum X is n-connected, then the Hurewicz map $\pi_{n+1}(X) \to H_{n+1}(X)$ is an isomorphism. This follows since if X has no cells of dimension $\leq n$ then the Hurewicz map $\pi_{n+1}(X) \to H_{n+1}(X)$ is the direct limit of the Hurewicz isomorphisms $\pi_{n+1+k}(X_k) \to H_{n+1+k}(X_k)$, hence is also an isomorphism.

Proof the Proposition: We can follow the same procedure as for CW complexes, constructing the desired CW spectrum Y and a map $Y \to X$ inducing isomorphisms on all homotopy groups by inductive process. To start, choose maps $S^{n+1+k_\alpha} \to X_{k_\alpha}$ representing generators of $\pi_{n+1}(X)$. These give a map of spectra $\bigvee_\alpha S^{n+1}_\alpha \to X$ inducing a surjection on π_{n+1}. Next choose generators for the kernel of this surjection and represent these generators by nullhomotopic maps from suitable suspensions of S^{n+1} to the corresponding suspensions of $\bigvee_\alpha S^{n+1}_\alpha$. Use these maps to attach cells to the wedge of spheres, producing a map $Y^1 \to X$ that induces an isomorphism on π_{n+1}. Now repeat the process for π_{n+2} and each successive higher homotopy group. □

Notice that if X has finitely generated homotopy groups, then we can choose the CW spectrum Y to be of finite type. Thus a connective CW spectrum with finitely generated homotopy groups is homotopy equivalent to a connective spectrum of finite type.

Cofibration Sequences

We have defined the mapping cylinder M_f for a map of CW spectra $f : X \to Y$, and the mapping cone C_f can be constructed in a similar way, by first passing to a strict map on a cofinite subspectrum X' and then taking the mapping cones of the maps $f_n : X'_n \to Y_n$. For an inclusion $A \hookrightarrow X$ the mapping cone can be written as $X \cup CA$. We would like to say that the quotient map $X \cup CA \to X/A$ collapsing CA is a homotopy equivalence, but first we need to specify what X/A means for a spectrum X and subspectrum A. In order for the quotients X_n/A_n to form a CW spectrum we need to assume that A is a closed subspectrum of X, meaning that if a cell of an X_n has an iterated suspension lying in A_{n+k} for some k, then the cell is itself in A_n. Any subspectrum is cofinal in its closure, the subspectrum consisting of cells of X having some suspension in A, so in case A is not closed we can first pass to its closure before taking the quotient X/A.

When A is closed in X the quotient map $X \cup CA \to X/A$ is a strict map consisting of the quotient maps $X_n \cup CA_n \to X_n/A_n$, which are homotopy equivalences of CW complexes. Whitehead's theorem for CW spectra then implies that the map $X \cup CA \to X/A$ is a homotopy equivalence of spectra. (This could also be proved directly.)

Thus for a pair (X, A) of CW spectra we have a cofibration sequence just like the one for CW complexes:

$$A \hookrightarrow X \longrightarrow X \cup CA \longrightarrow \Sigma A \hookrightarrow \Sigma X \longrightarrow \cdots$$

This implies that, just as for CW complexes, there is an associated long exact sequence

$$[A,Y] \longleftarrow [X,Y] \longleftarrow [X/A,Y] \longleftarrow [\Sigma A,Y] \longleftarrow [\Sigma X,Y] \longleftarrow \cdots$$

But unlike for CW complexes, there is also an exact sequence

$$[Y,A] \longrightarrow [Y,X] \longrightarrow [Y,X/A] \longrightarrow [Y,\Sigma A] \longrightarrow [Y,\Sigma X] \longrightarrow \cdots$$

To derive this it suffices to show that $[Y,A]\to[Y,X]\to[Y,X\cup CA]$ is exact. The composition of these two maps is certainly zero, so to prove exactness consider a map $f:Y\to X$ which becomes nullhomotopic after we include X in $X\cup CA$. A nullhomotopy gives a map $CY\to X\cup CA$ making a commutative square with f in the following diagram:

$$\begin{array}{ccccccccc}
Y & \xrightarrow{1\!\!1} & Y & \longrightarrow & CY & \longrightarrow & \Sigma Y & \xrightarrow{1\!\!1} & \Sigma Y \\
\vdots & & \downarrow f & & \downarrow & & \downarrow & & \downarrow \Sigma f \\
A & \xrightarrow{i} & X & \longrightarrow & X\cup CA & \longrightarrow & \Sigma A & \xrightarrow{\Sigma i} & \Sigma X
\end{array}$$

We can then automatically fill in the next two vertical maps to make homotopy-commutative squares. We observed earlier that the suspension map $[Y,A]\to[\Sigma Y,\Sigma A]$ is an isomorphism, so we can take the map $\Sigma Y\to\Sigma A$ in the diagram to be a suspension Σg for some $g:Y\to A$. Commutativity of the right-hand square gives $\Sigma f \simeq (\Sigma i)(\Sigma g) = \Sigma(ig)$, and this implies that $f\simeq ig$ since suspension is an isomorphism. This gives the desired exactness.

If we were dealing with spaces instead of spectra, the analog of the exactness of $[Y,A]\to[Y,X]\to[Y,X/A]$ would be the exactness of $[Y,F]\to[Y,E]\to[Y,B]$ for a fibration $F\to E\to B$. This exactness follows immediately from the homotopy lifting property. Thus when one is interested in homotopy properties of spectra, cofibrations can also be regarded as fibrations. For a cellular map $f:A\to X$ of CW spectra with mapping cone C_f, the sequence $[Y,\Sigma^{-1}(C_f)]\to[Y,A]\to[Y,X]$ is exact, so $\Sigma^{-1}C_f$ can be thought of as the fiber of f.

The second long exact sequence associated to a cofibration, in the case of a pair $(A\vee B,A)$, has the form

$$\cdots \longrightarrow [Y,A] \longrightarrow [Y,A\vee B] \longrightarrow [Y,B] \longrightarrow \cdots$$

and this sequence splits, so we deduce that the natural map $[Y,A\vee B]\to[Y,A]\oplus[Y,B]$ is an isomorphism. By induction this holds more generally for wedge sums of finitely many factors.

Cohomology and Eilenberg-MacLane Spectra

The long exact sequences we have constructed can be extended indefinitely in both directions since spectra can always be desuspended. In the case of the first long exact sequence this means that for a fixed spectrum Y the functors $h^i(X) = [\Sigma^{-i}(X),Y]$ define a reduced cohomology theory on the category of CW spectra. The wedge axiom $h^i(\bigvee_\alpha X_\alpha) = \prod_\alpha h^i(X_\alpha)$ is obvious.

