Algebraic Topology 3

Achim Krause

May 21, 2026

These are notes for an Algebraic Topology 3 course taught at Universitetet i Oslo in the spring semester 2026. They will be updated frequently over the course of the semester.

Contents

1 Introduction

The plan for this course is to cover key results from stable homotopy theory. Our goals will include:

- Spectra and their relationship to generalized (co)homology theories (Brown representability and related results)
- The Steenrod algebra and its dual
- The Adams spectral sequence
- Thom spectra

- The homotopy groups of MU
- The nilpotence and periodicity theorems of chromatic homotopy theory

While all of these constructions and results are very classical, the language in which they are discussed has undergone some serious revolutions, most notably the development of the theory of ∞-categories. While ∞-category theory has a reputation of being quite technical, it brings with it some serious advantages and makes several of the above stories much more conceptual (e.g. the symmetric-monoidal structure on spectra, the construction of spectral sequences, and the Thom isomorphism theorem.) On the other hand, the reputation of ∞-category theory as a very technical machinery is not entirely unwarranted, and a fully rigorous and self-contained introduction into that language would easily take a semester-long course on its own.

At the same time, there is a serious gap between "how one thinks about working in ∞-categories" and "how one writes these things down rigorously". While the latter is extensively documented in various excellent resources, the former is best learned in context. So the priorities of this course are as follows:

1. Cover the above topics from stable homotopy theory
2. ... while not shying away from the slick modern ways to express and prove some of these results, to see some ∞-category language in action
3. ... but avoiding to turn this into a technical introductory course to ∞-categories (not because I don't think of these as valuable, but because I feel several great resources for those exist already, [15], [9], [8])

For example, I am planning to explain for a bit what ∞-categories and functors between them are, and how to think about limits and colimits in that context (since this will give some very elegant perspectives on Thom spectra and the Thom isomorphism), but avoid proving all their properties.

2 Why ∞-categories?

Idea: In an ordinary category, we have for every pair of objects a *set* of morphisms. In an ∞-category, we want a *space* of morphisms $\mathrm{Map}(X, Y)$.

- Every 1-category should give an ∞-category where the mapping space is discrete.
- There should be an ∞-category $\mathcal{S}$ of "nice spaces" (e.g. CW complexes), where the morphism spaces are mapping spaces.
- For any space X and any points $x, y \in X$ there is a space of paths $P_{x,y}$, and there should be an ∞-category with objects the points of X and morphism spaces these path spaces.

From an ∞-category $\mathcal{C}$, we may pass back to an ordinary category $\mathrm{Ho}(\mathcal{C})$, with the same objects and $\mathrm{Hom}_{\mathrm{Ho}(\mathcal{C})}(X, Y) = \pi_0 \, \mathrm{Map}_{\mathcal{C}}(X, Y)$. In our examples, this gives

- For an 1-category viewed as ∞-category with discrete mapping spaces, $\mathrm{Ho}(\mathcal{C})$ recovers the original 1-category.
- $\mathrm{Ho}(\mathcal{S})$ is the ordinary homotopy category of CW complexes.
- Passing from a space X to the "path ∞-category" discussed above and passing back to the homotopy category, we obtain a category whose objects are the points of X, and where $\mathrm{Hom}(x, y)$ is the set of *homotopy classes* of paths from x to y. This is the so-called "fundamental groupoid" of X, since for $x_0 \in X$, $\mathrm{Hom}(x_0, x_0)$ in that category is exactly $\pi_1(X, x_0)$, the fundamental group. The ∞-categorical version is sometimes called "fundamental ∞-groupoid".

There are different ways one could try to make a theory of ∞-categories precise. The above discussion suggests that one might try to do this in terms of "categories enriched in topological spaces". That works, but brings some technical problems with it, for example in the third example above, since composition of paths is not associative, only up to homotopy. The most successful approach is therefore a different one: *quasicategories*, pioneered by Joyal and fully developed by Lurie. Here the idea is to encode an ∞-category as simplicial set $\mathcal{C}_\bullet$ with a "horn filling" condition - one can still extract mapping spaces (more precisely, simplicial sets) $\mathrm{Map}(x, y)$ from every pair of objects, together with a homotopy-associative composition, but the whole setup is flexible enough that all three examples can directly be written down:

- A 1-category can be viewed as ∞-category by forming the simplicial set whose n-simplices are diagrams $x_0 \xrightarrow{f_0} \ldots \xrightarrow{f_{n-1}} x_n$ in $\mathcal{C}$. (The *nerve* of $\mathcal{C}$.)
- $\mathcal{S}$ can be obtained from the actual space-enriched category of CW complexes by a construction called "homotopy-coherent nerve" where 0-simplices are CW complexes, 1-simplices are continuous maps, 2-simplices are tuples of spaces X, Y, Z, maps $f : X \to Y$, $g : Y \to Z$, $h : X \to Z$ and a homotopy between h and $g \circ f$, and higher-dimensional simplices encode "homotopies of homotopies".
- The ∞-category where mapping spaces are path spaces is obtained surprisingly elegantly as $\mathrm{Sing}(X)$, the same simplicial set at the one used to define homology. Here 1-simplices are literally paths, and the higher-dimensional simplices encode homotopies of paths indirectly.

Functors between ∞-categories are just morphisms of simplicial sets. For two ∞-categories $\mathcal{C}$ and $\mathcal{D}$, there is a functor category $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ whose n-simplices are $\mathrm{Hom}_{\mathrm{sSet}}(\Delta^n \times \mathcal{C}, \mathcal{D})$. We think of morphisms in $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ as natural transformations.

There are a lot of technical challenges when setting up the theory. For example, the fact that $\mathrm{Map}_{\mathcal{C}}(x,y)$ is functorial in both arguments (i.e. may be viewed as functor $\mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{S}$) is not easy to prove at all. We will try to stick to an informal understanding of ∞-category theory for these notes.

2.1 Groupoids and spaces

A morphism in $\mathcal{C}$ can be thought of either as a functor $\Delta^1 \to \mathcal{C}$ or as a point in $\mathrm{Map}_{\mathcal{C}}(x,y)$ for objects $x, y \in \mathcal{C}$. It is invertible if it becomes invertible in $\mathrm{Ho}(\mathcal{C})$. We call x, y equivalent if there is an invertible morphism between them. We call $\mathcal{C}$ an ∞-groupoid if all morphisms are invertible, and for an ∞-category $\mathcal{C}$, there exists a *groupoid core* $\mathcal{C}^{\simeq}$ which is universal with respect to functors from groupoids into $\mathcal{C}$ (each $\mathcal{D} \to \mathcal{C}$ factors through $\mathcal{C}^{\simeq}$). We can think of $\mathrm{Fun}(\mathcal{C},\mathcal{D})^{\simeq}$ as a version of the functor category where we only allow invertible natural transformations. A functor $F : \mathcal{C} \to \mathcal{D}$ is an equivalence if there exists an inverse $G : \mathcal{D} \to \mathcal{C}$ such that both composites $F \circ G$ and $G \circ F$ are equivalent to $\mathrm{id}_{\mathcal{D}}$ and $\mathrm{id}_{\mathcal{C}}$ in $\mathrm{Fun}(\mathcal{D},\mathcal{D})^{\simeq}$ and $\mathrm{Fun}(\mathcal{C},\mathcal{C})^{\simeq}$.

Theorem 2.1 ("Homotopy hypothesis")**.** Sing *provides an equivalence between CW complexes up to homotopy and ∞-groupoids up to equivalence.*

This is intentionally vaguely phrased, but can be made precise as a combination of two things: Sing provides an equivalence between ∞-categories of CW complexes and of Kan complexes, and ∞-groupoids are actually exactly Kan complexes by a theorem of Joyal.

So we can think of the above fundamental ∞-groupoid construction as allowing us to view spaces (and their homotopy theory) as a special case of ∞-category theory, namely groupoids.

Example 2.2. Let G be a group. The nerve of the 1-groupoid with one object $*$ and $\mathrm{Hom}(*,*) = G$ is an ∞-groupoid. It is equivalent to the fundamental ∞-groupoid of the space BG. (This comes from the above equivalence and the description of BG as geometric realisation of the nerve of G.)

2.2 Diagrams and (co)limits

We now think about differently shaped diagrams. For any simplicial set K, we can still make sense of $\mathrm{Fun}(K,\mathcal{C})$ just as maps of simplicial sets. We think of a functor $K \to \mathcal{C}$ as a *K-shaped diagram* in $\mathcal{C}$, and have the following examples:

- For $K = \Delta^0$ (the point), a functor $\Delta^0 \to \mathcal{C}$ is just an object of $\mathcal{C}$.
- For $K = \Delta^1$, a functor $\Delta^1 \to \mathcal{C}$ is just a morphism of $\mathcal{C}$.
- For $K = \Delta^1 \amalg_{\Delta^0} \Delta^1$, two 1-simplices glued source-to-source, a functor $K \to \mathcal{C}$ is a pair of morphisms which share the source.

- For $K = \Delta^1 \times \Delta^1$, a functor $K \to \mathcal{C}$ can be thought of as a "commutative square" in $\mathcal{C}$, but in the following sense: It consists of objects and morphisms
$$\begin{array}{ccc} x_{00} & \xrightarrow{f_0} & x_{01} \\ \downarrow{\scriptstyle g_0} & & \downarrow{\scriptstyle g_1} \\ x_{10} & \xrightarrow{f_1} & x_{11} \end{array}$$
but then the data on the 2-simplices in $\Delta^1 \times \Delta^1$ can be thought of as providing a *homotopy* in $\mathrm{Map}(x_{00}, x_{11})$ between $g_1 \circ f_0$ and $f_1 \circ g_0$. So where we would have *equalities* in a commutative diagram in 1-category theory, we have instead homotopies in ∞-category theory, and these are part of the data of the diagram.

Now recall the universal property of pushouts in 1-categories: The pushout of a diagram $Y \leftarrow X \to Z$ comes with a commutative diagram

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ Z & \longrightarrow & P, \end{array}$$

which is universal in the sense that for each other commutative diagram like this (with another object P' replacing P), there is a unique $P \to P'$ making everything commute. Said differently, postcomposing the above fixed diagram with P defines a bijection between $\mathrm{Hom}(P, P')$ and commutative diagrams with P' in the lower right corner.

Definition 2.3. Let $Y \leftarrow X \to Z$ be a diagram in an ∞-category. A pushout consists of a commutative square

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ Z & \longrightarrow & P, \end{array}$$

such that postcomposing defines an equivalence between the mapping space $\mathrm{Map}(P, P')$ and the space of commutative diagrams

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ Z & \longrightarrow & P'. \end{array}$$

The latter can be made precise as a simplicial set (contained in $\mathrm{Fun}(\Delta^1 \times \Delta^1, \mathcal{C})$).

Example 2.4. Let $Y \leftarrow X \rightarrow Z$ be a diagram in $\mathcal{S}$. Then a pushout is given by

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ Z & \longrightarrow & Z \amalg_{X\times 0} (X \times [0,1]) \amalg_{X\times 1} Y, \end{array}$$

with the obvious morphisms and homotopy coming from the canonical map from $X \times [0,1]$.

Using the interpretation of commutative squares as tuples of maps and a homotopy again, the universality can be sketched as follows: A commutative diagram with P' in the lower right corner is exactly maps from Y and Z and an X-indexed homotopy, so a map from $Y \amalg_{X\times 0} (X \times [0,1]) \amalg_{X\times 1} Z$.

If at least one of the maps $X \to Y$ and $X \to Z$ is a nice enough embedding, this "homotopy pushout" is homotopy equivalent to the usual pushout, but its universal property makes the latter more well-behaved (for example it is invariant under equivalences of the diagram we started with). One great special case is given by the case where $Y, Z = \mathrm{pt}$. Here the ∞-categorical pushout comes out as suspension ΣX (whereas the 1-categorical pushout would be pt (or $\mathrm{pt} \amalg \mathrm{pt}$ if $X = \emptyset$).

Dually, we have the concept of *pullbacks*, whose universal property says that the space of maps $\mathrm{Map}(P', P)$ for a pullback P of $X \to Z \leftarrow Y$ agrees with the space of commutative diagrams

$$\begin{array}{ccc} P' & \longrightarrow & X \\ \downarrow & & \downarrow \\ Y & \longrightarrow & Z \end{array}$$

Example 2.5. Let $X \to Z \leftarrow Y$ be a diagram in $\mathcal{S}$. Then a pullback is given by *the space of diagrams*

$$\begin{array}{ccc} \mathrm{pt} & \longrightarrow & X \\ \downarrow & & \downarrow \\ Y & \longrightarrow & Z \end{array}$$

itself. Informally, this is the space of triples of a point in X, a point in Y, and a path between their images in Z. In the special case $X = Y = \mathrm{pt}$ and both maps $\mathrm{pt} \to Z$ given by the same basepoint of Z, this is therefore the space of loops at the basepoint in Z, the *loop space* ΩZ.

