

Spectral Sequences from Sequences of Spectra: Towards the Spectrum of the Category of Spectra

Dylan Wilson

March 23, 2013

1 The Adams Spectral Sequences

As is well known, it is our manifest destiny as 21st century algebraic topologists to compute homotopy groups of spheres.

This noble venture began even before the notion of homotopy was around. In 1931, Hopf¹ was thinking about a map he had encountered in geometry from S^3 to S^2 and wondered whether or not it was essential. He proved that it was by considering the linking of the fibers. After Hurewicz developed the notion of higher homotopy groups this gave the first example, aside from the self-maps of spheres, of a non-trivial higher homotopy group. Hopf classified maps from S^3 to S^2 and found they were given in a manner similar to degree, generated by the Hopf map, so that

$$\pi_3(S^2) = \mathbb{Z}$$

In modern-day language we would prove the nontriviality of the Hopf map by the following argument.

Consider the cofiber of the map $S^3 \rightarrow S^2$. By construction this is $\mathbb{C}P^2$. If the map were nullhomotopic then the cofiber would be homotopy equivalent to a wedge $S^2 \vee S^3$. But the cup-square of the generator in $H^2(\mathbb{C}P^2)$ is the generator of $H^4(\mathbb{C}P^2)$, so this can't happen.

This gives us a general procedure for constructing essential maps $\phi : S^{2n-1} \rightarrow S^n$. Cook up fancy CW-complexes built of two cells, one in dimension n and another in dimension $2n$, and show that the square of the bottom generator is the top generator. Working over $\mathbb{Z}/2$ we note that, if α is the bottom generator and β the top generator, then $\alpha^2 = h(\phi)\beta$ where $h(\phi) \in \mathbb{Z}/2$. This is called the *Hopf invariant*.

The big idea is that this is really a question about Steenrod squares: the squaring operation is a special type of Steenrod square, and so we're really asking for

$$Sq^n \alpha = \beta$$

Since the cohomology of this complex is concentrated only in degrees n and $2n$, the first obvious obstruction is decomposability of the Steenrod square: if the operation factors as $Sq^n = Sq^i Sq^j$ with $i, j > 0$ then $Sq^n \alpha = 0$ since $Sq^j \alpha$ is already zero.

Another way of saying what we've done is to notice that for any complex, X , $H^*(X)$ has the structure of a module over the Steenrod algebra. In our case we know that $H^*(X)$, as a vector space, just looks like $\Sigma^n \mathbb{F}_2 \oplus \Sigma^{2n} \mathbb{F}_2$. By considering the cofiber sequence


$$S^n \longrightarrow X \longrightarrow S^{2n}$$

and applying cohomology we see that we're looking for extensions

$$0 \longrightarrow \Sigma^{2n} \mathbb{F}_2 \longrightarrow H^*(X) \longrightarrow \Sigma^n \mathbb{F}_2 \longrightarrow 0$$

¹This is an unrelated story but I couldn't resist telling it: As a young student, Freudenthal went to Hopf and told him, "I would like you to be my advisor, here is my thesis," and promptly handed him a fully realized dissertation on non-compact topological groups.

and we want them to be nontrivial. From this point of view it's clear that the only candidates look like



Of course, not all of these candidates are realized in homotopy theory. The reason is that there are certain *secondary* cohomology operations that exist in cases when the Steenrod squares vanish. It turns out that, when $j \geq 4$, the Steenrod squares Sq^{2^j} are decomposable in terms of secondary cohomology operations, and Adams used this fact to prove that the only Hopf invariant one maps are the ones constructed from \mathbb{R} , \mathbb{C} , \mathbb{H} , and \mathbb{O} .

Roughly speaking, Adams created a filtration on the homotopy groups of spheres by declaring that an element is in filtration degree s if it cannot be detected by cohomology operations of the r th kind for $r < s$. This filtration leads to a spectral sequence beginning in algebra and ending in the 2-component of the stable homotopy groups of spheres:

$$\text{Ext}_{\mathcal{A}}^s(\Sigma^t \mathbb{F}_2, \mathbb{F}_2) \Rightarrow {}_2\pi_{t-s}(S)$$

This is called the Adams spectral sequence, or as Adams liked to call it, the ASS. The goal of the first half of this talk is to give some of the gory details of its construction and see why it's not so bad, and then give a generalization called the Adams-Novikov spectral sequence.

1.1 Murder and Mayhem

Scientists and philosophers used to argue that what separates humans from other animals is our ability to create tools. This is clearly false: Macaques in Thailand have been recorded using hair as floss, and octopuses often use coconut shells as battle armor. No; what separates humans from animals is our ability to create spectral sequences. To date, no other species on Earth has been observed using spectral sequences.

In what is possibly the greatest thesis of all time², J.P. Serre showed the world just how great these gadgets can be. It's important to realize that the state of affairs prior to Serre's thesis was very primitive: we knew the 0-stem and the 1-stem of the stable homotopy groups of spheres, and that was about it. His work was the equivalent of introducing guns to a society of Neanderthals.