In particular, we have a cohomology theory associated to the Eilenberg-MacLane spectrum $K = K(G,m)$ with $K_n = K(G,m+n)$, and this coincides with ordinary cohomology:

Proposition 2.4. *There are natural isomorphisms $H^m(X;G) \approx [X, K(G,m)]$ for all CW spectra X.*

The proof of the analogous result for CW complexes given in §4.3 of [AT] works equally well for CW spectra, and is in fact a little simpler since there is no need to talk about loopspaces since spectra can always be desuspended. It is also possible to give a direct proof that makes no use of generalities about cohomology theories, analogous to the direct proof for CW complexes. One takes the spaces $K_n = K(G, m+n)$ to have trivial $(m+n-1)$-skeleton, and then each cellular map $f: X \to K$ gives a cellular cochain c_f in X with coefficients in $\pi_m(K) = G$ sending an m-cell of X to the element of $\pi_m(K)$ determined by the restriction of f to this cell. One checks that this association $f \mapsto c_f$ satisfies several key properties: The cochain c_f is always a cocycle since f extends over $(m+1)$-cells; every cellular cocycle occurs as c_f for some f; and $c_f - c_g$ is a coboundary iff f is homotopic to g.

The identification $H^m(X;G) = [X, K(G,m)]$ allows cohomology operations to be defined for cohomology groups of spectra by taking compositions of the form $X \to K(G,m) \to K(H,k)$. Taking coefficients in $\mathbb{Z}_p$, this gives an action of the Steenrod algebra $\mathcal{A}$ on $H^*(X)$, making $H^*(X)$ a module over $\mathcal{A}$. This uses the fact that composition of maps of spectra satisfies the distributivity properties $f(g+h) = fg + fh$ and $(f+g)h = fh + gh$, the latter being valid when h is a suspension, which is no loss of generality if we are only interested in homotopy classes of maps. For spectra X of finite type this definition of an $\mathcal{A}$-module structure on $H^*(X)$ agrees with the definition using the usual $\mathcal{A}$-module structure on the cohomology of spaces and the identification of $H^*(X)$ with the inverse limit $\varprojlim H^{*+n}(X_n)$ since Steenrod operations are stable under suspension.

For use in the Adams spectral sequence we need a version of the splitting $[Y, A \vee B] = [Y,A] \oplus [Y,B]$ for certain infinite wedge sums. Here the distinction between infinite direct sums and infinite direct products becomes important. For an infinite wedge sum $\bigvee_\alpha X_\alpha$ the group $[Y, \bigvee_\alpha X_\alpha]$ can sometimes be the direct sum $\bigoplus_\alpha [Y, X_\alpha]$, for example if Y is a finite CW spectrum. This follows from the case of finite wedge sums by a direct limit argument since the image of any map $Y \to \bigvee_\alpha X_\alpha$ lies in the wedge sum of only finitely many factors by compactness. However, we will need cases when Y is not finite and $[Y, \bigvee_\alpha X_\alpha]$ is instead the direct product $\prod_\alpha [Y, X_\alpha]$. There is always a natural map $[Y, \bigvee_\alpha X_\alpha] \to \prod_\alpha [Y, X_\alpha]$ whose coordinates are obtained by composing with the projections of $\bigvee_\alpha X_\alpha$ onto its factors.

Proposition 2.5. *The natural map $[X, \bigvee_i K(G,n_i)] \to \prod_i [X, K(G,n_i)]$ is an isomorphism if X is a connective CW spectrum of finite type and $n_i \to \infty$ as $i \to \infty$.*

Proof: When X is finite the result is obviously true since we can omit the factors $K(G,n_i)$ with n_i greater than the maximum dimension of cells of X without affecting either $[X,\bigvee_i K(G,n_i)]$ or $\prod_i[X,K(G,n_i)]$. For the general case we use a limiting argument, expressing X as the union of its skeleta X^k, which are finite. Let $h^*(X)$ be the cohomology theory associated to the spectrum $\bigvee_i K(G,n_i)$, so $h^n(X) = [\Sigma^{-n}X, \bigvee_i K(G,n_i)]$. There is a short exact sequence

$$0 \longrightarrow \varprojlim{}^1 h^{n-1}(X^k) \longrightarrow h^n(X) \xrightarrow{\lambda} \varprojlim h^n(X^k) \longrightarrow 0$$

whose derivation for CW complexes in Theorem 3F.8 of [AT] applies equally well to CW spectra. The term $\varprojlim h^n(X^k)$ is just the product $\prod_i[X,K(G,n_i)]$ from the finite case, since the inverse limit of the finite products is the infinite product. So it remains to show that the $\varprojlim^1$ term vanishes.

We will use the Mittag-Leffler criterion, which says that $\varprojlim^1 G_k$ vanishes for a sequence of homomorphisms of abelian groups $\cdots \longrightarrow G_2 \xrightarrow{\alpha_2} G_1 \xrightarrow{\alpha_1} G_0$ if for each k the decreasing chain of subgroups of G_k formed by the images of the compositions $G_{k+n} \to G_k$ is eventually constant once n is sufficiently large. This holds in the present situation since the images of the maps $H^i(X^{k+n};G) \to H^i(X^k;G)$ are independent of n when $k+n > i$. (When $G = \mathbb{Z}_p$ these cohomology groups are finite so the groups G_k are all finite and the Mittag-Leffler condition holds automatically.)

The proof of the Mittag-Leffler criterion was relegated to the exercises in [AT], so here is a proof. Recall that $\varprojlim G_k$ and $\varprojlim^1 G_k$ are defined as the kernel and cokernel of the map $\delta : \prod_k G_k \to \prod_k G_k$ given by $\delta(g_k) = (g_k - \alpha_{k+1}(g_{k+1}))$, or in other words as the homology groups of the two-term chain complex

$$0 \to \prod_k G_k \xrightarrow{\delta} \prod_k G_k \to 0$$

Let $H_k \subset G_k$ be the image of the maps $G_{k+n} \to G_k$ for large n. Then α_k takes H_k to H_{k-1}, so the short exact sequences $0 \to H_k \to G_k \to G_k/H_k \to 0$ give rise to a short exact sequence of two-term chain complexes and hence a six-term associated long exact sequence of homology groups. The part of this we need is the sequence $\varprojlim^1 H_k \to \varprojlim^1 G_k \to \varprojlim^1 (G_k/H_k)$. The first of these three terms vanishes since the maps $\alpha_k : H_k \to H_{k-1}$ are surjections, so it suffices to show that the third term vanishes. For the sequence of quotients G_k/H_k the associated groups 'H_k' are zero, so it is enough to check that $\varprojlim^1 G_k = 0$ when the groups H_k are zero. In this case δ is surjective since a given sequence (g_k) is the image under δ of the sequence obtained by adding to each g_k the sum of the images in G_k of $g_{k+1}, g_{k+2}, \cdots$, a finite sum if $H_k = 0$. □

2.2 The Spectral Sequence

Having established the basic properties of CW spectra that we will need, we begin this section by filling in details of the sketch of the construction of the Adams spectral sequence given in the introduction to this chapter. Then we examine the spectral sequence as a tool for computing stable homotopy groups of spheres.