Other universal properties make sense as well (in fact, colimits and limits over arbitrary diagram shapes make sense), for example an *initial object* in an ∞-category $\mathcal{C}$ is a $c \in \mathcal{C}$ with $\mathrm{Map}_{\mathcal{C}}(c, d)$ contractible for each $d \in \mathcal{C}$, and dually terminal objects. In $\mathcal{S}$, the empty space $\emptyset$ is initial and the 1-point space pt is terminal. In the category of pointed spaces $\mathcal{S}_*$, pt is both terminal and initial, a so-called *zero object*.

Coproducts and products also make sense, and are what one would think in our examples: In $\mathcal{S}$, coproducts are given by disjoint union and products by cartesian products. In $\mathcal{S}_*$, coproducts are given by wedge sums and products by cartesian products again.

3 Spectra

3.1 Spectra and cohomology theories

For pointed spaces X and Y, since both of the spaces $\mathrm{Map}_{\mathcal{S}_*}(\Sigma X, Y)$ and $\mathrm{Map}_{\mathcal{S}_*}(X, \Omega Y)$ are equivalent to the space of diagrams

$$\begin{array}{ccc} X & \longrightarrow & \mathrm{pt} \\ \downarrow & & \downarrow \\ \mathrm{pt} & \longrightarrow & Y, \end{array}$$

there is a natural equivalence $\mathrm{Map}_{\mathcal{S}_*}(\Sigma X, Y) \simeq \mathrm{Map}_{\mathcal{S}_*}(X, \Omega Y)$. (This expresses that Σ is left adjoint to Ω.) Passing to π_0 (i.e. homotopy classes), this means $[\Sigma X, Y] \cong [X, \Omega Y]$. In particular (for $X = S^n$), we learn $\pi_n(\Omega Y) \cong \pi_{n+1}(Y)$, so Ω "shifts homotopy groups down".

Definition 3.1. A *spectrum* is a sequence of pointed spaces X_n, together with equivalences $X_n \simeq \Omega X_{n+1}$.

So a spectrum consists of an "underlying space" X_0 together with a sequence of choices of deloopings. Ω is not invertible, that is there is no unique (or even canonical) way to choose X_{n+1} from X_n.

Example 3.2. For an abelian group A, $\Omega K(A, n+1)$ is itself a $K(A, n)$, so there exists a sequence of equivalences $K(A, n) \simeq \Omega K(A, n+1)$. This gives a spectrum usually denoted HA (and called *Eilenberg-MacLane spectrum for the abelian group* A).

A map between spectra consists of a sequence of pointed maps $f_n : X_n \to Y_n$ and *choices of homotopies* making the diagrams

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

commute. It is easy to define a space of such maps (as a space of diagrams), but composing such maps provides some difficulty since composing homotopies is not associative. It is however not hard to *directly* construct an ∞-category Sp of spectra as a full subcategory of a suitable functor category (encoding the spaces X_n and the maps $X_n \to \Omega X_{n+1}$ into one diagram), for details see [13, Section 1.4].

Remark 3.3. Since via $X_n \simeq \Omega X_{n+1}$, X_{n+1} determines X_n (uniquely in the ∞-categorical sense), we may equally well extend the indexing to all $n \in \mathbb{Z}$. We can still think of X_0 as the underlying space, the X_{-1}, X_{-2} etc. will just be its loop spaces.

Since we may think of a spectrum X as an "infinite delooping" of X_0, the forgetful functor $\mathrm{Sp} \to \mathcal{S}_*$ which informally takes X to X_0 is often denoted Ω^∞, and more generally the forgetful functor which informally takes X to X_n by $\Omega^{\infty-n}$. Note that

$$\Omega^{\infty-n}X \simeq \Omega(\Omega^{\infty-n-1}X),$$

and so

$$\pi_{k+n}(\Omega^{\infty-n}X) \cong \pi_{k+n+1}(\Omega^{\infty-n-1}X)$$

as long as $k+n \geq 0$. The value of this group for large enough n is simply denoted $\pi_k(X)$. Note that for nonnegative k, $\pi_k(X) = \pi_k(\Omega^\infty X)$, but this is also defined for negative k, we just have to look in one of the later deloopings then. A uniform way to say this is $\pi_k X = \pi_0 \Omega^{\infty+k} X$, with the indexing extended to $\mathbb{Z}$ as explained above.

Definition 3.4. For a spectrum E and a pointed space X, we write

$$\widetilde{E}^k(X) := [X, \Omega^{\infty-k}E]$$

where we write $[X,Y] := \pi_0 \operatorname{Map}_{\mathcal{S}_*}(X,Y)$ for the set of pointed homotopy classes. Note that in this case, this is an abelian group, as the target is a double loop space

Example 3.5.

$$\widetilde{E}^0(S^k) = \widetilde{E}^{-k}(S^0) = \pi_k(E).$$

Lemma 3.6. *For a spectrum E, the $\widetilde{E}^*(-)$ form a (generalized) cohomology theory on pointed spaces.*

Proof. Homotopy invariance holds by construction (since these are functors on the ∞-category $\mathcal{S}_*$), and additivity follows since $\operatorname{Map}_{\mathcal{S}_*}(-,Y)$ turns coproducts into products and π_0 preserves those. We also have suspension isomorphisms $\widetilde{E}^k(X) \simeq \widetilde{E}^{k+1}(\Sigma X)$ by construction, since $[\Sigma X, \Omega^{\infty-k-1}E] \simeq [X, \Omega(\Omega^{\infty-k-1}E)] \simeq [X, \Omega^{\infty-k}E]$, with the former isomorphism coming from the adjunction between Σ and Ω, and the latter isomorphism coming from the structure maps in the spectrum.

For a map of pointed spaces $A \to X$, we may form the *cofiber*, a pointed space X/A, as pushout

$$\begin{array}{ccc} A & \longrightarrow & X \\ \downarrow & & \downarrow \\ \mathrm{pt} & \longrightarrow & X/A. \end{array}$$

(If $A \to X$ comes from a good inclusion of actual topological spaces, this X/A will be equivalent to the actual quotient, but this thing always exists.) Applying $\mathrm{Map}_{\mathcal{S}_*}(-, \Omega^{\infty-k}E)$, we get a pullback diagram

$$\begin{array}{ccc} \mathrm{Map}_{\mathcal{S}_*}(X/A, \Omega^{\infty-k}E) & \longrightarrow & \mathrm{Map}_{\mathcal{S}_*}(X, \Omega^{\infty-k}E) \\ \downarrow & & \downarrow \\ \mathrm{pt} & \longrightarrow & \mathrm{Map}_{\mathcal{S}_*}(A, \Omega^{\infty-k}E), \end{array}$$

which on π_0 gives rise to an exact sequence

$$\widetilde{E}^k(X/A) \to \widetilde{E}^k(X) \to \widetilde{E}^k(A).$$

We may apply the same construction to $X \to X/A$, which, by pasting the pushout squares

$$\begin{array}{ccccc} A & \longrightarrow & X & \longrightarrow & \mathrm{pt} \\ \downarrow & & \downarrow & & \downarrow \\ \mathrm{pt} & \longrightarrow & X/A & \longrightarrow & (X/A)/X, \end{array}$$

has cofiber $(X/A)/X \simeq \Sigma A$, yielding an exact sequence

$$\widetilde{E}^k(\Sigma A) \to \widetilde{E}^k(X/A) \to \widetilde{E}^k(X)$$

or

$$\widetilde{E}^{k-1}(A) \to \widetilde{E}^k(X/A) \to \widetilde{E}^k(X).$$

Iterating and combining these gives the long exact sequence. □

Example 3.7. For the Eilenberg-MacLane spectrum HA, the corresponding cohomology theory is (reduced) singular cohomology $\widetilde{H}^*(X; A)$.

Theorem 3.8 (Brown representability). *Let h^* be a generalized cohomology theory. Then there exists a spectrum E and a natural isomorphism $\widetilde{E}^* \to h^*$.*

The proof is based on a more technical representability result for individual functors:

Lemma 3.9. *Let $\mathcal{S}_{*,\geq 1}$ be the full subcategory of $\mathcal{S}_*$ on pointed connected spaces, and consider a functor $h : \mathcal{S}_{*,\geq 1}^{\mathrm{op}} \to \mathrm{Set}_*$ with the following properties:*

1. *h takes coproducts to products (arbitrary indexing set, so this includes $h(\mathrm{pt}) \cong \mathrm{pt}$).*
2. *For a pushout square*

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ C & \longrightarrow & D \end{array}$$

of pointed connected spaces, we have that in the square

$$\begin{array}{ccc} h(D) & \longrightarrow & h(B) \\ \downarrow & & \downarrow \\ h(C) & \longrightarrow & h(A) \end{array}$$

the canonical map from $h(D)$ to the actual pullback is surjective. For $C = \mathrm{pt}$, $D \simeq B/A$ and this specializes to exactness of $h(B/A) \to h(B) \to h(A)$, in the sense that the first term maps surjectively to the preimage of the basepoint under the last map.

Then there exists $K \in \mathcal{S}_{*,\geq 1}$ with a natural isomorphism $[-, K] \to h$.

Proof. One can think of the pushout condition as a kind of nonadditive Mayer-Vietoris sequence, since if h took values in abelian groups, it would just correspond to $h(D) \to h(B) \oplus h(C) \to h(A)$ being exact in the middle. The sequence can be extended to the left, with group structures showing up "step by step" and exactness being interpreted in a subtle way:

1. From the pinch map $\Sigma A \to \Sigma A \vee \Sigma A$, $h(\Sigma A)$ inherits a group structure. The pushout $D \vee_{B \vee C} D$ is equivalent to $D \vee \Sigma A$, and the second map $D \to D \vee \Sigma A$ induces a $h(\Sigma A)$ group action on $h(D)$. Surjectivity of $h(D) \times h(\Sigma A) \cong h(D \vee_{B \vee C} D) \to h(D) \times_{h(B) \times h(C)} h(D)$ (another application of the pushout condition) means that two elements of $h(D)$ with the same image in $h(B) \times h(C)$ differ by acting with an element of $h(\Sigma A)$. So $h(\Sigma A)$ acts transitively on the fibers of $h(D) \to h(B) \times_{h(A)} h(C)$.

2. Analogously, the pushout $(D \vee_{B \vee C} D) \vee_{D \vee D} (D \vee_{B \vee C} D)$ is equivalent to $\Sigma B \vee \Sigma C$. Identifying $D \vee_{B \vee C} D \simeq D \vee \Sigma A$, the corresponding application of the pushout condition implies that for $d \in h(D)$ and two $a_1, a_2 \in h(\Sigma A)$ whose action on d produces the same element in $h(D)$, a_1 and a_2 differ by an element in the image of the group homomorphism $h(\Sigma B \vee \Sigma C) \to h(\Sigma A)$. That is, the orbits of the $h(\Sigma A)$-action on $h(D)$ all have stabilizer exactly the image of $h(\Sigma B) \times h(\Sigma C)$.

3. At the next position, we simply learn that the kernel of $h(\Sigma B) \times h(\Sigma C) \to h(\Sigma A)$ is the image of $h(\Sigma D) \to h(\Sigma B) \times h(\Sigma C)$, and after that the groups are even abelian and the Mayer-Vietoris sequence continues as expected.

The strategy will now be to construct connected K with a natural transformation $[-, K] \to h$ which induces an isomorphism on all S^i for $i \geq 1$. Let us first observe that this suffices, i.e. that in this case the natural transformation is actually an isomorphism on all $X \in \mathcal{S}_{*,\geq 1}$. To do this, consider the collection of all X such that this map is an isomorphism on $\Sigma^i X$ for all $i \geq 0$. By assumption, this contains spheres. Since both sides have Mayer-Vietoris sequences in the sense explained above, a diagram chase analogous to the 5-lemma shows that the collection of all such X is closed under pushouts. It is also closed under infinite coproducts. If $\ldots X_n \to X_{n+1} \to \ldots$ is a sequence of spaces in this collection, their colimit can be written as pushout

$$\begin{array}{ccc} \bigvee_{n\geq 0} X_n & \longrightarrow & \bigvee_{n\geq 0} X_{2n} \\ \downarrow & & \downarrow \\ \bigvee_{n\geq 0} X_{2n+1} & \longrightarrow & \operatorname{colim} X_n, \end{array}$$

where the top horizontal map is given by the identity on even indices and by the map $X_{2n-1} \to X_{2n}$ on odd indices, and the left vertical map analogously. So this class is also closed under sequential colimits. Describing every object as CW complex, we see that the class of objects contains everything as desired.