Here is how Serre used the Leray-Serre spectral sequence to compute homotopy groups, in the simplest example. He uses the fact that, up to that point, the only reliable way to compute homotopy groups was to compute homology groups and then apply the Hurewicz theorem. The trouble is that the Hurewicz theorem is no good after the first nonzero homotopy group. So the idea is to kill this homotopy group and proceed from there.

Working mod 2, we know that $\pi_n S^n$ looks like $\mathbb{Z}/2$ and this is detected in cohomology, so we have a map

$$S^n \rightarrow K(\mathbb{Z}/2, n)$$

inducing an isomorphism on π_k for $k \leq n$. This means that the homotopy fiber, F , has the property that $\pi_k F = 0$ for $k \leq n$ and $\pi_{n+1} F \cong \pi_{n+1}(S^n)$ since Eilenberg-MacLane spaces have no higher homotopy groups. So now we just need to compute the $(n+1)$ st homology group of F and we'll be done by Hurewicz. Well, in this case we can recover from the spectral sequence an exact sequence

$$0 \longrightarrow H_{n+2}(K(\mathbb{Z}/2), n) \longrightarrow H_{n+1}(F) \longrightarrow 0$$

(for $n > 2$) and so we are reduced to computing $H^{n+2}(K(\mathbb{Z}/2))$, which is do-able and we see it's just a $\mathbb{Z}/2$, which is what we expect.

From here, one would kill off the lowest homotopy group of F via a map to an Eilenberg-MacLane space, and then compute with the spectral sequence to see what happens, etc. etc.

²Tied with René Thom's thesis

The thing that makes this process annoying is that this process is really only streamlined when we're killing off one copy of $\mathbb{Z}/2$ at a time, or at least that's what Serre did at first.

Adams, while thinking about the Hopf invariant one problem as we discussed in the introduction, decided to formalize this process as follows. Working in the stable setting, let X be a finite spectrum and suppose we want to compute its homotopy groups. Then X has some cohomology, and we can choose generators which are represented by maps $X \rightarrow \Sigma^k H\mathbb{F}_p$ into shifts of Eilenberg-MacLane spectra. We can collect these together into a map

$$X \xrightarrow{f_0} K_0$$

where K_0 is a locally finite wedge of suspensions of Eilenberg-MacLane spectra (sometimes called *generalized Eilenberg-MacLane spectra*). Then we can consider the homotopy fiber, X_1 . This captures all the information about X that cohomology couldn't see right off the bat. But now we can iterate this procedure: cohomology *can* see something about X_1 , and so we can get a map

$$X_1 \rightarrow K_1$$

giving a surjective map on cohomology, where K_1 is a locally finite wedge of suspensions of Eilenberg-MacLane spaces. Continuing we get a diagram

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X_s & \longrightarrow & \cdots & \longrightarrow & X_2 & \longrightarrow & X_1 & \longrightarrow & X_0 = X \\ & & \downarrow & & & & \downarrow & & \downarrow & & \downarrow \\ & & K_s & & & & K_2 & & K_1 & & K_0 \end{array}$$

Now, applying π_* we get an (unrolled) exact couple

$$\begin{array}{ccccccccccccccc} \cdots & \longrightarrow & \pi_* X_{s+2} & \longrightarrow & \pi_* X_{s+1} & \longrightarrow & \pi_* X_s & \longrightarrow & \cdots & \longrightarrow & \pi_* X_2 & \longrightarrow & \pi_* X_1 & \longrightarrow & \pi_* X_0 = X \\ & & \downarrow & & \downarrow & & \downarrow & & & & \downarrow & & \downarrow & & \downarrow \\ & & \pi_* K_{s+2} & & \pi_* K_{s+1} & & \pi_* K_s & & & & \pi_* K_2 & & \pi_* K_1 & & \pi_* K_0 \end{array}$$

(Note: Dashed arrows in the original image indicate a more complex relationship between the $\pi_* X$ and $\pi_* K$ terms, but the above represents the main structure.)

The game is that we totally understand $\pi_* K_s$, and so we want to figure out which elements come from $\pi_* X_s$ so that we can lift and push them all the way to an element in $\pi_* X$. Coming from $\pi_* X_s$ is the same as vanishing along the map

$$\pi_* K_s \rightarrow \pi_* X_{s+1}$$

by exactness. We can collect evidence for this by pushing down to $\pi_* K_{s+1}$ and checking if we get zero. If we do then wherever our potential element landed in $\pi_* X_{s+1}$ had to come from $\pi_* X_{s+2}$ and then we continue up along the tower.

If we always get the answer zero, then our element mapping to $\pi_* X_{s+1}$ factors through the map from the inverse limit $\lim \pi_* X_s$. In good cases this is zero, and then we're good. Elements that climb all the way up the tower like this are called *permanent cycles*.

Of course, the obvious way to get elements coming from $\pi_* X_s$ is to use elements of $\pi_* K_{s-1}$ and push up, or they could come from even further down. This is redundant information, so we'd like to mod out by these guys, which we call *boundaries* or *eventual boundaries*.

At the end of the day we'll have collected the data of

$$\frac{\text{permanent cycles}}{\text{eventual boundaries}}$$

and this should be isomorphic to the associated graded object coming from the filtration of $\pi_* X$ (or its p -component).