Constructing the Spectral Sequence

We will be dealing throughout with CW spectra that are connective and of finite type. This assures that all homotopy and cohomology groups are finitely generated. The coefficient group for cohomology will be $\mathbb{Z}_p$ throughout, with p a fixed prime. A comment on notation: We will no longer have to consider the spaces X_n that make up a spectrum X, so we will be free to use subscripts to denote different spectra, rather than the spaces in a single spectrum.

Let X be a connective CW spectrum of finite type. We construct a diagram

$$\begin{array}{ccccccccccc} X & \longrightarrow & K_0 & & \longrightarrow & & K_1 & & \longrightarrow & & K_2 & \longrightarrow & \cdots \\ & & & \searrow & & \nearrow & & \searrow & & \nearrow & & & \\ & & & & K_0/X = X_1 & & & & K_1/X_1 = X_2 & & K_2/X_2 = X_3 & & \end{array}$$

in the following way. Choose generators α_i for $H^*(X)$ as an $\mathcal{A}$-module, with at most finitely many α_i's in each group $H^k(X)$. These determine a map $X \to K_0$ where K_0 is a wedge of Eilenberg-MacLane spectra, and K_0 has finite type. Replacing the map $X \to K_0$ by an inclusion, we form the quotient $X_1 = K_0/X$. This is again a connective spectrum of finite type, so we can repeat the construction with X_1 in place of X. In this way the diagram is constructed inductively. Note that even if X is the suspension spectrum of a finite complex, as in the application to stable homotopy groups of spheres, the subsequent spectra X_s will no longer be of this special form.

The associated diagram of cohomology

$$\begin{array}{ccccccccccccc} 0 & \longleftarrow & H^*(X) & \longleftarrow & H^*(K_0) & \longleftarrow & H^*(K_1) & \longleftarrow & H^*(K_2) & \longleftarrow & & \cdots \\ & & & & & \nwarrow\swarrow & & \nwarrow\swarrow & & \nwarrow\swarrow & & & \\ & & & & & H^*(X_1) & & H^*(X_2) & & H^*(X_3) & & & \\ & & & & & \swarrow\nwarrow & & \swarrow\nwarrow & & \swarrow\nwarrow & & & \\ & & & & 0 & & 0 & & 0 & & 0 & & \end{array}$$

then gives a resolution of $H^*(X)$ by free $\mathcal{A}$-modules, by Proposition 2.5.

Now we fix a finite spectrum Y and consider the functors $\pi_t^Y(Z) = [\Sigma^t Y, Z]$. Applied to the cofibrations $X_s \to K_s \to X_{s+1}$ these give long exact sequences forming a staircase diagram

$$\begin{array}{ccccccccc}
 & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
\longrightarrow & \pi^Y_{t+1}X_s & \longrightarrow & \pi^Y_{t+1}K_s & \longrightarrow & \pi^Y_{t+1}X_{s+1} & \longrightarrow & \pi^Y_{t+1}K_{s+1} & \longrightarrow \ \pi^Y_{t+1}X_{s+2} \longrightarrow \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
\longrightarrow & \pi^Y_{t}X_{s-1} & \longrightarrow & \pi^Y_{t}K_{s-1} & \longrightarrow & \pi^Y_{t}X_{s} & \longrightarrow & \pi^Y_{t}K_{s} & \longrightarrow \ \pi^Y_{t}X_{s+1} \longrightarrow \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
\longrightarrow & \pi^Y_{t-1}X_{s-2} & \longrightarrow & \pi^Y_{t-1}K_{s-2} & \longrightarrow & \pi^Y_{t-1}X_{s-1} & \longrightarrow & \pi^Y_{t-1}K_{s-1} & \longrightarrow \ \pi^Y_{t-1}X_{s} \longrightarrow \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow &
\end{array}$$

so we have a spectral sequence, the Adams spectral sequence. The spectrum Y plays a relatively minor role in what follows, and the reader is free to take it to be the spectrum S^0 so that $\pi^Y_t(Z) = \pi_t(Z)$. The groups $\pi^Y_t(Z)$ are finitely generated when Z is a connective spectrum of finite type, as one can see by induction on the number of cells of Y.

There is another way of describing the construction of the spectral sequence which provides some additional insight, although it involves nothing more than a change in notation really. Let $X^n = \Sigma^{-n}X_n$ and $K^n = \Sigma^{-n}K_n$. Then the earlier horizontal diagram starting with X can be rewritten as a vertical tower as at the right. The spectra K^n are again wedges of Eilenberg-MacLane spectra, so this tower is reminiscent of a Postnikov tower. Let us call it an **Adams tower** for X. The staircase diagram can now be rewritten in the following form:

$$\begin{array}{ccc}
 & \vdots & \\
 & \downarrow & \\
X^2 & \longrightarrow & K^2 \\
 & \downarrow & \\
X^1 & \longrightarrow & K^1 \\
 & \downarrow & \\
X & \longrightarrow & K^0
\end{array}$$

$$\begin{array}{ccccccccc}
 & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
\longrightarrow & \pi^Y_{t-s+1}X^s & \longrightarrow & \pi^Y_{t-s+1}K^s & \longrightarrow & \pi^Y_{t-s}X^{s+1} & \longrightarrow & \pi^Y_{t}K^{s+1} & \longrightarrow \ \pi^Y_{t-s-1}X^{s+2} \longrightarrow \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
\longrightarrow & \pi^Y_{t-s+1}X^{s-1} & \longrightarrow & \pi^Y_{t-s+1}K^{s-1} & \longrightarrow & \pi^Y_{t-s}X^{s} & \longrightarrow & \pi^Y_{t}K^{s} & \longrightarrow \ \pi^Y_{t-s-1}X^{s+1} \longrightarrow \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
\longrightarrow & \pi^Y_{t-s+1}X^{s-2} & \longrightarrow & \pi^Y_{t-s+1}K^{s-2} & \longrightarrow & \pi^Y_{t-s}X^{s-1} & \longrightarrow & \pi^Y_{t}K^{s-1} & \longrightarrow \ \pi^Y_{t-s-1}X^{s} \longrightarrow \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow &
\end{array}$$

This has the small advantage that the groups π^Y_i in each column all have the same index i.