We now construct K with the desired natural transformation. Assume first we have constructed a K_n with a natural transformation $[-, K_n] \to h(-)$ which induces an isomorphism on $S^1, \ldots, S^{n-1}$ and a surjection on S^n. Let α_i be generators for the kernel of

$$\pi_n(K_n) \cong [S^n, K_n] \to h(S^n),$$

and construct K'_n as the pushout

$$\begin{array}{ccc} \bigvee S^n & \xrightarrow{(\alpha_i)} & K_n \\ \downarrow & & \downarrow \\ \mathrm{pt} & \longrightarrow & K'_n. \end{array}$$

The natural transformation $[-, K_n] \to h$ corresponds (by Yoneda) to an element of $h(K_n)$. Since $h(K'_n)$ maps surjectively to the "kernel" (preimage of basepoint) of $h(K_n) \to h(\bigvee S^n)$, it extends to a natural transformation $[-, K'_n] \to h$. Now let β_j be generators of $h(S^{n+1})$ and let $K_{n+1} = K'_n \vee \bigvee S^{n+1}$, with the resulting natural transformation $[-, K_{n+1}] \to h$ (which by Yoneda corresponds again to an element of $h(K_{n+1}) \cong h(K'_n) \times \prod h(S^{n+1})$.) Now $[-, K_{n+1}] \to h$ is an isomorphism on $S^1, \ldots, S^n$ and surjective on S^{n+1}.

Finally, we may write the colimit K of $K_0 \to K_1 \to K_2 \to \ldots$ as a pushout,

$$\begin{array}{ccc} \bigvee_{n\geq 0} K_n & \longrightarrow & \bigvee_{n\geq 0} K_{2n} \\ \downarrow & & \downarrow \\ \bigvee_{n\geq 0} K_{2n+1} & \longrightarrow & \operatorname{colim} K_n =: K. \end{array}$$

Applying h, we find an element of $h(K)$ which restricts to the given elements of $h(K_n)$, i.e. a natural transformation $[-, K] \to h$ which restricts to the given natural transformations $[-, K_n] \to h$. This is now an isomorphism on S^i for all $i \geq 1$, so $[-, K] \to h$ is an isomorphism on all of $\mathcal{S}_{*,\geq 1}$. □

Proof of Brown representability. For a cohomology theory h^*, we find pointed connected spaces E'_n with natural isomorphisms $[-, E'_n] \to h^n$ on $\mathcal{S}_{*,\geq 1}$. Since

$$[X, \Omega E'_{n+1}] \cong [\Sigma X, E'_{n+1}] \cong h^{n+1}(\Sigma X) \cong h^n(X),$$

if we set $E_n = \Omega E'_{n+1}$, we get natural isomorphisms $[-, E_n] \cong h^n$ on all of $\mathcal{S}_*$. Now the suspension isomorphism gives natural isomorphisms

$$[-, E_n] \simeq h^n \simeq h^{n+1}(\Sigma(-)) \simeq [\Sigma(-), E_{n+1}] \simeq [-, \Omega E_{n+1}],$$

thus by Yoneda an equivalence $E_n \simeq \Omega E_{n+1}$. The E_n hence assemble into a spectrum with equivalences $\widetilde{E}^* \simeq h^*$. □

Warning 3.10. Even though Brown representability tells us that there is a close connection between spectra and generalized cohomology theories, it is *not* true that homotopy classes of maps between spectra are the same as natural transformations of cohomology theories. It is true that every natural transformation between cohomology theories lifts to a map of representing spectra (in particular, if two spectra represent isomorphic cohomology theories, they are equivalent, so one has a certain uniqueness of the representing spectrum guaranteed by Brown representability), but the map

$$[E, F] \to \{\text{nat. trafos } \widetilde{E}^* \to \widetilde{E}^*\}$$

is in general not injective. This is the phenomenon of so-called *phantom maps*.

3.2 Suspension spectra

For a pointed space X, we have $\mathrm{Map}_{\mathcal{S}_*}(X, \Omega^\infty E) \simeq \mathrm{Map}_{\mathcal{S}_*}(\Sigma^n X, \Omega^{\infty-n} E)$, from iterating the adjunction between Σ and Ω.

If we momentarily talk about slightly more general objects and let a *prespectrum* be a sequence of Y_k with maps $Y_k \to \Omega Y_{k+1}$ which are not necessarily equivalences, we have for every n a prespectrum with spaces

$$\Omega^n \Sigma^n X, \Omega^{n-1} \Sigma^n X, \ldots, \Sigma^n X, \mathrm{pt}, \mathrm{pt}, \mathrm{pt}, \ldots.$$

A map from such a prespectrum to a spectrum is completely determined by the map $\Sigma^n X \to \Omega^{\infty-n} E$, so the space of maps from this prespectrum to E is $\mathrm{Map}_{\mathcal{S}_*}(\Sigma^n X, \Omega^{\infty-n})$. We also have a map from the n-th such prespectrum to the $(n+1)$-st, determined by the map $\Sigma^n X \to \Omega \Sigma^{n+1} X$, and inducing the equivalence

$$\mathrm{Map}_{\mathcal{S}_*}(\Sigma^{n+1}, \Omega^{\infty-n-1} E) \to \mathrm{Map}_{\mathcal{S}_*}(\Sigma^n, \Omega^{\infty-n} E).$$

Taking the colimit over n, we obtain a spectrum again, since at the n-th stage already the structure maps until the n-th are equivalences, and Ω commutes with sequential colimits (this is a special case of a basic fact about $\mathcal{S}$, that finite limits commute with filtered colimits). This spectrum is denoted $\Sigma^\infty X$, and by construction satisfies

$$\mathrm{Map}_{\mathrm{Sp}}(\Sigma^\infty X, E) \simeq \mathrm{Map}_{\mathcal{S}_*}(X, \Omega^\infty E).$$

A priori we have only constructed this objectwise, but from basic facts about adjoint functors, one can conclude that this uniquely determines a functor $\Sigma^\infty : \mathcal{S}_* \to \mathrm{Sp}$ left adjoint to Ω^∞.

For *unpointed spaces*, we have a functor $\mathcal{S} \to \mathcal{S}_*$ which freely adds a basepoint. This is left adjoint to the forgetful functor, and we define $\Sigma^\infty_+ : \mathcal{S} \to \mathrm{Sp}$ as the composite. Composing adjoints, we see

$$\mathrm{Map}_{\mathrm{Sp}}(\Sigma^\infty_+ X, E) \simeq \mathrm{Map}_{\mathcal{S}}(X, \Omega^\infty E).$$

What does $\Sigma^\infty X$ look like? From the above construction, its k-th space is given by $\mathrm{colim}_n\, \Omega^{n-k}\Sigma^n X$ (note that this works for $k \in \mathbb{Z}$). Also note that the notation works out beautifully: We have

$$\Omega^{\infty-m}\Sigma^\infty X = \mathrm{colim}_n\, \Omega^{n-m}\Sigma^n X.$$

Since homotopy groups commute with filtered colimits of spaces,

$$\pi_k\Sigma^\infty X = \pi_0\Omega^{\infty+k}\Sigma^\infty X = \mathrm{colim}_n\, \pi_0(\Omega^{n+k}\Sigma^n X) = \mathrm{colim}_n\, \pi_{n+k}(\Sigma^n X).$$

These are exactly the stable homotopy groups of the space X. By the Freudenthal theorem, the suspension homomorphisms between the groups in the diagram on the right become isomorphisms for n large enough, so the colimit is actually just given by $\pi_{n+k}(\Sigma^n X)$ for n from this stable range.

Definition 3.11. The sphere spectrum $\mathbb{S}$ is the suspension spectrum $\Sigma^\infty S^0$.

Its underlying space is $\Omega^\infty\Sigma^\infty S^0 \simeq \mathrm{colim}_n\, \Omega^n\Sigma^n S^0 \simeq \mathrm{colim}\, \Omega^n S^n$. Its homotopy groups are given by the stable homotopy groups of spheres, $\pi_k\mathbb{S} \cong \pi_{n+k}S^n$ for n large enough.

More generally, we can write $\mathbb{S}^n = \Sigma^\infty S^n$ and have

$$[\mathbb{S}^n, E] \cong [S^n, \Omega^\infty E] \cong \pi_n E,$$

so homotopy groups are still corepresented by spheres in Sp. But as we saw above, we also have negative degree homotopy groups of spectra. Do we also have negative-dimensional spheres?

3.3 Stability

The spectrum Σ^∞ pt is just the spectrum whose n-th space is pt for every n. It is easy to see that it is both terminal and initial in Sp, so Sp has a zero object (like $\mathcal{S}_*$), and we typically write $\Sigma^\infty\,\mathrm{pt} = 0$. Analogous to $\mathcal{S}_*$, we can also define a "loop" functor $\Omega : \mathrm{Sp} \to \mathrm{Sp}$ in terms of pullbacks

$$\begin{array}{ccc} \Omega X & \longrightarrow & 0 \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & X, \end{array}$$

now in the category Sp. Since spectra are defined in terms of $\Omega : \mathcal{S}_* \to \mathcal{S}_*$, which is a limit and hence commutes with limits, one can check that limits of spectra are computed levelwise (on the underlying pointed spaces), in particular $\Omega : \mathrm{Sp} \to \mathrm{Sp}$ is just given by applying Ω levelwise to the underlying spectrum. Since in a spectrum, $X_n \simeq \Omega X_{n+1}$, one can see that $\Omega : \mathrm{Sp} \to \mathrm{Sp}$ is equivalent to an index shift! In particular (indexing our spectra over $\mathbb{Z}$), it is an equivalence.

We also have a functor $\Sigma : \mathrm{Sp} \to \mathrm{Sp}$ defined through pushouts

$$\begin{array}{ccc} X & \longrightarrow & 0 \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & \Sigma X, \end{array}$$

and as in $\mathcal{S}_*$, one gets that they are adjoint, $\mathrm{Map}_{\mathrm{Sp}}(\Sigma X, Y) \simeq \mathrm{Map}_{\mathrm{Sp}}(X, \Omega Y)$. But since Ω is invertible, this is also naturally equivalent to $\mathrm{Map}_{\mathrm{Sp}}(\Omega^{-1} X, Y)$, so the Yoneda lemma tells us that $\Sigma \simeq \Omega^{-1}$. So Σ can be identified with the index shift "up", and we will often just write Σ^{-1} for Ω and more generally Σ^k for $k \in \mathbb{Z}$ for the suitable power of Σ or Ω.

(Technically, we don't a priori know that Sp has pushouts, but one can reverse this argument to check that $\Omega^{-1} X$ satisfies the right universal property to be the pushout in question.)

The fact that Σ is invertible lets us define $\mathbb{S}^n = \Sigma^n \mathbb{S}$ for all $n \in \mathbb{Z}$ (and since Σ^∞ commutes with colimits, this agrees with the earlier definition for $n \geq 0$), and have

$$[\mathbb{S}^n, E] \cong [\mathbb{S}, \Sigma^{-n} E] \cong \pi_n(E),$$

so the spheres corepresent all homotopy groups.

A nicely symmetric way to express the fact that Σ and Ω are inverse to each other is that a square

$$\begin{array}{ccc} X & \longrightarrow & 0 \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & Y \end{array}$$

in Sp is a pullback diagram if and only if it is a pushout diagram. By some clever manipulations of pushouts and pullbacks, one can prove from this that more generally, any square of spectra

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ Z & \longrightarrow & W \end{array}$$

is a pushout if and only if it is a pullback. One says Sp is *stable*.

Definition 3.12. An ∞-category is *stable* if it has pushouts, pullbacks, a zero object, and the property that a square is a pushout diagram if and only if it is a pullback diagram.

There are a number of further consequences that follow from this:

1. Finite coproducts and products agree. That is, the canonical map $X \amalg Y \to X \times Y$ which uses identity and zero maps (the unique maps through the zero object) are equivalences. One simply writes $X \oplus Y$ for coproducts and products.

2. Morphisms can be added: For $f, g : X \to Y$, we can form $f + g : X \to X \oplus X \xrightarrow{f \oplus g} Y \oplus Y \to Y$. Maps between finite (co)products can be described as matrices of maps, with composition compatible with the usual description of matrix multiplication (using the above sum). The monoid structure on $[X, Y]$ given by $+$ is a group, so we may also subtract morphisms.