Now, applying H^* to the original diagram we get a resolution of $H^* X$ by free modules over the Steenrod algebra; this is supposed to convince us that it's plausible to believe the following theorem:

Theorem 1.1.1 (Adams). *Let X be a finite spectrum. Then there is a spectral sequence converging to the p -component of $\pi_* X$, given by*

$$E_2^{s,t} = \text{Ext}_{\mathcal{A}}^s(\Sigma^t H^* X, \mathbb{F}_p) \Rightarrow \pi_{t-s} X$$

where \mathcal{A} is the mod p Steenrod algebra.

Our goal for the remainder of this half is to convince you that such a spectral sequence is a natural idea to invent, and that one can do it in a very clean way in much more generality leading to spectral sequences of this form associated to any reasonable cohomology theory, natural both in the choice of X and the choice of cohomology theory.

1.2 Brita and Overpriced Airlines, or Where Spectral Sequences Come From

For this talk we will take the point of view that spectral sequences come from filtered gadgets, and that's how we should think of them morally.³ Everything I'm about to say works for anything that feels like the category of spectra (nice stable model categories, stable ∞ -categories, derived categories of sheaves, etc.)

A filtered object is just a \mathbb{Z} -shaped diagram, X ,

$$\cdots \longrightarrow X(-1) \xrightarrow{f^0} X(0) \xrightarrow{f^1} X(1) \longrightarrow \cdots$$

We're going to do things dual to the way you usually see them, since that's what we'll need for later. That is, instead of looking at cofibers and trying to produce maps to the homotopy colimit of this diagram, we'll look at fibers and try to get at the homotopy limit of this diagram.

Define an exact couple

$$\begin{array}{ccc} & \pi_* \text{fib}(f^*) & \\ & \swarrow \quad \searrow & \\ \pi_* X(*) & \xrightarrow{\quad} & \pi_* X(*) \end{array}$$

where the maps are...

$$\begin{aligned} \pi_{p+q} \text{fib}(f^p) &\rightarrow \pi_{p+q} X(p-1) \\ \pi_{p+q} X(p) &\rightarrow \pi_{p+q} X(p+1) \\ \pi_{p+q} X(p) &\rightarrow \pi_{p+q-1} \text{fib}(f^p) \end{aligned}$$

all coming from the various long exact sequences associated to the (co)fiber sequences.

The differentials are then given in the usual way by taking derived exact couples. The theorem is then...

Theorem 1.2.1. *In nice cases (for example $X(n) = 0$ for $n \gg 0$ and certain Mittag-Leffler conditions are satisfied), there is a convergent spectral sequence*

$$E_1^{p,q} = \pi_{p+q} \text{fib}(f^p) \Rightarrow \pi_{p+q}(\lim X)$$

And d_1 is given by the composition

$$\pi_{p+q} \text{fib}(f^p) \rightarrow \pi_{p+q} X(p-1) \rightarrow \pi_{p+q-1} \text{fib}(f^{p-1})$$

Remark 1.2.2. The dual statement, where we are taking cofibers and trying to get at the homotopy colimit, works under fewer hypothesis since the colimit functor is exact on abelian groups. In the general setting of a stable ∞ -category one has to figure out how limits or colimits play with the t -structure, and so the asymmetry for the category of spectra comes from the fact that the heart of this category is the category of abelian groups which plays nice with colimits but not limits.

³In other talks we might want to emphasize the Grothendieck spectral sequence and say that spectral sequences come from deriving in the wrong order.

So at this point we know just what to do with filtered gadgets, but we may ask ourselves: if spectral sequences come from filtered things, where do filtered things come from? Unlike many questions of this form that preschoolers ask, this one has a good answer: A lot of them come from (co)simplicial things.

Here's the idea: given a simplicial object, I can filter its geometric realization by its skeleta. Or, closer to what we'll be doing, given a cosimplicial object I can filter its totalization by coskeleta. This is not mysterious at all: cosimplicial objects are functors

$$\Delta \rightarrow \mathcal{C}$$

and the category Δ has an obvious filtration

$$\Delta_{\leq 0} \rightarrow \Delta_{\leq 1} \rightarrow \cdots \rightarrow \Delta_{\leq n} \rightarrow \cdots \rightarrow \Delta$$

where $\Delta_{\leq n}$ is the subcategory generated by $[0], \dots, [n]$. Given a cosimplicial object $X : \Delta \rightarrow \mathcal{C}$ we may form its n th coskeleton (another cosimplicial object) by restricting to $\Delta_{\leq n}$ and then taking the Kan extension. Explicitly, $\text{cosk}_n X : \Delta \rightarrow \mathcal{C}$ is given by

$$(\text{cosk}_n X)(k) = \lim_{\Delta_{\leq n}^{k/}} X(j)$$

where $\Delta_{\leq n}^{k/}$ denotes the category of objects $[j] \in \Delta_{\leq n}$ together with maps $k \rightarrow j$ (in Δ). We have maps

$$\text{cosk}_{n+1} X \rightarrow \text{cosk}_n X$$

and these fit into a filtration

$$\cdots \longrightarrow \text{cosk}_2 X \longrightarrow \text{cosk}_1 X \longrightarrow \text{cosk}_0 X \longrightarrow 0$$

Taking totalizations we get a filtered diagram in \mathcal{C}

$$\cdots \longrightarrow D(-2) \longrightarrow D(-1) \longrightarrow D(0) \longrightarrow 0$$

and the (ho)limit of this diagram is the totalization of X . This described a functor

$$\text{Fun}(\Delta, \mathcal{C}) \rightarrow \text{Fun}(\mathbb{Z}_{\leq 0}, \mathcal{C})$$

It actually turns out this is an equivalence (when \mathcal{C} is a stable ∞ -category, etc.), though we won't need this here.