The E_1 and E_2 terms of the spectral sequence are easy to identify. Since K_s is a wedge of Eilenberg-MacLane spectra $K_{s,i}$, elements of $[Y, K_s]$ are tuples of elements of $H^*(Y)$, one for each summand $K_{s,i}$, in the appropriate group $H^{n_i}(Y)$. Since $H^*(K_s)$ is free over $\mathcal{A}$ this means that the natural map $[Y, K_s] \to \mathrm{Hom}^0_{\mathcal{A}}(H^*(K_s), H^*(Y))$ is an isomorphism. Here Hom^0 denotes homomorphisms that preserve degree, i.e., dimension. Replacing Y by $\Sigma^t Y$, we obtain a natural identification

$$[\Sigma^t Y, K_s] = \mathrm{Hom}^0_{\mathcal{A}}(H^*(K_s), H^*(\Sigma^t Y)) = \mathrm{Hom}^t_{\mathcal{A}}(H^*(K_s), H^*(Y))$$

where the superscript t denotes homomorphisms that lower degree by t. Thus if we set $E_1^{s,t} = \pi^Y_t(K_s)$, we have $E_1^{s,t} = \mathrm{Hom}^t_{\mathcal{A}}(H^*(K_s), H^*(Y))$.

The differential $d_1 : \pi^Y_t(K_s) \to \pi^Y_t(K_{s+1})$ is induced by the map $K_s \to K_{s+1}$ in the resolution of X constructed earlier. This implies that the E_1 page of the spectral

sequence consists of the complexes

$$0 \longrightarrow \mathrm{Hom}^t_{\mathcal{A}}(H^*(K_0), H^*(Y)) \longrightarrow \mathrm{Hom}^t_{\mathcal{A}}(H^*(K_1), H^*(Y)) \longrightarrow \cdots$$

The homology groups of this complex are by definition $\mathrm{Ext}^{s,t}_{\mathcal{A}}(H^*(X), H^*(Y))$, so we have $E_2^{s,t} = \mathrm{Ext}^{s,t}_{\mathcal{A}}(H^*(X), H^*(Y))$.

Theorem 2.6. *For X a connective CW spectrum of finite type, this spectral sequence converges to $\pi^Y_*(X)$ modulo torsion of order prime to p. In other words,*

(a) *For fixed s and t the groups $E_r^{s,t}$ are independent of r once r is sufficiently large, and the stable groups $E_\infty^{s,t}$ are isomorphic to the quotients $F^{s,t}/F^{s+1,t+1}$ for the filtration of $\pi^Y_{t-s}(X)$ by the images $F^{s,t}$ of the maps $\pi^Y_t(X_s) \to \pi^Y_{t-s}(X)$, or equivalently the maps $\pi^Y_{t-s}(X^s) \to \pi^Y_{t-s}(X)$.*

(b) *$\bigcap_n F^{s+n,t+n}$ is the subgroup of $\pi^Y_{t-s}(X)$ consisting of torsion elements of order prime to p.*

Thus we are filtering $\pi^Y_{t-s}(X)$ by how far its elements pull back in the Adams tower. Unlike in the Serre spectral sequence this filtration is potentially infinite, and in fact will be infinite if $\pi^Y_{t-s}(X)$ contains elements of infinite order since all the terms in the spectral sequence are finite-dimensional $\mathbb{Z}_p$ vector spaces. Namely $E_1^{s,t} = \mathrm{Hom}^t_{\mathcal{A}}(H^*(K_s), H^*(Y))$ is certainly a finite-dimensional $\mathbb{Z}_p$ vector space, so $E_r^{s,t}$ is as well.

Throughout the proof we will be dealing only with connective CW spectra of finite type, so we make this a standing hypothesis that will not be mentioned again.

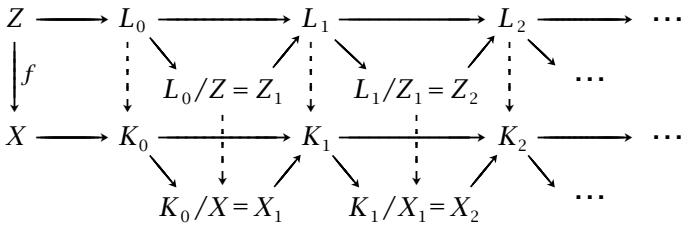
A key ingredient in the proof will be an analog for spectra of the algebraic lemma (Lemma 3.1 in [AT]) used to show that Ext is independent of the choice of free resolution. In order to state this we introduce some terminology. A sequence of maps of spectra $Z \to L_0 \to L_1 \to \cdots$ will be called a **complex** on Z if each composition of two successive maps is nullhomotopic. If the L_i's are wedges of Eilenberg-MacLane spectra $K(\mathbb{Z}_p, m_{ij})$ we call it an **Eilenberg-MacLane complex**. A complex for which the induced sequence $0 \leftarrow H^*(Z) \leftarrow H^*(L_0) \leftarrow \cdots$ is exact is a **resolution** of Z.

Lemma 2.7. *Suppose we are given the solid arrows in a diagram*

$$\begin{array}{ccccccccc} Z & \longrightarrow & L_0 & \longrightarrow & L_1 & \longrightarrow & L_2 & \longrightarrow & \cdots \\ \downarrow f & & \vdots f_0 & & \vdots f_1 & & \vdots f_2 & & \\ X & \longrightarrow & K_0 & \longrightarrow & K_1 & \longrightarrow & K_2 & \longrightarrow & \cdots \end{array}$$

where the first row is a resolution and the second row is an Eilenberg-MacLane complex. Then the dashed arrows can be filled in by maps $f_i : L_i \to K_i$ forming homotopy-commutative squares.

Proof: Since the compositions in a complex are nullhomotopic we may start with an enlarged diagram



where the triangles are homotopy-commutative. The map $X \to K_0$ is equivalent to a collection of classes $\alpha_j \in H^*(X)$. Since $H^*(L_0) \to H^*(Z)$ is surjective by assumption, there are classes $\beta_j \in H^*(L_0)$ mapping to the classes $f^*(\alpha_j) \in H^*(Z)$. These β_j's give a map $f_0 : L_0 \to K_0$ making a homotopy-commutative square with f. This square induces a map $L_0/Z \to K_0/X$ making another homotopy commutative square. The exactness property of the upper row implies that the map $H^*(L_1) \to H^*(L_0/Y)$ is surjective, so we can repeat the argument with Z and X replaced by $Z_1 = L_0/Z$ and $X_1 = K_1/X$ to construct the map f_1, and so on inductively for all the f_i's. □

Proof of Theorem 2.6: First we show statement (b). As noted earlier, all the terms $E_1^{s,t} = \mathrm{Hom}_{\mathcal{A}}^t(H^*(K_s), H^*(Y))$ in the staircase diagram are $\mathbb{Z}_p$ vector spaces, so by exactness all the vertical maps in the diagram are isomorphisms on non-p torsion. This implies that the non-p torsion in $\pi_{t-s}^Y(X)$ is contained in $\bigcap_n F^{s+n,t+n}$.