3. A pushout square of the form

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & Z \end{array}$$

   is called a *cofiber sequence*, a pullback square of this form a fiber sequence. We write $X \to Y \to Z$ (even though the nullhomotopy is part of the data!) Stability implies that cofiber sequences are fiber sequences and vice versa.

4. We can complete any morphism $X \xrightarrow{f} Y$ to a cofiber sequence $X \to Y \to \mathrm{cofib}(f)$ by forming a pushout, and to a fiber sequence $\mathrm{fib}(f) \to X \to Y$ by forming a pullback. Since cofiber sequences are fiber sequences, we can recover X as fiber of $Y \to \mathrm{cofib}(f)$, and analogously Y as cofiber of $\mathrm{fib}(f) \to X$.

5. If we iterate forming cofibers, we get cofiber sequences $Y \to \mathrm{cofib}(f) \to \Sigma X$, $\mathrm{cofib}(f) \to \Sigma X \xrightarrow{\Sigma f} \Sigma Y$, continuing in a 3-periodic way with more and more suspensions. Analogously iterating fibers gives a sort of 3-periodic sequence of fiber sequences with powers of Ω.

6. Pushout/pullback diagrams

$$\begin{array}{ccc} X & \xrightarrow{i} & Y \\ \downarrow{\scriptstyle j} & & \downarrow{\scriptstyle f} \\ Z & \xrightarrow{g} & W \end{array}$$

   correspond to cofiber/fiber sequences $X \xrightarrow{(i,j)} Y \oplus Z \xrightarrow{\begin{pmatrix} f \\ -g \end{pmatrix}} W$.

These are in a sense the ∞-categorical analogue of abelian categories, with cofiber/fiber roughly corresponding to cokernel/kernel, and (co)fiber sequences roughly corresponding to short exact sequences. Note that there are no left-exact or right-exact sequences here, only exact ones, and relatedly the analogue of the fourth observation (that the fiber of the cofiber of $f : X \to Y$ recovers X) only works for monomorphisms in an abelian category.

Stability of Sp tells us that contrary to pullbacks, pushouts of spectra are more subtle. For example, since a pushout diagram in Sp is also a pullback diagram, and pullbacks are formed underlying, on underlying spaces a pushout diagram of spectra looks like a pullback diagram in $\mathcal{S}_*$ (and these typically are not pushout diagrams!) Similarly, products of spectra are just formed as products on the underlying spaces, but coproducts are not just coproducts on the underlying spaces. For example, finite coproducts agree with finite products by stability, so $\Omega^\infty(X \oplus Y) = \Omega^\infty X \times \Omega^\infty Y$.

What we can say is this:

1. Ω^∞ turns (co)fiber sequences in Sp into *fiber* sequences in $\mathcal{S}_*$, and more generally pushout/pullback squares into pullback squares.

2. Σ^∞ turns cofiber sequences in $\mathcal{S}_*$ (sequences of the form $A \to X \to X/A$) into (co)fiber sequences in Sp, and more generally pushout squares into pushout/pullback squares. This can be seen in different ways, but the easiest is to use that Σ^∞ is a left adjoint functor, hence preserves arbitrary colimits.

A consequence of the first fact is that for a (co)fiber sequence of spectra, we have a fiber sequence on underlying spaces, hence a long exact sequence of homotopy groups. More generally, for a pushout/pullback square

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ Z & \longrightarrow & W \end{array}$$

of spectra, we get a pullback square on underlying spaces and hence a long exact ("Mayer-Vietoris") sequence

$$\ldots \to \pi_n(X) \to \pi_n(Y) \oplus \pi_n(Z) \to \pi_n(W) \to \pi_{n-1}(X) \to \ldots$$

3.4 Mapping spectra

Since $\Omega : \mathrm{Sp} \to \mathrm{Sp}$ is constructed as a limit, $\mathrm{Map}_{\mathrm{Sp}}(X, \Omega Y) \simeq \Omega \, \mathrm{Map}_{\mathrm{Sp}}(X, Y)$ (where the latter Ω is the one in $\mathcal{S}_*$). Since Ω is invertible, this means that the spaces

$$\mathrm{Map}_{\mathrm{Sp}}(X, \Omega^{-n} Y)$$

form a sequence of deloopings of $\mathrm{Map}_{\mathrm{Sp}}(X, Y)$, i.e. form a spectrum $\mathrm{map}_{\mathrm{Sp}}(X, Y)$ whose underlying space is $\mathrm{Map}_{\mathrm{Sp}}(X, Y)$. More precisely, we have:

Lemma 3.13. *If $\mathcal{C}$ is a stable ∞-category, the mapping space functor* $\mathrm{Map}_{\mathcal{C}} : \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{S}_*$ *factors through* $\mathrm{map}_{\mathcal{C}} : \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathrm{Sp}$. *We have*

$$\Omega^{\infty - n} \, \mathrm{map}_{\mathcal{C}}(X, Y) = \mathrm{Map}_{\mathcal{C}}(X, \Sigma^n Y).$$

Example 3.14. We have

$$\begin{aligned}\pi_k \operatorname{map}_{\mathrm{Sp}}(\Sigma^\infty X, E) &\simeq \pi_0 \Omega^{\infty+k} \operatorname{map}_{\mathrm{Sp}}(\Sigma^\infty X, E) \\ &\simeq \pi_0 \operatorname{Map}_{\mathrm{Sp}}(\Sigma^\infty X, \Sigma^{-k} E) \qquad \simeq \pi_0 \operatorname{Map}_{\mathrm{Sp}}(X, \Omega^{\infty+k} E) \simeq \widetilde{E}^{-k}(X)\end{aligned}$$

meaning *the homotopy groups of* $\operatorname{map}(\Sigma^\infty X, E)$ *are the* E*-cohomology groups of the space* X. More generally, we may define $E^*(Y)$ for a spectrum Y as $\pi_{-*} \operatorname{map}_{\mathrm{Sp}}(Y, E)$, with this definition $E^*(\Sigma^\infty X) = \widetilde{E}^*(X)$.

We may also write $\pi_k \operatorname{map}_{\mathrm{Sp}}(X, Y)$ for spectra X, Y as $[X, \Sigma^{-k} Y]$ or $[\Sigma^k X, Y]$, i.e. homotopy classes up to a shift.

From universal properties, one can prove that map is *exact* in both arguments:

Lemma 3.15. *If* $X \to Y \to Z$ *is a (co)fiber sequence of spectra, and* W *some other spectrum, then*

1. $\operatorname{map}(W, X) \to \operatorname{map}(W, Y) \to \operatorname{map}(W, Z)$ *is a (co)fiber sequence,*
2. $\operatorname{map}(Z, W) \to \operatorname{map}(Y, W) \to \operatorname{map}(X, W)$ *is a (co)fiber sequence.*

Proof. We check the second part, the first is completely analogous. By functoriality, we definitely have a sequence and a nullhomotopy, in particular we have an induced map

$$\operatorname{map}(Z, W) \to \operatorname{fib}(\operatorname{map}(Y, W) \to \operatorname{map}(X, W)),$$

which we need to check is an equivalence. A map of spectra is an equivalence if it is so on all underlying spaces, so it suffices to apply $\Omega^{\infty-n}$. So it suffices to check that

$$\Omega^{\infty-n} \operatorname{map}(Z, W) \to \Omega^{\infty-n} \operatorname{map}(Y, W) \to \Omega^{\infty-n} \operatorname{map}(X, W)$$

is a fiber sequence in $\mathcal{S}_*$. Since $\Omega^{\infty-n} \operatorname{map}(Z, W) \simeq \operatorname{Map}_{\mathrm{Sp}}(Z, \Sigma^n W)$, this sequence is obtained by applying $\operatorname{Map}_{\mathrm{Sp}}(-, \Sigma^n W)$ to the pushout square

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & Z, \end{array}$$

but $\operatorname{Map}_{\mathrm{Sp}}(-, \Sigma^n W)$ turns pushouts into pullbacks by universal properties, so we do in fact obtain a fiber sequence. □

Since $\Sigma^\infty A \to \Sigma^\infty X \to \Sigma^\infty(X/A)$ is a (co)fiber sequence, and gives another (co)fiber sequence after applying $\operatorname{map}(-, E)$, the resulting long exact sequence on homotopy groups recovers the long exact sequence for $\widetilde{E}^*(-)$.

3.5 Spectra and homology theories

As explained above, $E^*(X) = \pi_{-*}(\mathrm{map}(X, E))$ gives a very satisfying explanation of the long exact sequence in cohomology (and even generalizes it to spectrum input, instead of just pointed spaces). Can we similarly write *homology groups* as homotopy groups of a spectrum? Since they should be covariant, it is tempting to write $\pi_* \mathrm{map}(E, X)$, but that has the wrong behavior. For example, $E = HA$ should still correspond to singular homology, but for example $[H\mathbb{F}_p, \mathbb{S}] = 0$ turns out to be different from $\widetilde{H}_0(S^0; \mathbb{F}_p) = \mathbb{F}_p$ (even though the former is quite difficult to show!).

The correct construction requires another piece of structure on Sp:

Lemma 3.16. *There exists a unique (in the ∞-categorical sense) functor $\otimes$: $\mathrm{Sp} \times \mathrm{Sp} \to \mathrm{Sp}$ which preserves colimits in both variables, and makes the diagram*

$$\begin{array}{ccc} \mathcal{S}_* \times \mathcal{S}_* & \xrightarrow{\wedge} & \mathcal{S}_* \\ \big\downarrow{\scriptstyle \Sigma^\infty \times \Sigma^\infty} & & \big\downarrow{\scriptstyle \Sigma^\infty} \\ \mathrm{Sp} \times \mathrm{Sp} & \xrightarrow{\otimes} & \mathrm{Sp} \end{array}$$

This is not the strongest one can say, for example one would want this $\otimes$ to satisfy some form of associativity and commutativity, and indeed one can make precise what it means to have a symmetric-monoidal structure on an ∞-category, and gets a stronger version of the above Lemma. For our purposes, the above is enough for now.

The underlying space of $X \otimes Y$ is difficult to describe in terms of the underlying spaces of X and Y (a little bit like the underlying set of a tensor product of abelian groups is difficult to describe directly), but from the above lemma one can extract a number of very helpful properties of $\otimes$:

1. $\otimes$ preserves (co)fiber sequences in both variables separately, so if $X \to Y \to Z$ is a (co)fiber sequence and W some other spectrum, then also

$$X \otimes W \to Y \otimes W \to Z \otimes W$$

   is a (co)fiber sequence, and analogously with W on the other side.

2. $\mathbb{S}$ acts as unit for $\otimes$, that is we have natural equivalences $\mathbb{S} \otimes X \simeq X$ and $X \otimes \mathbb{S} \simeq X$.

Definition 3.17. For a spectrum E and a pointed space X, we write

$$\widetilde{E}_n(X) := \pi_n(E \otimes \Sigma^\infty X),$$

and more generally $E_n(X) = \pi_n(E \otimes X)$ for a spectrum X.

Lemma 3.18. *$\widetilde{E}_*(X)$ is a homology theory, with "coefficients" $\widetilde{E}_*(S^0) \cong \pi_*(E)$.*

Proof. Homotopy invariance is clear. Additivity follows from the fact that $E\otimes -$ preserves coproducts and that $\pi_n : \mathrm{Sp} \to \mathrm{Ab}$ preserves coproducts as well. The latter follows for finite coproducts from the fact that they are also products, and $[\mathbb{S}^n, -]$ preserves those by universal properties. For general coproducts it follows by an argument with filtered colimits, using that filtered colimits of spectra are formed underlying ($\Omega^{\infty+n}$ preserves filtered colimits), that $\pi_0 : \mathcal{S}_* \to \mathrm{Set}$ preserves them, and that they are also formed underlying in Ab ($\mathrm{Ab} \to \mathrm{Set}$ preserves and detects them).

The long exact sequence arises from the long exact sequence in homotopy associated to the (co)fiber sequence

$$E \otimes \Sigma^\infty A \to E \otimes \Sigma^\infty X \to E \otimes \Sigma^\infty (X/A).$$

Finally, $\widetilde{E}_*(S^0) \simeq \pi_*(E \otimes \mathbb{S}) \simeq \pi_*(E)$. □

Example 3.19. 1. HA gives a homology theory which satisfies all the Eilenberg-Steenrod axioms (including the dimension axiom), so it coincides with $\widetilde{H}_*(-; A)$.