Applying our theorem we get a spectral sequence, but in this case we can identify the E_2 term a little differently. There is an obvious way to associate a cochain complex to a cosimplicial object, just by taking the alternating sum of the coface maps. The claim is that the cohomology of this cochain complex can be canonically identified with the E_2 -term of our spectral sequence.

Proposition 1.2.3. *With notation as above, we have that*

$$E_2^{p,q} = H^p(\pi_q(D(*)))$$

in the spectral sequence associated to $D()$.*

Proof. This is rather technical, see ([7], Rmk 1.2.4.4) for details on the dual statement about simplicial objects, which implies this one since it is proven for all stable ∞ -categories. \square

It follows that...

Theorem 1.2.4. *There is a functorial assignment of a spectral sequence to each cosimplicial object of \mathcal{C} converging (in good cases) to its homotopy limit (totalization) with a canonical identification of the E_2 term as given in the previous proposition.*

Remark 1.2.5. Since people don't seem to like negative numbers, this spectral sequence is usually re-indexed so that we have

$$E_2^{s,t} \Rightarrow \pi_{t-s} \text{Tot}(X)$$

with $s, t \geq 0$.

1.3 Adams-Novikov Spectral Sequence and Examples

We are finally ready to construct the Adams-Novikov spectral sequence. The idea is easily motivated by geometry. If we are given a map of spaces, $E \rightarrow B$, we can try to study cohomology of sheaves on B by descending information about cohomology of sheaves on E . The way we do this is to construct a simplicial space with n simplices given by $E^{\times_B^{n+1}}$, the $(n+1)$ -fold fiber product of E over B . If the map E is nice enough (say, a fibration, or an open cover expressed as such, etc.), then the homotopy colimit (geometric realization) of this simplicial space gives back B . Applying, for example, cochains, we get a cosimplicial cochain complex, and since the category of cochains is stable and nice we can apply the methods of the previous section to get a spectral sequence and hope that it's useful. In the case of an open cover, this recovers the familiar Mayer-Vietoris spectral sequence. A similar construction works for coefficients in any sheaf.

If we think of our spaces as being affine schemes, then the given map is induced by a map of rings

$$S \rightarrow R$$

and the iterated fiber products are just iterated tensors over S , so we get a cosimplicial gadget that has $R^{\otimes_S^{n+1}}$ for its $(n+1)$ -cosimplices. Cohomology of sheaves corresponds to Ext groups of various modules, and the spectral sequence is obtained by tensoring this cosimplicial object with the module in question (in the derived sense), and then applying the methods of the previous section.

In our case, we are working in the category of spectra, which is, in a very precise sense, the category of modules over the sphere spectrum, S . When we ask what this category looks like through the eyes of some other (commutative) ring spectrum, E , we can proceed by considering the map

$$S \rightarrow E$$

and forming the cosimplicial spectrum E^\bullet with $E^n = E \wedge \cdots \wedge E$, where this has $(n+1)$ terms. If we want to study the homotopy groups of some S -module, then we can smash this cosimplicial object with the spectrum we care about and then study the spectral sequence.

Of course, for arbitrary ring spectra, E , this won't do us any good unless we can solve two problems:

1. Identify the E_2 -term.
2. Describe what the spectral sequence converges to.

By the previous section we know that the E_2 term is given as the cohomology of the cochain complex

$$0 \longrightarrow \pi_*(E \wedge X) \longrightarrow \pi_*(E \wedge E \wedge X) \longrightarrow \pi_*(E \wedge E \wedge E \wedge X) \longrightarrow \cdots$$

The first term is just E_*X . When X is a sphere the second term is E_*E , and in general the terms *want* to look like tensor powers of E_*E over E_* together with E_*X . This is only true, however, with some conditions:

Proposition 1.3.1. *Let E be a ring spectrum such that E_*E is flat over E_* . Then there is a natural isomorphism*

$$\pi_*(E^{\wedge^n} \wedge X) \cong E_*E^{\otimes n} \otimes_{E_*} E_*X$$

Proof. We have a natural map from the right hand side to the left. By the flatness assumption both sides are homology theories, and the result is true by definition when $X = S$ so the result holds by the uniqueness theorem for homology theories. \square

Example 1.3.2. $H\mathbb{Z}/p$, BP , and MU all satisfy this property. On the other hand, $H\mathbb{Z}$, bo , bu , and MSU all fail to satisfy this condition. There's still a spectral sequence, it's just more difficult to get at the E_2 -term. Note, however, that we can always say what the E_2 -term is when $X = S$ is the sphere spectrum.