To prove the opposite inclusion we first do the special case that $\pi_*(X)$ is entirely p-torsion. These homotopy groups are then finite since we are dealing only with connective spectra of finite type. We construct a special Eilenberg-MacLane complex (not a resolution) of the form $X \to L_0 \to L_1 \to \cdots$ in the following way. Let $\pi_n(X)$ be the first nonvanishing homotopy group of X. Then let L_0 be a wedge of $K(\mathbb{Z}_p, n)$'s with one factor for each element of a basis for $H^n(X)$, so there is a map $X \to L_0$ inducing an isomorphism on H^n. This map is also an isomorphism on H_n, so on π_n it is the map $\pi_n(X) \to \pi_n(X) \otimes \mathbb{Z}_p$ by the Hurewicz theorem, which holds for connective spectra. After converting the map $X \to L_0$ into an inclusion, the cofiber $Z_1 = L_0/X$ then has $\pi_i(Z_1) = 0$ for $i \le n$ and $\pi_{n+1}(Z_1)$ is the kernel of the map $\pi_n(X) \to \pi_n(L_0)$, which has smaller order than $\pi_n(X)$. Now we repeat the process with Z_1 in place of X to construct a map $Z_1 \to L_1$ inducing the map $\pi_{n+1}(Z_1) \to \pi_{n+1}(Z_1) \otimes \mathbb{Z}_p$ on π_{n+1}, so the cofiber $Z_2 = L_1/Z_1$ has its first nontrivial homotopy group $\pi_{n+2}(Z_2)$ of smaller order than $\pi_{n+1}(Z_1)$. After finitely many steps we obtain Z_{n+k} with $\pi_{n+k}(Z_{n+k}) = 0$ as well as all the lower homotopy groups. At this point we switch our attention to $\pi_{n+k+1}(Z_{n+k})$ and repeat the steps again. This infinite process yields the complex $X \to L_0 \to L_1 \to \cdots$.

It is easier to describe what is happening in this complex if we look at the associated tower $\cdots \to Z^2 \to Z^1 \to X$ where $Z^k = \Sigma^{-k} Z_k$. Here the first map $Z^1 \to X$ induces an isomorphism on all homotopy groups except π_n, where it induces an inclusion of a proper subgroup. The same is true for the next map $Z^2 \to Z^1$, and after finitely

many steps this descending chain of subgroups $\pi_n(Z^k)$ becomes zero and we move on to $\pi_{n+1}(X)$, eventually reducing this to zero, and so on up the tower, killing each $\pi_i(X)$ in turn. Thus for each i the groups $\pi_i(Z^k)$ are zero for all sufficiently large k. The same is true for the groups $\pi_i^Y(Z^k)$ when Y is a finite spectrum, since a map $\Sigma^i Y \to Z^k$ can be homotoped to a constant map one cell at a time if all the groups $\pi_j(Z^k)$ vanish for j less than or equal to the largest dimension of the cells of $\Sigma^i Y$.

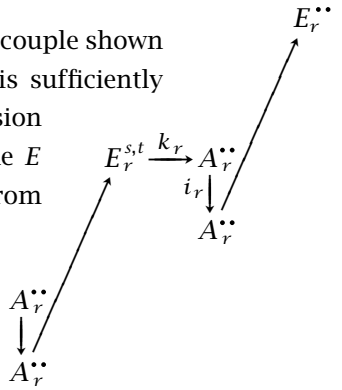
By the lemma the complex used to define the spectral sequence maps to the complex we have just constructed. This is equivalent to a map of towers, inducing a commutative diagram

$$\begin{array}{ccccccc}
\cdots & \longrightarrow & \pi_i^Y(X^2) & \longrightarrow & \pi_i^Y(X^1) & \longrightarrow & \pi_i^Y(X) \\
 & & \downarrow & & \downarrow & & \| \\
\cdots & \longrightarrow & \pi_i^Y(Z^2) & \longrightarrow & \pi_i^Y(Z^1) & \longrightarrow & \pi_i^Y(X)
\end{array}$$

If an element of $\pi_i^Y(X)$ pulled back arbitrarily far in the first row, it would also pull back arbitrarily far in the second row, but we have just seen this is impossible. Hence $\bigcap_n F^{s+n,t+n}$ is empty, which proves (b) in the special case that $\pi_*(X)$ is all p-torsion.

In the general case let α be an element of $\pi_n^Y(X)$ whose order is either infinite or a power of p. Then there is a positive integer k such that α is not divisible by p^k, meaning that α is not p^k times any element of $\pi_n^Y(X)$. Consider the map $X \xrightarrow{p^k} X$ obtained by adding the identity map of X to itself p^k times using the abelian group structure in $[X,X]$. This map fits into a cofibration $X \xrightarrow{p^k} X \longrightarrow Z$ inducing a long exact sequence $\cdots \to \pi_i(X) \xrightarrow{p^k} \pi_i(X) \to \pi_i(Z) \to \cdots$ where the map p^k is multiplication by p^k. From exactness it follows that $\pi_*(Z)$ consists entirely of p-torsion. By the lemma the map $X \to Z$ induces a map from the given Adams tower on X to a chosen Adams tower on Z. The map $\pi_n^Y(X) \to \pi_n^Y(Z)$ sends α to a nontrivial element $\beta \in \pi_n^Y(Z)$ by our choice of α and k, using exactness of $\pi_n^Y(X) \xrightarrow{p^k} \pi_n^Y(X) \to \pi_n^Y(Z)$. If α pulled back arbitrarily far in the tower on X then β would pull back arbitrarily far in the tower on Z. This is impossible by the special case already proved. Hence (b) holds in general.

To prove (a) consider the portion of the r^{th} derived couple shown in the diagram at the right. We claim first that if r is sufficiently large then the vertical map i_r is injective. For nontorsion and non-p-torsion this follows from exactness since the E columns are $\mathbb{Z}_p$ vector spaces. For p-torsion it follows from part (b) that a term $A_r^{s,t}$ contains no p-torsion if r is sufficiently large since $A_r^{s,t}$ consists of the elements of $A_1^{s,t}$ that pull back $r-1$ units vertically.