2. $\mathbb{S}$ somewhat tautologically gives $\mathbb{S}_*(X) = \pi_*(\Sigma^\infty X)$, that is, the stable homotopy groups of spaces form a generalized homology theory.

Contrary to cohomology theories, the question of representing an abstract homology theory by a spectrum is much more subtle.

We summarize the results of this section as follows:

1. Sp is a stable ∞-category, which means among other things that $[X, Y]$ has an abelian group structure, that finite coproducts and finite products coincide, and that cofiber sequences and fiber sequences coincide.

2. Spectra correspond to (generalized) cohomology theories via $\pi_* \operatorname{map}(-, E)$ and also give rise to (generalized) homology theories via $\pi_*(E \otimes -)$.

3. Some spectra also arise as suspension spectra from spaces, via $\Sigma^\infty : \mathcal{S}_* \to \mathrm{Sp}$. This gives rise to a spectrum $\Sigma^\infty X$ whose E-(co)homology coincides with the E-(co)homology of the space X, and whose homotopy is given by the stable homotopy groups of X.

4. Cofiber sequences of spaces give rise to (co)fiber sequences on suspension spectra, which lead to long exact sequences on all those things.

4 The t-structure on spectra

As with unstable homotopy theory, homotopy groups control everything:

Theorem 4.1 ("Whitehead for spectra")**.** *If a map $X \to Y$ of spectra induces isomorphisms on each π_n, it is an equivalence.*

Proof. It suffices to check that for each n, $\Omega^{\infty-n}X \to \Omega^{\infty-n}Y$ is an equivalence. This follows from the usual Whitehead theorem, but we need to be a bit careful, since that only applies to connected spaces. So write $(\Omega^{\infty-n-1}X)_0$ for the basepoint component of $\Omega^{\infty-n-1}X$, then by the usual Whitehead theorem

$$(\Omega^{\infty-n-1}X)_0 \to (\Omega^{\infty-n-1}Y)_0$$

is an equivalence (since their π_i for $i \geq 1$ is just π_{i+n+1} of the respective spectra), but after applying Ω one more time, we have

$$\Omega^{\infty-n}X \to \Omega^{\infty-n}Y,$$

since Ω only depends on the basepoint component anyways. This finishes the proof. $\square$

So it makes sense to try and organize spectra by their homotopy groups. We have $\pi_k(\mathbb{S}) = 0$ for $k < 0$ and $\pi_0(\mathbb{S}) \cong \mathbb{Z}$, since $\pi_{m+k}(S^m) = 0$ for $n < 0$ and $\pi_m(S^m) \cong \mathbb{Z}$ for $m \geq 1$.

Definition 4.2.

1. A spectrum X is *connective* if $\pi_k(X) = 0$ for $k < 0$. More generally, X is *n-connective* if $\pi_k(X) = 0$ for $k < n$.
2. A spectrum X is *coconnective* if $\pi_k(X) = 0$ for $k > 0$. More generally, X is *n-coconnective* if $\pi_k(X) = 0$ for $k > n$.

So the Eilenberg-MacLane spectrum HA is both connective and coconnective. $\mathbb{S}$ is connective (but not coconnective), and more generally any suspension spectrum is connective.

Objects of $\mathcal{S}$ have cell structures. We can analogously speak of cell structures in Sp:

Definition 4.3. A cell structure on a spectrum X consists of a sequence X_n ($n \in \mathbb{Z}$) with $X_n = 0$ for n small enough, together with sets Γ_n and pushout squares

$$\begin{array}{ccc} \bigoplus_{\Gamma_n} \mathbb{S}^{n-1} & \longrightarrow & X_{n-1} \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & X_n, \end{array}$$

as well as an equivalence $\operatorname{colim}_{n\in\mathbb{N}} X_n \simeq X$.

(Here, we're officially retiring the notation where a subscript on a spectrum refers to the underlying space, in order to free it up for talking about sequences of spectra.)

Note that this is completely analogous to cell structures on spaces, with 0 replacing the contractible D^n (as we're taking homotopy pushouts anyways). With a little bit of care regarding basepoints, one can show that a cell structure on a space X gives rise to a cell structure on $\Sigma^\infty X$, but we can prove something stronger directly:

Lemma 4.4. *Let X be k-connective for some k. Then X has a cell structure with cells only in degrees $\geq k$.*

Proof. The proof is essentially the same as for CW approximation, but stability lets us write it a little bit more elegantly. Instead of trying to attach cells to 0 to approximate X, let us try to first attach cells to X to make it 0.

Specifically, start with $Y_k = X$, and then build a sequence of Y_n with each Y_n n-connective. In each step, we choose a set Γ_n of generators of $\pi_n(Y_n)$, and represent those by maps $\mathbb{S}^n \to Y_n$. The universal property of the coproduct gives a map $\bigoplus_{\Gamma_n} \mathbb{S}^n \to Y_n$ which on π_n induces the surjective map

$$\bigoplus_{\Gamma_n} \mathbb{Z} \to \pi_n(Y_n).$$

So if we let Y_{n+1} be the cofiber, the long exact sequence shows that it is $(n+1)$-connective, and we continue inductively. Since homotopy groups commute with sequential colimits, $\operatorname{colim}_n Y_n$ has trivial homotopy groups and is 0.

We now define spaces $X_n = \operatorname{fib}(X \to Y_{n+1})$. A useful lemma in stable ∞-categories (reminiscent of a lemma in homological algebra) says that given a natural transformation of (co)fiber sequences, the sequence obtained by taking fibers (or cofibers) is also a (co)fiber sequence. We may apply this to the bottom two rows of the following diagram to obtain the cofiber sequence at the top:

$$\begin{array}{ccccc}
\bigoplus_{\Gamma_n} \mathbb{S}^{n-1} & \longrightarrow & X_{n-1} & \longrightarrow & X_n \\
\downarrow & & \downarrow & & \downarrow \\
0 & \longrightarrow & X & \longrightarrow & X \\
\downarrow & & \downarrow & & \downarrow \\
\bigoplus_{\Gamma_n} \mathbb{S}^{n} & \longrightarrow & Y_n & \longrightarrow & Y_{n+1}.
\end{array}$$

Also colimits of (co)fiber sequences are still (co)fiber sequences (since we can express (co)fiber sequences in terms of pushouts and these commute with other types of colimits), and applying this to the (co)fiber sequence $X_n \to X \to Y_{n+1}$ gives a (co)fiber sequence $\operatorname{colim}_n X_n \to X \to 0$, hence an equivalence $\operatorname{colim}_n X_n \simeq X$. Together, this gives a cell structure on X. □

This is the key step in the following important lemma:

Lemma 4.5. *For a spectrum X, the following are equivalent:*

1. *X is k-connective.*

2. *X admits a cell structure with cells in degrees $\geq k$.*

3. *X is contained in the smallest full subcategory of* Sp *which contains $\mathbb{S}^n$ for $n \geq k$ and is closed under colimits and* extensions*: If $A \to B \to C$ is a cofiber sequence, and A, C are contained in this subcategory, then also B is.*

Proof. We have just done the step 1⇒2. For 2⇒3 it is enough to observe that in a cell structure, X_n is an extension of X_{n-1} by

$$X_{n-1} \to X_n \to \bigoplus_{\Gamma_n} \mathbb{S}^n,$$

by "rotating" the cofiber sequence $\bigoplus_{\Gamma_n} \mathbb{S}^{n-1} \to X_{n-1} \to X_n$ (passing to the cofiber on the right gives the suspension of the object on the left). The sums of spheres that appear here are built from $\mathbb{S}^n$ for $n \geq k$ under coproducts (a type of colimit), and X is obtained from the X_n by another colimit.

For the 3⇒1 step, we need to use that all colimits can be decomposed into a sequence of coproducts and pushouts (we have seen this explicitly for colimits over the poset $\mathbb{N}$ earlier.) Taking this for granted, we need to check that the k-connective objects contain the spheres $\mathbb{S}^n$ for $n \geq k$ (clear from connectivity of $\mathbb{S}$), are closed coproducts and pushouts (clear from compatibility of homotopy groups with arbitrary coproducts and the long exact sequence) and extensions (also clear from the long exact sequence). □

We are now ready to state and prove the main feature of the notions of connectivity and coconnectivity.

Definition 4.6. Write $\mathrm{Sp}_{\geq k}$ for the full subcategory of Sp on k-connective spectra, and $\mathrm{Sp}_{\leq k}$ for the full subcategory of k-coconnective spectra.

Definition 4.7. A *t-structure* on a stable ∞-category $\mathcal{C}$ consists of subcategories $\mathcal{C}_{\geq k}, \mathcal{C}_{\leq k}$ for each k with the following properties:

1. $\mathcal{C}_{\geq k+1} \subseteq \mathcal{C}_{\geq k}$, $\mathcal{C}_{\leq k-1} \subseteq \mathcal{C}_{\leq k}$, and Σ gives equivalences $\mathcal{C}_{\geq k} \to \mathcal{C}_{\geq k+1}$ and $\mathcal{C}_{\leq k} \to \mathcal{C}_{\leq k+1}$.
2. For $X \in \mathcal{C}_{\geq k}$ and $Y \in \mathcal{C}_{\leq k-1}$, we have $[X, Y] \simeq 0$, and in fact $\mathrm{Map}(X, Y) \simeq \mathrm{pt}$.
3. Conversely, if $Z \in \mathcal{C}$ satisfies $[X, Z] \simeq 0$ for all $X \in \mathcal{C}_{\geq k}$, then $Z \in \mathcal{C}_{\leq k-1}$, and if $Z \in \mathcal{C}$ satisfies $[Z, Y] \simeq 0$ for all $Y \in \mathcal{C}_{\leq k-1}$, $Z \in \mathcal{C}_{\geq k}$.
4. For every Z and k, there exists a (co)fiber sequence $X \to Z \to Y$ with $X \in \mathcal{C}_{\geq k}$ and $Y \in \mathcal{C}_{\leq k-1}$.

(These are of course a little redundant, for example the first property allows us to reduce the other ones to the case $k = 0$.)

Lemma 4.8. *The subcategories* $\mathrm{Sp}_{\geq k}$, $\mathrm{Sp}_{\leq k}$ *define a t-structure on* Sp.

Proof. Since Σ and its inverse Ω just shift homotopy groups around, and $\mathrm{Sp}_{\geq k}$ and $\mathrm{Sp}_{\leq k}$ are characterized in terms of those, the first statement is clear.

For the second statement, observe that since all the $\Sigma^i X$ for $i \geq 0$ are also in $\mathrm{Sp}_{\geq k}$, the second statement is equivalent to the a priori stronger statement that $\mathrm{map}(X, Y)$ has trivial homotopy groups in degrees ≥ 0 (as these are given by $[X, Y]$), and this is equivalent to $\mathrm{Map}_{\mathrm{Sp}}(X, Y)$ being contractible (by Whitehead). The former version is clearly closed under extensions in X, the latter under colimits. So it suffices to check that the class of all X with $[-, Y] \simeq 0$ contains the spheres $\mathbb{S}^n$ with $n \geq k$, but this is clear since $\pi_n(Y) = 0$ by definition of coconnectivity.

We prove the third statement last and first check the fourth. Given Z, we may (as in the proof of the existence of cell structures above) form a sequence of Y_n with $\pi_i(Y_n) = 0$ for $i = k, \ldots, n-1$, starting with $Y_k = Z$, where Y_{n+1} arises from Y_n by taking a cofiber of a map $\bigoplus_{\Gamma_n} \mathbb{S}^n \to Y_n$ surjective on π_n. The colimit Y then has the property that it is $k-1$-coconnective and $Z \to Y$ induces an isomorphism on π_i for $i \leq k-1$. Letting X be the fiber, X is k-connective by the long exact sequence.