The other problem is more subtle, we need to identify $\text{holim } E^\bullet$. We know that E -local spectra are closed under homotopy limits, and that $E \wedge X$ is E -local for any X , so we get a map

$$\phi : L_E X \rightarrow \text{holim } E^\bullet$$

At this point people usually say “under nice conditions this map is an isomorphism.” We’ll give those conditions in just a moment after a brief

Definition 1.3.3. let R be a ring. Then the *core* of R is the equalizer

$$cR \rightarrow R \rightrightarrows R \otimes R$$

Bousfield proved that the only possible cores that one can get are

- (I) The rings $\mathbb{Z}[J^{-1}]$, where J is a set of primes.
- (II) The rings \mathbb{Z}/n , $n \geq 2$.
- (III) The rings $\mathbb{Z}[J^{-1}] \times \mathbb{Z}/n$, where every prime factor of n appears in J .
- (IV) The cores of $\mathbb{Z}[J^{-1}] \times \prod_{p \in K} \mathbb{Z}/p^{e(p)}$ where J and K are infinite sets of primes with $K \subset J$ and $e(p)$ is positive.

In the first two cases (which is what most often occurs in nature), we’re set, as the next proposition shows:

Proposition 1.3.4. *[[4]] Let E be connective commutative ring spectrum, and X connective. Then the map ϕ above is an equivalence if one of the following holds:*

1. $c\pi_0 E \cong \mathbb{Z}[J^{-1}]$
2. $c\pi_0 E \cong \mathbb{Z}/n$, $n \geq 2$.

Moreover, in this case, $L_E X \cong L_R X$ where $R = c\pi_0 E$.

In the other two cases of cores the result fails already for $X = H\mathbb{Z}$.

Putting all of this together we have...

Theorem 1.3.5 ([4]). *Let E be a connective, commutative ring spectrum satisfying one of the conditions of Proposition 1.3.4, X a connective spectrum, and $R = c\pi_0 E$. Then there is a strongly convergent, multiplicative spectral sequence*

$$E_r^{s,t} \Rightarrow \pi_{t-s} L_R X$$

which is natural in both X and E , and takes fiber sequences in the X -variable to long exact sequences of spectral sequences. If, moreover, $E_* E$ is flat over E_* , then the E_2 -term is given by

$$E_2^{s,t} = \text{Ext}_{E_* E}^s(\Sigma^t E_*, E_* X)$$

and the multiplication here agrees with the multiplication in the spectral sequence.

In the sequel we will focus mainly on the case when $E = BP$ or $E = H\mathbb{Z}/p$.

Example 1.3.6. In the classical Adams spectral sequence at the prime 2, we’ve already seen that the Hopf invariant one theorem can be cast as a problem about differentials. In [5], Browder recast the Kervaire invariant one problem in terms of differentials in the ASS, and this was how it was solved.

Remark 1.3.7. The ANSS for BP_* is strictly better at detecting homotopy groups in the following sense. The map $BP \rightarrow H\mathbb{F}_p$ induces a map of spectral sequences with the following property: if an element of $\pi_* S$ is detected by the ASS in filtration s , then it will be detected in the ANSS in filtration *at most* s ; if an element of $\pi_* S$ is detected by the ANSS in filtration s then the same element is detected in the ASS in filtration *at least* s . The reason for this is actually not complicated, and follows from the following invariant description of the filtration for the ANSS.

Theorem 1.3.8. *Let E be a spectrum satisfying the assumptions necessary for an ANSS. Then an element $\alpha : S \rightarrow X$ is in filtration $\geq s$ if it can be factored into s maps each of which is null after smashing the target with E .*

2 The Chromatic Spectral Sequence

With the construction of the Adams-Novikov spectral sequence we have efficiently organized the way that algebraic information relates to homotopy-theoretic information, and so, to even get started, we need to at least make sure we understand how to deal with the algebra.

Once again, this is a job for spectral sequences. Hopf algebroids are a lot like groups or groupoids, and the starting point of the ANSS is essentially group cohomology. Now, we don't really know much about computing group cohomology: but we do know how to compute the cohomology of abelian groups! So the game is to try and reduce calculations to this one.

The first instance of this is probably Peter May's work on cohomology of modules over the Steenrod algebra. He filtered \mathcal{A}_* by copowers of the unit coideal (which is just $\mathcal{A}_*/\mathbb{F}_p$) and showed that this gives a spectral sequence whose E_2 -term is an *Ext* group over the associated graded which it's easy to see is primitively generated. He then developed an efficient way of computing cohomology over such things by using the fact that primitively generated Hopf algebras are universal enveloping algebras for the Lie algebra of primitives.

In modern terminology we might phrase what he did in the following manner. The dual Steenrod algebra represents an affine group scheme, namely the one given by automorphisms of the mod 2 additive formal group law. So comodules over this Hopf algebra correspond precisely to quasi-coherent sheaves over some algebro-geometric object called BG . From the grading we get an action of \mathbb{G}_m on BG , and we consider equivariant quasi-coherent sheaves with respect to this grading. Now we notice that BG has a natural filtration given by and this gives the desired spectral sequence.