Since i_r is injective for large r, the preceding map k_r is zero, so the differential d_r starting at $E_r^{s,t}$ is zero for large r. The differential d_r mapping to $E_r^{s,t}$ is also zero for large r since it

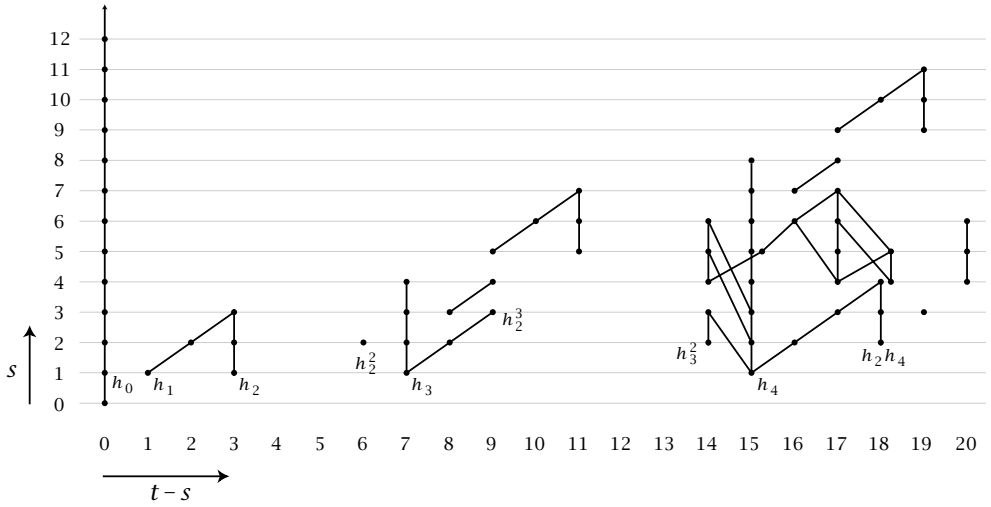
originates at a zero group, as all the terms in each E column of the initial staircase diagram are zero below some point. Thus $E_r^{s,t} = E_{r+1}^{s,t}$ for r sufficiently large.

Since the map k_r starting at $E_r^{s,t}$ is zero for large r, exactness implies that for large r the group $E_r^{s,t}$ is the cokernel of the vertical map in the lower left corner of the diagram. This vertical map is just the inclusion $F^{s+1,t+1} \hookrightarrow F^{s,t}$ when r is large, so the proof of (a) is finished. □

Stable Homotopy Groups of Spheres

For a first application of the Adams spectral sequence let us consider the special case that was one of the primary motivations for its construction, the problem of computing stable homotopy groups of spheres. Thus we take X and Y both to be S^0, in the notation of the preceding section. We will focus on the prime $p = 2$, but we will also take a look at the $p = 3$ case as a sample of what happens for odd primes.

Fixing p to be 2, here is a picture of an initial portion of the E^2 page of the spectral sequence (the musical score to the harmony of the spheres?):



The horizontal coordinate is $t - s$ so the i^{th} column is giving information about π_i^s. Each dot represents a $\mathbb{Z}_2$ summand in the E^2 page, so in this portion of the page there are only two positions with more than one summand, the $(15, 5)$ and $(18, 4)$ positions. Referring back to the staircase diagram, we see that the differential d_r goes one unit to the left and r units upward. The nonzero differentials are drawn as lines sloping upward to the left. For $t - s \leq 20$ there are thus only six nonzero differentials, but if the diagram were extended farther to the right one would see many more nonzero differentials, quite a jungle of them in fact.

For example, in the $t - s = 15$ column we see six dots that survive to E^∞, which says that the 2 torsion in π_{15}^s has order 2^6. In fact it is $\mathbb{Z}_{32} \times \mathbb{Z}_2$, and this information about extensions can be read off from the vertical line segments which indicate multi-

plication by 2 in π_*^S. So the fact that this column has a string of five dots that survive to E^∞ and are connected by vertical segments means that there is a $\mathbb{Z}_{32}$ summand of π_{15}^S, and the other $\mathbb{Z}_2$ summand comes from the remaining dot in this column. In the $t-s=0$ column there is an infinite string of connected dots, corresponding to the fact that $\pi_0^S=\mathbb{Z}$, so iterated multiplication of a generator by 2 never gives zero. The individual dots in this column are the successive quotients $2^n\mathbb{Z}/2^{n+1}\mathbb{Z}$ in the filtration of $\mathbb{Z}$ by the subgroups $2^n\mathbb{Z}$.

The line segments sloping upward to the right indicate multiplication by the element h_1 in the $(1,1)$ position of the diagram. We have drawn them mainly as a visual aid to help tie together some of the dots into recognizable patterns. There is in fact a graded multiplication in each page of the spectral sequence that corresponds to the composition product in π_*^S. (This is formally like the multiplication in the Serre spectral sequence for cohomology.) For example in the $t-s=3$ column we can read off the relation $h_1^3=4h_2$. To keep the diagram uncluttered we have not used line segments to denote any other nonzero products, such as multiplication by h_2, which is nonzero in a number of cases.

The $s=1$ row of the E^2 page consists of just the elements h_i in the position $(2^i-1,1)$. These are related to the Hopf invariant, and in particular h_1, h_2, and h_3 correspond to the classical Hopf maps. The next one, h_4 does not survive to E^∞, and in fact the differential $d_2h_4=h_0h_3^2$ is the first nonzero differential in the spectral sequence. It is easy to see why this differential must be nonzero: The element of π_{14}^S corresponding to h_3^2 must have order 2 by the commutativity property of the composition product, since h_3 has odd degree, and there is no other term in the E^2 page except h_4 that could kill $h_0h_3^2$. No h_i for $i>4$ survives to E^∞ either, but this is a harder theorem, equivalent to Adams' theorem on the nonexistence of elements of Hopf invariant one.

There are only a few differentials to the left of the $t-s=14$ column that could be nonzero since d_r goes r units upward and $r\geq 2$. It is easy to use the derivation property $d(xy)=x(dy)+(dx)y$ to see that these differentials must vanish. For the element h_1, if we had $d_rh_1=h_0^{r+1}$ then we would have $d(h_0h_1)=h_0^{r+2}$ nonzero as well, but $h_0h_1=0$. The only other differential which could be nonzero is d_2 on the element h_1h_3 in the $t-s=8$ column, but d_2h_1 and d_2h_3 both vanish so $d_2(h_1h_3)=0$.