Finally, we prove the third statement. Assume Z satisfies $[X, Z] \simeq 0$ for all $X \in \mathrm{Sp}_{\geq k}$. Using the fourth statement, choose a cofiber sequence

$$X \to Z \to Y$$

with X k-connective and Y $(k-1)$-coconnective. By assumption, the map $X \to Z$ is nullhomotopic. But by the long exact sequence it also induces isomorphisms in π_i for $i \geq k$. So these homotopy groups vanish and Z itself is $k-1$-coconnective. The other statement works analogously. $\square$

What makes t-structures useful is that the decomposition $X \to Z \to Y$ is actually functorial in Z:

Proposition 4.9. *There exist functors*

1. $\tau_{\geq n} : \mathrm{Sp} \to \mathrm{Sp}_{\geq n}$, *right adjoint to the inclusion,*
2. $\tau_{\leq n} : \mathrm{Sp} \to \mathrm{Sp}_{\leq n}$, *left adjoint to the inclusion,*

and natural (co)fiber sequences

$$\tau_{\geq n} Z \to Z \to \tau_{\leq n-1} Z.$$

Proof. For any Z, we have a (co)fiber sequence

$$X \to Z \to Y$$

with $X \in \mathrm{Sp}_{\geq n}$ and $Y \in \mathrm{Sp}_{\leq n-1}$. For any $W \in \mathrm{Sp}_{\geq n}$, applying $\mathrm{Map}(W, -)$ we obtain a fiber sequence

$$\mathrm{Map}(W, X) \to \mathrm{Map}(W, Z) \to \mathrm{Map}(W, Y),$$

of spaces, and since the rightmost term is contractible, we learn that

$$\mathrm{Map}(W, X) \simeq \mathrm{Map}(W, Z).$$

If we had such a right adjoint $\tau_{\geq n}$, then we would also have

$$\mathrm{Map}(W, \tau_{\geq n} Z) \simeq \mathrm{Map}(W, Z),$$

and so by the Yoneda lemma X and $\tau_{\geq n}Z$ would be equivalent. By reversing the logic a bit, one can use the Yoneda lemma to prove a standard fact (the "pointwise criterion for adjoints") which says that existence of such an X as above for each Z already guarantees existence of an adjoint functor. The other direction works similarly by mapping into an $n-1$-coconnective W.

(This is similar to how the universal property of limits or colimits characterizes them pointwise for each diagram, but they end up functorial in the diagram purely by universal properties.) □

Since an $(n-1)$-coconnective spectrum is in particular n-coconnective, the adjunction gives natural transformations $\tau_{\leq n} \to \tau_{\leq n-1}$ (and dually $\tau_{\geq n} \to \tau_{\geq n-1}$). We may also combine both truncations and write $\tau_{[a,b]} = \tau_{\leq b}\tau_{\geq a}X \simeq \tau_{\geq a}\tau_{\leq b}X$.

For a spectrum Z, one can think of the ($\mathbb{Z}$-indexed) sequence

$$\to \tau_{\leq n}Z \to \tau_{\leq n-1}Z \to \dots$$

as a stable version of the *Postnikov tower* of a space. The dual sequence

$$\to \tau_{\geq n}Z \to \tau_{\geq n-1}Z \to \dots$$

is called the *Whitehead tower*.

Proposition 4.10. *We have natural equivalences* $\operatorname{colim}_n \tau_{\geq n}X \simeq X$ *and* $X \simeq \lim_n \tau_{\leq n}X$*, that is a spectrum can be recovered both as colimit of its Whitehead tower and the limit of its Postnikov tower.*

Proof. The cofiber of the map $\operatorname{colim}_n \tau_{\geq n}X \to X$ agrees with $\operatorname{colim}_n \tau_{\leq n-1}X$ because colimits are exact. The first claim is equivalent to that spectrum being trivial. Since homotopy groups commute with filtered colimits, and in that diagram all homotopy groups eventually become zero, this follows from Whitehead.

The other statement analogously amounts to $\lim_n \tau_{\geq n+1}X$ vanishing. Homotopy groups do not commute with sequential limits, but almost: One can write this limit (over $\mathbb{N}$) as fiber of a map

$$\prod_n \tau_{\geq n+1}X \to \prod_n \tau_{\geq n+1}X,$$

(related to the expression of sequential colimits as pushouts in the proof of Lemma 3.9), and the resulting long exact sequence allows us to express π_n of a sequential limit in terms of π_n and π_{n+1} of the diagram, in a way where eventual vanishing of these groups still gives 0. □

We will now try to use this to understand how a spectrum is determined by its homotopy groups. By the Whitehead theorem for spectra, $X \simeq 0$ if all homotopy groups vanish. Next, we analyze the case with exactly one nonzero homotopy group.

Proposition 4.11. *Assume X is a spectrum which is n-coconnective and n-connective, and $\pi_n(X) = A$. Then $X \simeq \Sigma^n HA$.*

Proof. Applying a shift, we may assume $n = 0$ to simplify notation. Choose a 2-stage free resolution of the abelian group A, i.e. an exact sequence

$$0 \to F_1 \to F_0 \to A \to 0,$$

where $F_0 = \bigoplus_{\Gamma_0} \mathbb{Z}$ and $F_1 = \bigoplus_{\Gamma_1} \mathbb{Z}$ are free. Build $Z_0 = \bigoplus_{\Gamma_0} \mathbb{S}$ and observe that

$$\pi_0(Z_0) \cong \bigoplus_{\Gamma_0} \pi_0(\mathbb{S}) \cong F_0.$$

For $Z_1 = \bigoplus_{\Gamma_1} \mathbb{S}$, the universal property of coproducts allows us to identify a homotopy class $[Z_1, Y]$ with a Γ_1-tuple of elements of $\pi_0(Y)$, and hence gives a map

$$Z_1 \to Z_0$$

which on π_0 induces the map $F_1 \to F_0$.

Similarly, under the identification $\pi_0(X) \cong A$, we may choose a map $Z_0 \to X$ which on π_0 induces the map $F_0 \to A$. The composite of these maps $Z_1 \to Z_0 \to X$ induces the zero map on π_0 and (again because Z_1 is a sum of spheres) is thus nullhomotopic. We get a square

$$\begin{array}{ccc} Z_1 & \longrightarrow & Z_0 \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & X, \end{array}$$

and by the universal property of pushouts a map $\operatorname{cofib}(Z_1 \to Z_0) \to X$, which is an isomorphism on π_0 by construction. Since the target is 0-coconnective, it factors through a map

$$\tau_{\leq 0} \operatorname{cofib}(Z_1 \to Z_0) \to X,$$

which is an equivalence, because now it is an isomorphism on all homotopy groups.

What this argument shows is that *any* spectrum whose homotopy groups are concentrated in degree 0 and given by A there is actually equivalent to $\tau_{\leq 0} \operatorname{cofib}(Z_1 \to Z_0)$. In particular, this also applies to HA as defined earlier. □

In fact, the same strategy of proof also shows something about morphisms:

Proposition 4.12. $[HA, HB] \cong \operatorname{Hom}(A, B)$.

Proof. As before, we may pick a free resolution of A and identify HA with

$$\tau_{\leq 0} \operatorname{cofib}(Z_1 \to Z_0)$$

where the Z_i are the corresponding sums of spheres. From the adjunction,

$$[\tau_{\leq 0} \operatorname{cofib}(Z_1 \to Z_0), HB] \cong [\operatorname{cofib}(Z_1 \to Z_0), HB],$$

which sits in a long exact sequence

$$\ldots \to [\Sigma Z_1, HB] \to [\mathrm{cofib}(Z_1 \to Z_0), HB] \to [Z_1, HB] \to [Z_0, HB] \to \ldots.$$

Since Z_i is a sum of spheres indexed over generators of F_i, this simplifies to

$$0 \to [\mathrm{cofib}(Z_1 \to Z_0), HB] \to \mathrm{Hom}(F_1, B) \to \mathrm{Hom}(F_0, B) \to \ldots,$$

but the kernel of the middle morphism here is $\mathrm{Hom}(A, B)$ by left exactness of Hom, proving the claim. □

Corollary 4.13. *The functor* $\pi_0 : \mathrm{Sp} \to \mathrm{Ab}$ *restricts to an equivalence on the full subcategory* $\mathrm{Sp}_{\leq 0} \cap \mathrm{Sp}_{\geq 0}$. *Its inverse gives a fully faithful functor* $H : \mathrm{Ab} \to \mathrm{Sp}$ *taking an abelian group* A *to its Eilenberg-MacLane spectrum* HA.

In the language of t-structures, $\mathrm{Sp}_{\leq 0} \cap \mathrm{Sp}_{\geq 0}$ is the *heart* of the t-structure, $\mathrm{Sp}^{\heartsuit}$. The heart of a t-structure on a stable ∞-category is always an abelian (1-)category, and we have just shown that for Sp it is equivalent to Ab.

Remark 4.14. This allows us to really think of the 1-category of abelian groups as a full subcategory of the ∞-category of spectra Sp, i.e. of abelian groups as a special case of spectra, and spectra as a generalization of abelian groups, in a similar way like we may think of Set as a full subcategory of $\mathcal{S}$. Taking this perspective seriously, we will opt to usually drop the H from notation and for example write things like $\mathrm{map}(\mathbb{S}, \mathbb{Z})$ instead of $\mathrm{map}(\mathbb{S}, H\mathbb{Z})$.

We now observe that the fibers in the Postnikov tower $\mathrm{fib}(\tau_{\leq n}X \to \tau_{\leq n-1}X)$, have only a single homotopy group given by $\pi_n(X)$ in degree n. So we have fiber sequences

$$\Sigma^n \pi_n(X) \to \tau_{\leq n}X \to \tau_{\leq n-1}X.$$

Rotating the fiber sequence, we have a fiber sequence

$$\tau_{\leq n}X \to \tau_{\leq n-1}X \to \Sigma^{n+1}\pi_n(X)$$

If X is connective, $\tau_{\leq 0}X \simeq \pi_0(X)$ is an Eilenberg-MacLane space. The next term is given as a fiber

$$\tau_{\leq 1}X \to \pi_0(X) \to \Sigma^2\pi_1(X),$$

then

$$\tau_{\leq 2}X \to \tau_{\leq 1}X \to \Sigma^3\pi_2(X),$$

etc. So we can recover all the $\tau_{\leq n}X$ from the homotopy groups together with maps $\tau_{\leq n}X \to \Sigma^{n+2}\pi_{n+1}(X)$. We may think of the latter data also as cohomology classes in $H^{n+2}(\tau_{\leq n}X; \pi_{n+1}(X))$. These are known as the *k-invariants* of the spectrum X.

Remark 4.15. Instead of in an inductive description like this, a perspective of *coherent chain complexes* developed by Ariotta in [5] applied to the Postnikov tower allows us to also think of this data directly as consisting of

1. abelian groups $\pi_n(X)$
2. maps $\Sigma^{2n}\pi_n(X) \to \Sigma^{2n+2}\pi_{n+1}(X)$ (coming up to a shift from the Postnikov tower of $\tau_{[n,n+1]}(X)$)
3. nullhomotopies for the length 2 composites $\Sigma^{2n}\pi_n(X) \to \Sigma^{2n+4}\pi_{n+2}(X)$
4. homotopies between the two different nullhomotopies we obtain for the length 3 composites $\Sigma^{2n}\pi_n(X) \to \Sigma^{2n+6}\pi_{n+3}(X)$.
5. ...

that is, a spectrum is fully captured by a sort of homotopy-coherent "cochain complex" whose n-th term is $\Sigma^{2n}\pi_n(X)$.

The additional data here is contained in maps $HA \to \Sigma^2 HB$ (and the higher coherences also in spaces of the form $\mathrm{Map}(HA, \Sigma^n HB)$). While Corollary 4.13 tells us that $\mathrm{Map}(HA, HB)$ is equivalent to the discrete space $\mathrm{Hom}(A, B)$, and so the positive degree homotopy groups of $\mathrm{map}(HA, HB)$ vanish, $\mathrm{Map}(HA, \Sigma^n HB)$ also sees some of the *negative* degree homotopy groups of $\mathrm{map}(HA, HB)$, which we don't know anything about at this point.

Remark 4.16. To see that a stable ∞-category with t-structure is not determined by its heart, consider the *derived category* $\mathcal{D}(\mathbb{Z})$. This can be described in different ways, the easiest being as homotopy-coherent nerve of a category whose objects are chain complexes of free abelian groups, with mapping spaces defined in terms of chain homotopies. This is stable, with Σ and Ω corresponding to just an index shift of complexes. It has a t-structure with $\mathcal{D}(\mathbb{Z})_{\geq k}$ the full subcategory of complexes with homology concentrated in degrees $\geq k$, and $\mathcal{D}(\mathbb{Z})_{\leq k}$ analogously. Here, the heart is also identified with Ab (unwinding this statement gives the fundamental lemma of homological algebra), but $\mathcal{D}(\mathbb{Z})$ is not equivalent to Sp. The difference between the two is that for abelian groups A, B,

$$[A, \Sigma^n B]_{\mathcal{D}(\mathbb{Z})} \cong \mathrm{Ext}^n(A, B)$$

which is 0 if $n \neq 0, 1$. As we will see, $[A, \Sigma^n B]_{\mathrm{Sp}}$ is much bigger!