The reason for saying this in such a fancy way is to zoom in on a powerful computational idea:

We can study comodules over a Hopf algebroid by understanding something about the geometry of the underlying stack.

If none of these words mean anything to you, that's fine. Brian will say something about this later, but for now we'll move on with the goal of studying comodules over the Hopf algebroid (BP_*, BP_*BP) .

2.1 The Spectral Sequence Served 3 Ways

2.1.1 A la Miller-Ravenel-Wilson

We'll start with the way that the spectral sequence was originally created, which is a tad awkward in comparison to the next two methods but has the advantage of being completely algebraic.

We're interested in computing the cohomology of BP_* as a BP_*BP -comodule, we will henceforth denote this cohomology as

$$H^{s,t}(BP_*) = \text{Ext}_{BP_*BP}^s(\Sigma^t BP_*, BP_*)$$

Recall that, as a ring, we have

$$BP_* \cong \mathbb{Z}_{(p)}[v_1, v_2, \dots], \quad |v_i| = 2(p^i - 1).$$

The first thing we might look at is how much information we can get from the additive formal group law alone, which corresponds to looking at

$$BP_* \rightarrow p^{-1}BP_*$$

The cofiber of this map is just $BP_* \otimes \mathbb{Q}/\mathbb{Z} \cong BP_*/p^\infty$ (the easiest way to see this is that BP_* is flat over $\mathbb{Z}_{(p)}$ and we have an exact sequence of $\mathbb{Z}_{(p)}$ -modules

$$0 \rightarrow \mathbb{Z}_{(p)} \rightarrow \mathbb{Q} \rightarrow \mathbb{Q}/\mathbb{Z} = \mathbb{Z}/p^\infty \rightarrow 0$$

where $\mathbb{Z}/p^\infty = \text{colim } \mathbb{Z}/p^k$.) It follows that, in the derived category of comodules, we get

$$\Sigma^{-1}BP_*/p^\infty \rightarrow BP_* \rightarrow 0$$

with cofiber $p^{-1}BP_*$. Now we'd like a map of comodules $BP_*/p^\infty \rightarrow v_1^{-1}BP_*/p^\infty$ to get at the “height 1” information. This would only make sense if $v_1^{-1}BP_*/p^\infty$ has a comodule structure so that the localization map was a map of comodules, and it does:

Lemma 2.1.1. *Let $I_n = (p, \dots, v_{n-1})$. If M is an I_n -torsion BP_*BP -comodule, then $v_n^{-1}M$ admits the unique structure of a BP_*BP -comodule so that the localization map is a map of comodules.*

Proof. See, for example, [6]. □

Now we may iterate this procedure to form the following filtration of BP_* by comodules in the derived category:

$$\dots \longrightarrow \Sigma^{-3}BP_*/(p^\infty, v_1^\infty, v_2^\infty) \longrightarrow \Sigma^{-2}BP_*/(p^\infty, v_1^\infty) \longrightarrow \Sigma^{-1}BP_*/p^\infty \longrightarrow BP_*$$

where the cofibers are given by

$$\Sigma^{-n}v_n^{-1}BP_*/(p^\infty, \dots, v_{n-1}^\infty)$$

Since the homotopy colimit of this sequence is contractible, we get, in a manner similar to the methods of section 1.2, a spectral sequence converging to $H^{*,*}(BP^*)$ with E_1 term given by

$$E_1^{p,q,*} = H^{q,*}(v_p^{-1}BP_*/(p^\infty, \dots, v_{p-1}^\infty)) \Rightarrow H^{p+q,*}(BP_*)$$

with homological differentials

$$d_r : E_r^{p,q} \rightarrow E_r^{p+r, q-r+1}$$

notice that the shift changed an H^{p+q} to an H^q , this is one of the great computational features of the spectral sequence: computing lower cohomology groups is easier, so we can get at information in $H^{p+q}BP_*$ by considering stuff in H^q of some other module.

The big thing that makes this machine work is that the groups in this E_1 term are computable and sparse enough to make most questions about differentials completely formal:

Theorem 2.1.2 (Morava). *The natural map*

$$Ext_{BP_*BP}^{*,*}(BP_*, v_n^{-1}BP_*/(p, v_1, \dots, v_{n-1})) \rightarrow Ext_{K(n)_*K(n)}(K(n)_*, K(n)_*)$$

is an isomorphism. Moreover, the right hand side is a Poincaré duality algebra over $K(n)_$ of formal dimension n^2 .*

Using this result and leveraging up using Bocksteins Miller, Ravenel, and Wilson were able to compute much of the E_1 -term of the CSS and the vanishing result showed that most of the differentials vanished, making computations much easier.

Now, in a bit we'll state a few results proved with this spectral sequence, but first let's look at two other ways of arriving at this spectral sequence that might be more intuitive.

2.1.2 A la homotopy theory

We have constructed some cohomology theories, $E(n)$, and Ravenel has shown that there is an inclusion of Bousfield classes $\langle E(n) \rangle \leq \langle E(n+1) \rangle$ ⁴ This implies that we have a natural transformation of localizations

$$L_{E(n+1)} \rightarrow L_{E(n)}$$

which we may assemble into a tower

$$\dots \longrightarrow L_{E(n)} \longrightarrow \dots \longrightarrow L_{E(2)} \longrightarrow L_{E(1)} \longrightarrow L_{E(0)} \longrightarrow 0$$

⁴The Bousfield class $\langle E \rangle$ can be thought of as the collection of spectra X such that $E \wedge X$ is not contractible. This is analogous to the support of a module.