Computing the E^2 page of the spectral sequence is a mechanical process, as we will see, although its complexity increases rapidly as $t-s$ increases, so that even with computer assistance the calculations that have been made only extend to values of $t-s$ on the order of 100. Computing differentials is much harder, and not a purely mechanical process, and the known calculations only go up to $t-s$ around 60.

Let us first show that for computing $\mathrm{Ext}_{\mathcal{A}}^{s,t}(H^*(X),\mathbb{Z}_p)$ it suffices just to construct

a *minimal* free resolution of $H^*(X)$, that is, a free resolution

$$\cdots \longrightarrow F_2 \xrightarrow{\varphi_2} F_1 \xrightarrow{\varphi_1} F_0 \xrightarrow{\varphi_0} H^*(X) \longrightarrow 0$$

where at each step of the inductive construction of the resolution we choose the minimum number of free generators for F_i in each degree.

Lemma 2.8. *For a minimal free resolution, all the boundary maps in the dual complex*

$$\cdots \longleftarrow \mathrm{Hom}_{\mathcal{A}}(F_2, \mathbb{Z}_p) \longleftarrow \mathrm{Hom}_{\mathcal{A}}(F_1, \mathbb{Z}_p) \longleftarrow \mathrm{Hom}_{\mathcal{A}}(F_0, \mathbb{Z}_p) \longleftarrow 0$$

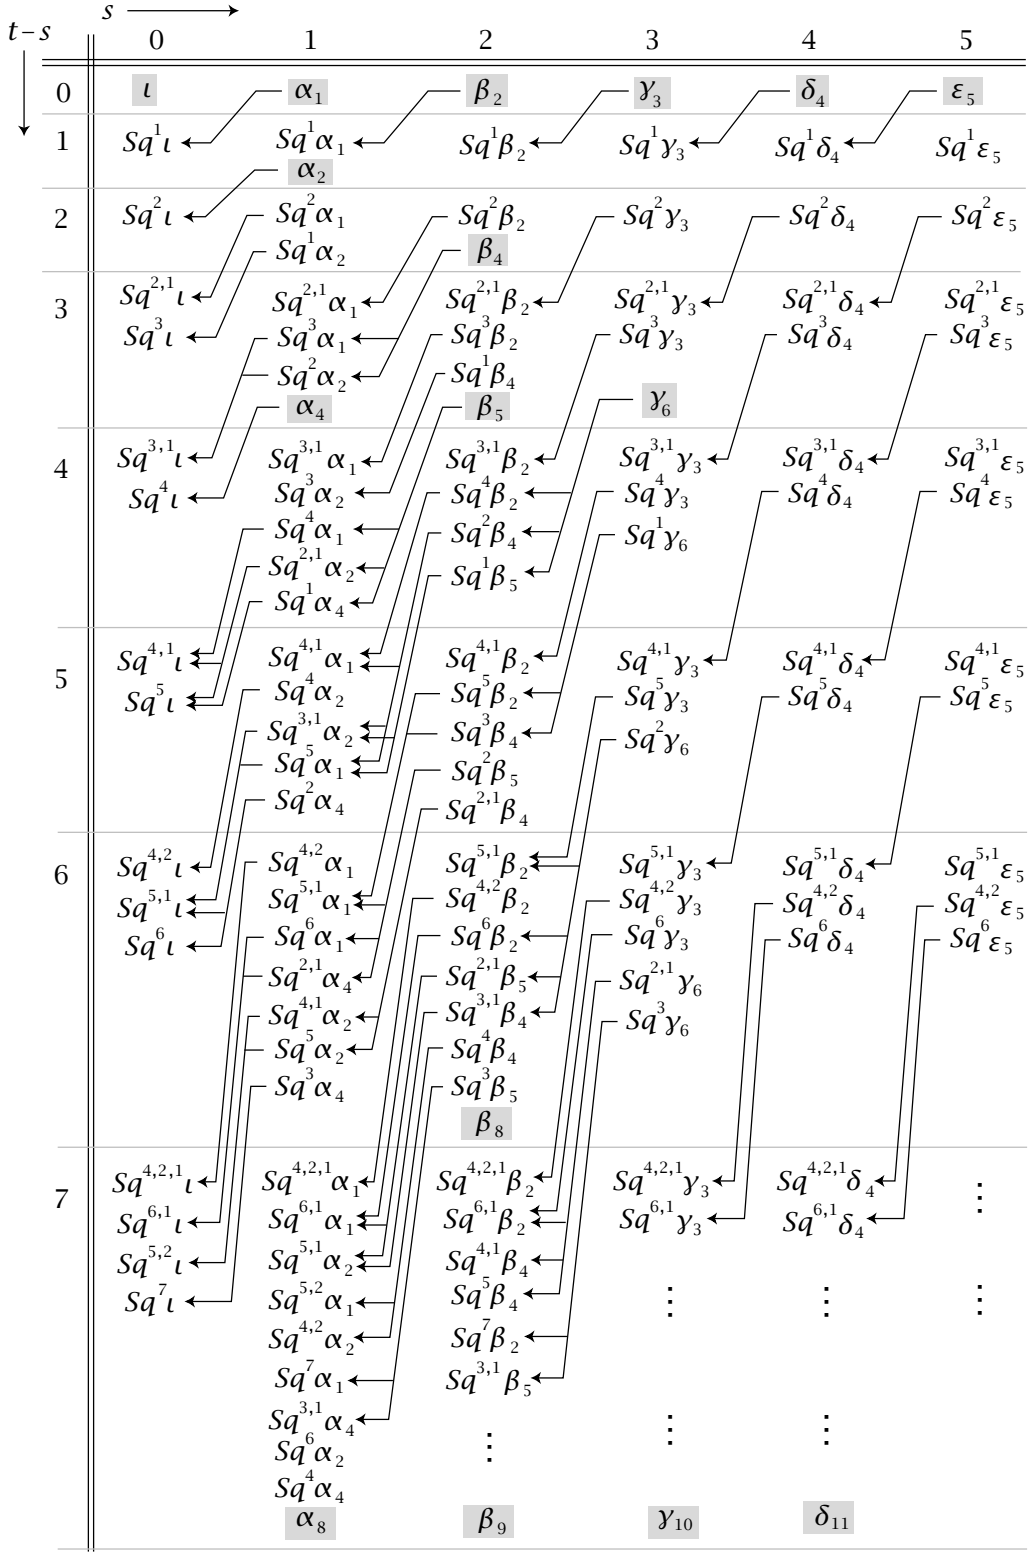
are zero, hence $\mathrm{Ext}^{s,t}_{\mathcal{A}}(H^*(X), \mathbb{Z}_p) = \mathrm{Hom}^t_{\mathcal{A}}(F_s, \mathbb{Z}_p)$.