Finally, we want to discuss how the t-structure interacts with map and $\otimes$:

Proposition 4.17. *1. If X is an n-connective spectrum and Y is an m-connective spectrum, then $X \otimes Y$ is $n+m$-connective.*

2. If X is an n-connective spectrum and Y is m-coconnective, then $\mathrm{map}(X, Y)$ *is $m-n$-coconnective.*

Proof. Since $\otimes$ commutes with Σ in both variables, and $\mathrm{map}(\Sigma X, Y) \simeq \Sigma^{-1}\mathrm{map}(X, Y)$, $\mathrm{map}(X, \Sigma Y) \simeq \Sigma\,\mathrm{map}(X, Y)$, we may apply shifts to reduce both statements to the cases $n = m = 0$. For fixed connective Y, the class of all X with $X \otimes Y$ connective clearly contains all $\mathbb{S}^i$ with $i \geq 0$ since $\mathbb{S}^i \otimes Y \simeq \Sigma^i Y$, is closed under extensions, and colimits. So it contains all connective X. Similarly, for fixed coconnective Y, the class of all X with $\mathrm{map}(X,Y)$ coconnective clearly contains all $\mathbb{S}^i$ since $\mathrm{map}(\mathbb{S}^i,Y) \simeq \Sigma^{-i}Y$, is closed under extensions, and colimits. The latter follows because we have $\mathrm{map}(\mathrm{colim}_I X_i, Y) \simeq \lim_I \mathrm{map}(X_i, Y)$ and coconnective objects are closed under limits, because they are characterized by maps from all connective objects into them being contractible. So, $\mathrm{map}(X,Y)$ is coconnective for all connective X. □

Example 4.18 (Hurewicz theorem)**.** We have a (co)fiber sequence

$$\tau_{\geq 1}\mathbb{S} \to \mathbb{S} \to \mathbb{Z}.$$

Tensoring with an n-connective spectrum X, we get a (co)fiber sequence

$$(\tau_{\geq 1}\mathbb{S}) \otimes X \to X \to \mathbb{Z} \otimes X,$$

with the left term being $n+1$-connective. The long exact sequence then shows that $\pi_n(X) \to \pi_n(\mathbb{Z} \otimes X) \cong H_n(X;\mathbb{Z})$ is an isomorphism.

More generally, for connective E, the same argument shows $E_n(X) \cong H_n(X;\pi_0(E))$ for n-connective X. We will soon see a more general relationship between generalized homology theories and usual singular homology of which this is a special case.

Corollary 4.19. *If $X \to Y$ is a map of bounded below spectra which induces an isomorphism on $H\mathbb{Z}$-homology, it is an equivalence.*

Proof. The cofiber is bounded below and has vanishing $H\mathbb{Z}$-homology. By Hurewicz, all of its homotopy groups are zero, so it vanishes. □

5 Filtered spectra and spectral sequences

We have seen that every spectrum can in a sense be cut up into abelian groups, but the gluing information lives in groups $[HA, \Sigma^n HB]$, i.e. in *cohomology of Eilenberg-MacLane spectra*. It turns out that this cohomology can be completely understood, and is the first stepping stone to more advanced computations in stable homotopy theory. Our next goal is to get a rough overview over how this is computed. The main tool for this (and all sorts of other computations in homotopy theory) are spectral sequences. We begin our motivation of these by considering the following object:

Definition 5.1. A *filtered spectrum* is a sequence of spectra

$$\ldots \to X_{n+1} \to X_n \to X_{n-1} \to \ldots.$$

We call $X_{-\infty} := \mathrm{colim}_n X_n$ the *underlying spectrum*, and call X_* *complete* if $\lim_n X_n \simeq 0$.

Filtered spectra form a category $\mathrm{FilSp} = \mathrm{Fun}(\mathbb{Z}^{\mathrm{op}}, \mathrm{Sp})$, where $\mathbb{Z}$ is viewed as a category using the partial order $\leq$.

For a filtered spectrum X, the spectra $\mathrm{gr}_n X := \mathrm{cofib}(X_{n+1} \to X_n)$ form the "associated graded spectrum" $\mathrm{gr}_* X$.

The basic question that will lead us to spectral sequences is the following: Suppose we understand the homotopy groups of the associated graded $\pi_*(\mathrm{gr}_* X)$, what can we say about the homotopy groups of the underlying spectrum $X_{-\infty} = \mathrm{colim}_n X_n$?

Consider first the case where $\mathrm{gr}_n X = 0$ for $n \neq 0$. This means that all the maps $X_{n+1} \to X_n$ are equivalences, unless $n = 0$. Completeness then means $0 \simeq \lim X_n \simeq X_1$, and the underlying spectrum $\mathrm{colim}_n X_n$ agrees with X_0. In that case, $\mathrm{gr}_0 X \simeq X_0 \simeq X_{-\infty}$ directly recovers the underlying spectrum.

If we have nonzero $\mathrm{gr}_1 X$ and $\mathrm{gr}_0 X$, a similar argument shows that $X_1 \simeq \mathrm{gr}_1 X$, and $X_0 \simeq X_{-\infty}$, and the definition of $\mathrm{gr}_0 X$ then yields a cofiber sequence

$$\mathrm{gr}_1 X \to X_{-\infty} \to \mathrm{gr}_0 X,$$

and now the relationship between $\pi_* X_{-\infty}$ and $\pi_*(\mathrm{gr}_* X)$ is given by a long exact sequence.

For more complicated associated graded, we have a more complicated relationship: A spectral sequence.

To spell this out, observe that from the (co)fiber sequences

$$X_{s+1} \to X_s \to \mathrm{gr}_s X$$

we get long exact sequences

$$\pi_n X_{s+1} \xrightarrow{i} \pi_n X_s \xrightarrow{p} \pi_n \mathrm{gr}_s X \xrightarrow{\partial} \pi_{n-1} X_{s+1} \to \dots,$$

which share some terms (each $\pi_n X_s$ appears in two of these). In particular, we may talk about iterates of the map i.

Definition 5.2. For a filtered spectrum X, we define $E^1_{n,s} := \pi_n(\mathrm{gr}_s(X))$, and subgroups

$$\begin{aligned} Z^r_{n,s} &= \{a \in E^1_{n,s} \mid \partial a \in \mathrm{im}(i^{r-1})\} \\ B^r_{n,s} &= \{a \in E^1_{n,s} \mid \exists b \text{ with } a = pb \text{ and } i^{r-1} b = 0\}. \end{aligned}$$

and $E^r_{n,s} = Z^r_{n,s}/B^r_{n,s}$.

The bigraded abelian group $E^r_{*,*}$ is referred to as the r'th page of the spectral sequence associated to the filtered spectrum.

It is clear that $B^r_{n,s} \subseteq B^{r+1}_{n,s}$, $Z^{r+1}_{n,s} \subseteq Z^r_{n,s}$, and all the B's are contained in all the Z's. In particular, $E^{r+1}_{n,s}$ is a subquotient (quotient of a subgroup) of $E^r_{n,s}$. Which subquotient it is is determined by a kind of higher connecting homomorphism

$$d_r : E^r_{n,s} \to E^r_{n-1,s+r}$$

explicitly defined by picking a representative $a \in Z^r_{n,s}$, choosing $b \in \pi_{n-1} X_{s+r}$ with $i^{r-1} b = \partial a$, and letting $d_r([a]) := [pb]$.

Proposition 5.3. *d_r is well-defined, and one has canonical isomorphisms $E^{r+1}_{n,s} \cong \ker(d_r)/\mathrm{im}(d_r)$. More specifically:*

1. $Z^{r+1}_{n,s} \subseteq Z^r_{n,s}$ *agrees with the preimage of* $\ker(d_r)$ *under* $Z^r_{n,s} \to E^r_{n,s}$.

2. $B^{r+1}_{n,s} \subseteq Z^r_{n,s}$ *agrees with the preimage of* $\operatorname{im}(d_r)$ *under* $Z^r_{n,s} \to E^r_{n,s}$.

Proof. This is a fun diagram chase, which we omit. □

So each page is the homology of the previous with respect to some differential.

Next, we analyze the relationship between the homotopy groups of the underlying spectrum $\pi_n(X_{-\infty})$ and the spectral sequence. As

$$\pi_n(X_{-\infty}) \cong \operatorname{colim}_s \pi_n(X_s),$$

it seems natural to organize elements in $\pi_n(X_{-\infty})$ by which X_s they come from:

Definition 5.4. For a filtered spectrum, the *abutment filtration* on $\pi_n(X_{-\infty})$ is given by

$$F^s\pi_n(X_{-\infty}) = \operatorname{im}(\pi_n(X_s) \to \pi_n(X_{-\infty})).$$

To make precise the relationship between the abutment filtration and the pages of the spectral sequence, we will make one additional simplifying assumption on the filtered spectrum.

Definition 5.5. Call a filtered spectrum X *locally finite* if it is complete and for every n, $\pi_n(\operatorname{gr}_s X) = 0$ both for large enough and small enough s.

Consequences include that the maps $\pi_n(X_{s+1}) \to \pi_n(X_s)$ become isomorphisms for small s and large s. An argument with completeness shows that then the $\pi_n(X_s)$ actually vanish for large s. We also have that for every n, $E^r_{n,s}$ is 0 for large enough and small enough s (and any r). As a consequence, $Z^r_{n,s}$, $B^r_{n,s}$ and $E^r_{n,s}$ stabilize for fixed (n, s) and large enough r. We write $Z^\infty_{n,s}$, $B^\infty_{n,s}$ and $E^\infty_{n,s}$ for the resulting groups.

Proposition 5.6. *If X is a locally finite filtered spectrum, the abutment filtration on each $\pi_n(X_{-\infty})$ is finite ($F^s\pi_n(X_{-\infty}) = 0$ for large s and $F^s\pi_n(X_{-\infty}) = \pi_n(X_{-\infty})$ for small s) and there are natural isomorphisms*

$$F^s\pi_n(X_{-\infty})/F^{s+1}\pi_n(X_{-\infty}) \cong E^\infty_{n,s}.$$

Proof. For an element $x \in F^s\pi_n(X_{-\infty})$, we may choose a preimage $y \in \pi_n(X_s)$. Two different preimages y, y' must become equal in some finite stage of the colimit, so there exists r with $i^{r-1}(y-y') = 0$. This means that $p(y-y') \in B^\infty_{n,s}$, and so the value of $p(y)$ in $E^1_{n,s}/B^\infty_{n,s}$ is well-defined. We also have $\partial p(y) = 0$, so $p(y)$ lies in $Z^\infty_{n,s}$, giving a homomorphism

$$F^s\pi_n(X_{-\infty}) \to E^\infty_{n,s}.$$

It is surjective, since an $a \in Z^\infty_{n,s}$ has the property that $\partial a \in \operatorname{im}(i^{r-1})$, i.e. a lifts arbitrarily far along the tower of i's. But since $\pi_{n+1}(X_{s+r})$ vanishes for large

r, this means that $\partial a = 0$ and so $a = p(y)$ for some y. This y maps to some element of $F^s\pi_n(X_{-\infty})$, so a lies in the image of the above homomorphism.

Finally, we check that the kernel is exactly $F^{s+1}\pi_n(X_{-\infty})$. Given an element in the kernel, it lifts to $y \in \pi_n(X_s)$ with $p(y) \in B^\infty_{n,s}$, say $B^r_{n,s}$. Then there exists $b \in \pi_n(X_s)$ with $p(b) = p(y)$ and $i^{r-1}(b) = 0$, and so $y - b$ is another representative of the same element of $F^s\pi_n(X_{-\infty})$ which is in the kernel of p. But exactness then tells us that there exists y' with $i(y') = y - b$, and so the original element of $F^s\pi_n(X_{-\infty})$ lies in $F^{s+1}\pi_n(X_{-\infty})$ as claimed. The fact that all of $F^{s+1}\pi_n(X_{-\infty})$ lies in the kernel follows similarly, but easier, from $p \circ i = 0$. □

To summarize, we have:

- For each filtered spectrum X constructed a spectral sequence $(E^r_{*,*}, d_r)$ with $E^1_{n,s} = \pi_n(\mathrm{gr}_s X)$.
- For a locally finite filtered spectrum proved *strong convergence* of the spectral sequence against $\pi_n(X_{-\infty})$: We provided a finite filtration on each of those groups, together with an identification of $E^\infty_{*,*}$ with the associated graded.

One usually writes $\pi_n(\mathrm{gr}_s X) \Rightarrow \pi_n(X_{-\infty})$ to indicate this relationship.