It is a theorem that the homotopy limit of this tower restricted to finite spectra is the same as p -localization (the chromatic convergence theorem). But now we see a filtration and we want a spectral sequence, and we get one with

$$E_1^{n,k} = \pi_k M_n X \Rightarrow \pi_{k-n} X_{(p)}$$

where X is a finite spectrum and $M_n X$ is the fiber of the map $L_{E(n)} X \rightarrow L_{E(n-1)} X$. These are called the *monochromatic layers*. This is the homotopy theoretic realization of the chromatic spectral sequence. Morava’s change of rings theorem then tells us that the BP_* -based ANSS for $M_n X$ takes the form

$$H_c^s(\mathbb{S}_n, (M_n E_n)_t(X))^{\text{Gal}(\mathbb{F}_p)} \Rightarrow \pi_{t-s}(M_n X)$$

and this is known to collapse when $p \gg n$.

The connection between the two points of view on the CSS so far is that these continuous cohomology group calculations also appear in the calculations for the Adams-Novikov E_2 -term coming from the first version of the CSS.

Example 2.1.3. $X = S$ and $2p - 2 \geq \max(n^2, 2n + 2)$ then the spectral sequence collapses and the groups $H_c^s(\mathbb{S}_n, \pi_t M_n E_n)^{\text{Gal}(\mathbb{F}_p)}$ vanish unless t is divisible by $2(p - 1)$. See [3] for a proof of this fact.

2.1.3 A la geometry

This last method of obtaining the chromatic spectral sequence is sort of a mixture of the first and second methods: it is algebraic in nature, but explicitly uses the height filtration of p -typical formal groups in a way that might be easier to remember. If the word “stack” fills you with fear, think “space”. If that still doesn’t help, skip this section, it’s unnecessary.

Throughout this discussion we are working at a prime, p .

The idea is that the moduli stack of p -typical formal groups, \mathcal{M} , has a filtration by height. That is, we may consider open substacks

$$\mathcal{M}_{<0} \subset \cdots \subset \mathcal{M}_{<n} \subset \mathcal{M}_{<n+1} \subset \cdots \subset \mathcal{M}$$

where $\mathcal{M}_{<n}$ consists of p -typical formal groups of height less than n . Now, this gives us a filtration on the derived category of coherent sheaves on \mathcal{M} . To see this, let $i_n : \mathcal{M}_{<n} \rightarrow \mathcal{M}$ denote the inclusion. Then we have an adjoint pair

$$\mathbf{R}(i_n)_* : D_{\text{coh}}(\mathcal{M}_{<n}) \rightleftarrows D_{\text{coh}}(\mathcal{M}) : i_n^*$$

And so we get natural transformations

$$\cdots \longrightarrow \mathbf{R}(i_n)_* i_n^* \longrightarrow \mathbf{R}(i_{n-1})_* i_{n-1}^* \longrightarrow \cdots \longrightarrow \mathbf{R}(i_1)_* i_1^* \longrightarrow \mathbf{R}(i_0)_* i_0^* \longrightarrow 0$$

This is neat, it means that every object automatically gives rise to a spectral sequence and this construction is functorial and maps preserve filtrations etc. etc. The question is then: what are the fibers of these maps? And there’s also the question of what the spectral sequence we get actually converges to, since we’re missing any information about things at height infinity (i.e. the additive formal group law over \mathbb{F}_p).

It turns out that, for coherent sheaves, this always converges to the right thing, and the E_1 term is obtained by looking at the local cohomology at *exactly* height n for each n . This gives a moral reason for the Morava change of rings result that we get in this case: we’re interested in the cohomology of some guy that lives on an infinitesimal neighborhood around the closed point at height n , so it makes sense that this is given by considering data related to Lubin-Tate deformations and the Morava stabilizer group.

2.2 Greek economy

Ok, it’s time we used all of this love to produce some nontrivial homotopy groups of spheres. The way to play is to combine two efforts: first use homotopy theory and geometry to create permanent cycles in the

ANSS that are in too low of a filtration to be hit by a differential, then use the CSS to show that these elements are nonzero in the ANSS and we're good.

The first step is done using v_n -self maps. Let's see how this goes in the known examples (this is sort of a review from Irina's talk, but now we have the insight of the ANSS).