Proof: Let $\mathcal{A}^+$ be the ideal in $\mathcal{A}$ consisting of all elements of strictly positive degree, or in other words the kernel of the augmentation map $\mathcal{A} \to \mathbb{Z}_p$ given by projection onto the degree zero part $\mathcal{A}^0$ of $\mathcal{A}$. Observe that $\mathrm{Ker}\,\varphi_i \subset \mathcal{A}^+ F_i$ since if we express an element $x \in \mathrm{Ker}\,\varphi_i$ in terms of a chosen basis for F_i as $x = \sum_j a_j x_{ij}$ with $a_j \in \mathcal{A}$, then if x is not in $\mathcal{A}^+ F_i$, some a_j is a nonzero element of $\mathcal{A}^0 = \mathbb{Z}_p$ and we can solve the equation $0 = \varphi_i(x) = \sum_j a_j \varphi_i(x_{ij})$ for $\varphi_i(x_{ij})$, which says that the generator x_{ij} was superfluous.

Since $\varphi_{i-1}\varphi_i = 0$, we have $\varphi_i(x) \in \mathrm{Ker}\,\varphi_{i-1}$ for each $x \in F_i$, so from the preceding paragraph we obtain a formula $\varphi_i(x) = \sum_j a_j x_{i-1,j}$ with $a_j \in \mathcal{A}^+$. Hence for each $f \in \mathrm{Hom}_{\mathcal{A}}(F_{i-1}, \mathbb{Z}_p)$ we have $\varphi_i^*(f(x)) = f\varphi_i(x) = \sum_j a_j f(x_{i-1,j}) = 0$ since $a_j \in \mathcal{A}^+$ and $f(x_{i-1,j})$ lies in $\mathbb{Z}_p$ which has a trivial $\mathcal{A}$-module structure. □

Let us describe how to compute $\mathrm{Ext}^{s,t}_{\mathcal{A}}(\mathbb{Z}_2, \mathbb{Z}_2)$ by constructing a minimal resolution of $\mathbb{Z}_2$ as an $\mathcal{A}$-module. An initial portion of the resolution is shown in the chart on the next page. For the first stage of the resolution $F_0 \to \mathbb{Z}_2$ we must take F_0 to be a copy of $\mathcal{A}$ with a generator ι in degree 0 mapping to the generator of $\mathbb{Z}_2$. This copy of $\mathcal{A}$ forms the first column of the table, which consists of the elements $Sq^I\iota$ as Sq^I ranges over the admissible monomials in $\mathcal{A}$. The kernel of the map $F_0 \to \mathbb{Z}_2$ consists of everything in the first column except ι, so we want the second column, which represents F_1, to map onto everything in the first column except ι. To start, we need an element α_1 at the top of the second column mapping to $Sq^1\iota$. (We will use subscripts to denote the degree t, so α_i will have degree $t = i$, and similarly for the later generators $\beta_i, \gamma_i, \cdots$.) Once we have α_1 in the second column, we also have all the terms $Sq^I\alpha_1$ for admissible I lower down in this column. To see what else we need in the second column we need to compute how the terms in the second column map to the first column. Since α_1 is sent to $Sq^1\iota$, we know that $Sq^I\alpha_1$ is sent to $Sq^I Sq^1\iota$. The product $Sq^I Sq^1$ will be admissible unless I ends in 1, in which case $Sq^I Sq^1$ will be 0 because of the Adem relation $Sq^1Sq^1 = 0$. In particular, $Sq^1\alpha_1$ maps to 0. This means we have to introduce a new generator α_2 to map to $Sq^2\iota$. Then $Sq^I\alpha_2$ maps to $Sq^I Sq^2\iota$ and we can use Adem relations to express this in terms of admissibles. For

example $Sq^1\alpha_2$ maps to $Sq^1Sq^2\iota = Sq^3\iota$ and $Sq^2\alpha_2$ maps to $Sq^2Sq^2\iota = Sq^3Sq^1\iota$.



Some of the simpler Adem relations, enough to do the calculations shown in the chart, are listed in the following chart.

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

Note that the relations for Sq^3Sq^i follow from the relations for Sq^2Sq^i and Sq^1Sq^i since $Sq^3 = Sq^1Sq^2$.

Moving down the $s = 1$ column we see that we need a new generator α_4 to map to $Sq^4\iota$. In fact it is easy to see that the only generators we need in the second column are α_{2^n}'s mapping to $Sq^{2^n}\iota$. This is because Sq^i is indecomposable iff $i = 2^n$, which implies inductively that every $Sq^I\iota$ except $Sq^{2^n}\iota$ will be hit by previously introduced terms, while $Sq^{2^n}\iota$ will not be hit.

Now we start to work our way down the third column, introducing the minimum number of generators necessary to map onto the kernel of the map from the second column to the first column. Thus, near the top of this column we need β_2 mapping to $Sq^1\alpha_1$, β_4 mapping to $Sq^3\alpha_1 + Sq^2\alpha_2$, and β_5 mapping to $Sq^4\alpha_1 + Sq^2Sq^1\alpha_2 + Sq^1\alpha_4$. One can see that things are starting to get more complicated here, and it is not easy to predict where new generators will be needed.

Subsequent columns are computed in the same way. Near the top, the structure of the columns soon stabilizes, each column looking just the same as the one before. This is fortunate since it is the rows, with $t - s$ constant, that we are interested in for computing π^s_{t-s}. The most obvious way to proceed inductively would be to compute each diagonal with t constant by induction on t, moving up the diagonal from left to right. However, this would require infinitely many computations to determine a whole row. To avoid this problem we can instead proceed row by row, moving across each row from left to right assuming that higher rows have already been computed. To determine whether a new generator is needed in the $(s, t-s)$ position we need to see whether the map from the $(s-1, t-s+1)$ position to the $(s-2, t-s+2)$ position is injective. These two positions are below the row we are working on, so we do not yet know whether any new generators are required in these positions, but if they are, they will have no effect on the kernel we are interested in since minimality implies that new generators always generate a subgroup that maps injectively. Thus we have enough information to decide whether new generators are needed in the $(s, t-s)$ position, and so the induction can continue.

The chart shows the result of carrying out the row-by-row calculation through the row $t - s = 5$. As it happens, no new generators are needed in this row or the

preceding one. In the next row $t-s=6$ one new generator β_8 will be needed, but the chart does not show the computations needed to justify this. And in the $t-s=7$ row four new generators α_8, β_9, γ_{10}, and δ_{11} will be needed. The reader is encouraged to do some of these calculations to get a real feeling for what is involved. Most of the work involves applying Adem relations, and then when the maps have been computed, their kernels need to be determined.