First there is the mod p Moore spectrum, which is the cofiber of $S \rightarrow S$ given by multiplication by p . This induces a boundary map

$$\partial_0 : \pi_i V(0) \rightarrow \pi_{i-1} S$$

We also have a complex $V(1)$ which is given by the cofiber sequence

$$\Sigma^{2(p-1)} V(0) \xrightarrow{\phi_1} V(0) \rightarrow V(1)$$

where ϕ_1 induces multiplication by v_1 in BP -homology. We can define elements in $\pi_* S$ by taking:

$$\alpha_t = \partial_0(\phi_1^t \ell)$$

where $\ell : S \rightarrow V(0)$ is the inclusion of the bottom cell. Now, $\phi_1^t \ell$ is detected by an element in $H^* BP_* V(1)$ in the ANSS for $\pi_* V(1)$, and so α_t is detected by the corresponding element in $H^{*+1} BP_* V(0)$ after taking the coboundary $\delta : H^* BP_* V(1) \rightarrow H^{*+1} BP_* V(0)$. Similarly we can construct elements β_t and γ_t , for large primes (see Irina's lecture), and see that these are detected by permanent cycles in the ANSS. These come from cofiber sequences

$$\Sigma^{2(p^2-1)} V(1) \xrightarrow{\phi_2} V(1) \rightarrow V(2)$$

$$\Sigma^{2(p^3-1)} V(2) \xrightarrow{\phi_3} V(2) \rightarrow V(3)$$

where the first maps induce multiplication by v_n , and are defined by

$$\beta_t = \partial_0 \partial_1(\phi_2^t \ell)$$

$$\gamma_t = \partial_0 \partial_1 \partial_2(\phi_3^t \ell)$$

What is left is to describe the elements in the ANSS algebraically and show that they are nonzero. We'll do the first step, and Rob will say something about the second step in his talk.

Applying BP_* to every cofiber sequence we've mentioned gives us exact sequences:

$$0 \rightarrow BP_* \xrightarrow{p} BP_* \rightarrow BP_*/p \rightarrow 0$$

$$0 \rightarrow BP_*/(p) \xrightarrow{v_1} BP_*/(p) \rightarrow BP_*/(p, v_1) \rightarrow 0$$

$$0 \rightarrow BP_*/(p, v_1) \xrightarrow{v_2} BP_*/(p, v_1) \rightarrow BP_*/(p, v_1, v_2) \rightarrow 0$$

$$0 \rightarrow BP_*/(p, v_1, v_2) \xrightarrow{v_3} BP_*/(p, v_1, v_2) \rightarrow BP_*/(p, v_1, v_2, v_3) \rightarrow 0$$

Now, $H^0(BP_*/I)$, for I an invariant ideal, consists of the elements invariant mod I .⁵ Here, an invariant ideal is one such that $I \cdot BP_* BP = BP_* BP \cdot I$ (we're using the left and right actions, respectively). An element, $a \in BP_*$, is invariant modulo I if $\eta_R a \equiv \eta_a \text{ mod } I \cdot BP_* BP$.

In our case, then, we have the following:

$$v_1^t \in H^0(BP_*/(p))$$

$$v_2^t \in H^0(BP_*/(p, v_1))$$

$$v_3^t \in H^0(BP_*/(p, v_1, v_2))$$

⁵This terminology makes the most sense when we think of such comodules as quasi-coherent sheaves on a stack given as the orbits of some scheme by a group action. Then orbits correspond to invariants, and the vocabulary is justified.

Notice that these elements represent $\phi_n^t \ell \in \pi_* V(n-1)$ since the bottom cell is the unit and we said that ϕ_n induces multiplication by v_n .

Using iterated coboundaries, corresponding to those constructed homotopically, we may define elements in $H^* BP_*$,

$$\alpha_t \in H^1 BP_*$$

$$\beta_t \in H^2 BP_*$$

$$\gamma_t \in H^3 BP_*$$

detecting the elements of the same name in homotopy theory.

These are the infamous Greek letter elements in homotopy theory, and the proof that these are nonzero in the ANSS goes by way of the CSS, as we already said. The original reference [8] is still probably the best.

I think I'll stop there.

Remark 2.2.1. The title mentions the “spectrum of the category of spectra” but I didn’t have time to say anything about it. Rob will talk about the periodicity theorem and this is proven using a result that classifies thick subcategories of the stable homotopy category. There is a way to reformulate such a classification in terms of classifying prime tensor-ideals inside a triangulated category, and the collection of these gadgets forms the “prime ideal spectrum” of the category. For more, see the very readable paper introducing these ideas: [2].

References

- [1] Adams, J.F. *Stable Homotopy and Generalised Homology*. Chicago, 1996.
- [2] Balmer, P. *The spectrum of prime ideals in tensor-triangulated categories*. Available on his website.
- [3] Behrens, M. *Congruences between modular forms given by the divided β family in homotopy theory*. Available on his website.
- [4] Bousfield, A.K. *The localization of spectra with respect to homology*. *Topology*. Vol. 18. pp. 257-281.
- [5] Browder, W. *The Kervaire invariant of framed manifolds and its generalization*. *The Annals of Mathematics*. 2nd Ser., Vol. 90, No. 1 (Jul., 1969), 157-186.
- [6] Goerss, Paul. *The Adams-Novikov Spectral Sequence and the Homotopy Groups of Spheres*. On his website.
- [7] Lurie, Jacob. *Higher Algebra*. Available for download.
- [8] Miller, H.R., Ravenel, D.C., and Wilson, W.S. *Periodic Phenomena in the Adams-Novikov Spectral Sequence*. *The Annals of Mathematics*. Second Series, Vol. 106, No. 3. (Nov. 1977).