Remark 5.7. Of course, there are useful things to say beyond the locally finite case, and in fact some examples we will consider later are not locally finite. The best thing that can be said is Boardman's notion of *conditional convergence* [6], which holds always for the spectral sequence of a filtered spectrum. The main difficulty addressed by it is that in general, $Z^\infty_{n,s} = \bigcap_r Z^r_{n,s}$ has to be replaced by a derived limit, leading to a derived E^∞ page and a derived version of the abutment filtration. Boardman then gives explicitly checkable conditions under which conditional convergence degenerates to strong convergence, with an ordinary abutment filtration (which is in general not finite, but complete).

5.1 The Atiyah-Hirzebruch spectral sequence

Almost all spectral sequences arise from filtered spectra (or filtered objects in a more general stable ∞-category). A lot of useful examples arise simply from constructing a filtration on a spectrum where the associated graded can be interpreted in some other way.

As a first example, consider the situation of a generalized cohomology $E^*(X)$. We interpreted these groups as homotopy groups of a spectrum $\mathrm{map}(\Sigma^\infty X, E)$. We can construct a filtration on this from the Whitehead tower of E: We can view

$$\ldots \to \tau_{\geq s} E \to \tau_{\geq s-1} E \to \ldots$$

as a filtered spectrum, i.e. a functor $\mathbb{Z}^{\mathrm{op}} \to \mathrm{Sp}$, and compose with the functor $\mathrm{map}(\Sigma^\infty X, -)$ to obtain a filtered spectrum

$$\ldots \to \mathrm{map}(\Sigma^\infty X, \tau_{\geq s} E) \to \mathrm{map}(\Sigma^\infty X, \tau_{\geq s-1} E) \to \ldots.$$

This is still complete (as map preserves limits in the second argument). Also, its underlying spectrum (i.e. the colimit) agrees with $\mathrm{map}(\Sigma^\infty X, E)$: The cofiber of the map

$$\mathrm{colim}_s \,\mathrm{map}(\Sigma^\infty X, \tau_{\geq s} E) \to \mathrm{map}(\Sigma^\infty X, E)$$

is

$$\mathrm{colim}_s \,\mathrm{map}(\Sigma^\infty X, \tau_{\leq s-1} E),$$

where the s'th term is $s-1$-coconnective. So the colimit vanishes, simply by virtue of its homotopy groups becoming 0. The s'th associated graded is (because map is exact in both variables) simply given by

$$\mathrm{map}(\Sigma^\infty X, \Sigma^s \pi_s E).$$

So we have found a filtration on $\mathrm{map}(\Sigma^\infty X, E)$ whose associated graded terms are $\mathrm{map}(\Sigma^\infty X, \Sigma^s \pi_s E)$. But their homotopy groups are just given by singular cohomology! Specifically,

$$E^1_{n,s} = \pi_n \,\mathrm{map}(\Sigma^\infty X, \Sigma^s \pi_s E) \cong H^{s-n}(X; \pi_s E).$$

This means there is a spectral sequence with E^1 page given by $H^*(X; \pi_* E)$, converging to $E^*(X)$ (up to some additional conditions, like E being bounded above, to ensure local finiteness, or a more subtle discussion of convergence).

The usual way to grade this is a little bit different, owing to a mix of differing grading conventions and the fact that there exist different constructions for the same spectral sequence where what we describe here as E^1 page is actually the E^2 page.

Theorem 5.8 (Atiyah-Hirzebruch spectral sequence)**.** *For a space X and a spectrum E there exists a conditionally convergent spectral sequence*

$$E_2^{p,q} = H^p(X; \pi_{-q} E) \Rightarrow E^{p+q}(X)$$

in cohomological Serre grading*: The r-differential goes*

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

and the abutment filtration on $E^k(X)$ is given by

$$\ldots \subseteq F^{s+1} E^k(X) \subseteq F^s E^k(X) \subseteq \ldots$$

with isomorphisms $F^s E^k(X)/F^{s+1}E^k(X) \cong E_\infty^{p,q}$ with $q = -s$ and $p = k + s$.

This grading arises from ours by letting $E_r^{p,q}$ be $E_{n,s}^{r-1}$ with $s = -q$ and $n = -p-q$. For this specific spectral sequence it is preferable to the general one, since it leads to a nice description of the first page. The important information to remember about any of these grading conventions is the bigrading of the d_r differentials ($(r, -r+1)$ here), and where $E^k(X)$ shows up in the E_∞ page: Whereas in our original (Adams) grading it was along the columns, here in Serre grading it is along the codiagonal $p + q = k$, in such a way that the rightmost one, $E_\infty^{k,0}$ here, is a subgroup of $E^k(X)$, and the leftmost one, $E_\infty^{0,k}$, a quotient.

Remark 5.9. The differentials in the Atiyah-Hirzebruch spectral sequence are closely related to the "extra information" required to reconstruct the Postnikov tower of a spectrum from its homotopy groups that we discussed earlier. For example, the d_2 differentials have an interpretation in terms of the maps $\pi_n(X) \to \Sigma^2\pi_{n+1}(X)$ that we saw back then.

Remark 5.10. If we instead apply $\Sigma^\infty X \otimes -$ to the Whitehead tower of E, we obtain a completely analogous spectral sequence for $E_*(X)$. Here, the usual grading is *homological Serre grading*, where

$$E^2_{p,q} \cong H_p(X; \pi_q(E)),$$

the differential d_r has bidegree $(-r, r-1)$, the associated graded of the abutment filtration on $E_k(X)$ is given by the groups $E^\infty_{p,q}$ with $p+q=k$, with $E^\infty_{0,k}$ the bottom subgroup and $E^\infty_{k,0}$ the top quotient.

For example, this gives a relationship between stable homotopy groups of a space, and singular homology with coefficients in $\pi_*\mathbb{S}$. There is a lot of helpful extra structure on the Serre spectral sequence, like relative versions, multiplicative structures, an interpretation of kernel and image of $H^*(B) \to H^*(X)$ and $H^*(X) \to H^*(F)$ in terms of the spectral sequence ("edge homomorphisms") and many more useful details. We will mention these as they come up, but not discuss these systematically.

5.2 The Serre spectral sequence

The Serre spectral sequence, first constructed by Serre in his PhD thesis, gives a relationship between the homology groups in a fiber sequence $F \to X \to B$ of pointed spaces. Even though this is manifestly an object of unstable homotopy theory (the relationship between homology groups for a fiber sequence of *spectra* is much simpler, as we saw earlier it is just a long exact sequence), we can construct it as essentially a twisted version of the Atiyah-Hirzebruch spectral sequence.

For this, we will use a fundamental fact about the ∞-category of spaces, which relates the slice category $\mathcal{S}_{/B}$ (whose objects are given by spaces over B, i.e. maps $E \to B$) and the functor category $\mathrm{Fun}(B, \mathcal{S})$ where we view B as ∞-groupoid:

Theorem 5.11 (Straightening/unstraightening for ∞-groupoids)**.** *Let $B \in \mathcal{S}$. There is an equivalence*

$$\mathrm{Fun}(B, \mathcal{S}) \xrightarrow{\simeq} \mathcal{S}_{/B}$$

where the functor from left to right is given by colim_B *together with an equivalence* $\mathrm{colim}_B \,\mathrm{const}_{\mathrm{pt}} \simeq B$.

The inverse can informally be thought of as taking an object $X \to B$ of $\mathcal{S}_{/B}$ to the functor $B \to \mathcal{S}$ which takes a point $b \in B$ to the fiber X_b over b given by the pullback $\{b\} \times_B X$. The functoriality can be thought as some kind of path lifting operation, where a path from b to b' in B induces a map $X_b \to X_{b'}$, in a homotopy-coherent way. The name of the theorem arises from the fact that one thinks of the inverse of the above functor as picking such path lifting maps, and hence "straightening" the information. The other direction is called "unstraightening".

Baked into the equivalence is the statement that we can write any space X over B as colimit of some functor $B \to \mathcal{S}$.

Now, both $\Sigma^\infty_+ : \mathcal{S} \to \mathrm{Sp}$ and $E \otimes - : \mathrm{Sp} \to \mathrm{Sp}$ preserve colimits. We thus learn:

$$E \otimes \Sigma^\infty_+ X \simeq \mathrm{colim}_{b\in B}\, E \otimes \Sigma^\infty_+ X_b$$

where we write X_b for the functor $B \to \mathcal{S}$. So the homology $E_*(X) = \pi_*(E \otimes \Sigma^\infty_+ X)$ can be described as homology of a colimit of a functor $B \to \mathrm{Sp}$. Before we continue, we need some facts about such functors:

Lemma 5.12. *Let E be a spectrum and $\mathrm{const}_E : B \to \mathrm{Sp}$ the constant functor with value E. Then there are natural equivalences*

$$\mathrm{colim}_B\, \mathrm{const}_E \simeq E \otimes \Sigma^\infty_+ B.$$

and

$$\lim_B \mathrm{const}_E \simeq \mathrm{map}(\Sigma^\infty_+ B, E)$$

Proof. We first consider the first statement. Both sides are natural in B, and hence give functors $\mathcal{S} \to \mathrm{Sp}$. Also, both sides preserve colimits. The right hand side because it is a composite of the colimit-preserving functors Σ^∞_+ and $E \otimes -$, and the left hand side by formal properties of colimits. The ∞-category of spaces has an important universal property: The functor

$$\mathrm{Fun}^{\mathrm{cocont}}(\mathcal{S}, \mathcal{C}) \xrightarrow{\mathrm{ev}_{\mathrm{pt}}} \mathcal{C}$$

is an equivalence for any $\mathcal{C}$ which has colimits. So colimit-preserving functors out of $\mathcal{S}$ are actually completely determined by what they do on pt. For $B = \mathrm{pt}$, both sides of the equivalence we want to prove just evaluate to E, so the universal property finishes the proof of the first statement. For the second one, we may analogously view both sides as colimit-preserving functors $\mathcal{S} \to \mathrm{Sp}^{\mathrm{op}}$, and again compare values at pt. □

The other statement we will need is the following:

Lemma 5.13. *If $B \xrightarrow{F} \mathrm{Sp}$ is a functor which takes values in the heart $\mathrm{Sp}_{\le 0} \cap \mathrm{Sp}_{\ge 0}$, then it factors as $B \xrightarrow{\pi_0 \circ F} \mathrm{Ab} \xrightarrow{H} \mathrm{Sp}$. In particular, if B is simply-connected, F is equivalent to the constant functor with value $\pi_0 F(b_0)$ for some choice of basepoint $b_0 \in B$.*

Proof. The first statement is clear since the heart is equivalent to Ab (via π_0 and H). For the second statement, observe that a functor from B into a 1-category factors through the homotopy category of B, which is the fundamental groupoid. If B is simply-connected, the fundamental groupoid is trivial. □

Remark 5.14. If B is not simply-connected, this proof still shows that functors from B into the heart of spectra are given by *local coefficient systems* on B, which are assignments that assign an abelian group to each point, and an isomorphism to each homotopy class of paths.

Theorem 5.15. *Let $F \to X \to B$ be a fiber sequence of pointed spaces, with B simply-connected. Then there is a filtration on $E \otimes \Sigma^\infty_+ X$ with*

$$\mathrm{gr}_n(E \otimes \Sigma^\infty_+ X) \simeq \Sigma^n(E_n(F) \otimes \Sigma^\infty_+ B).$$

The resulting conditionally convergent spectral sequence can be graded in such a way that it takes the form $H_p(B; E_q(F)) \Rightarrow E_{p+q}(X)$, specifically

$$E^2_{p,q} \cong H_p(B; E_q(F)),$$

the d_r differential has bidegree $(-r, r-1)$, and the associated graded of the abutment filtration on $E_k(X)$ is given by the $E^\infty_{p,q}$ with $p + q = k$.

Proof. Write $X_b : B \to \mathcal{S}$ for the straightening of $X \to B$. By naturality of the Whitehead tower, the

$$\tau_{\geq n}(E \otimes \Sigma^\infty_+ X_b)$$

form a sequence of functors $B \to \mathrm{Sp}$. Taking colimits, we obtain a filtered spectrum whose n-th term is

$$\mathrm{colim}_B\, \tau_{\geq n}(E \otimes \Sigma^\infty_+ X_b).$$

Now the colimit over n is $\mathrm{colim}_B\, E \otimes \Sigma^\infty_+ X_b \simeq E \otimes \Sigma^\infty_+ X$, since colimits commute. This filtration is also complete, since its n-th term is n-connective and so the sequential limit has trivial homotopy groups. Its n'th associated graded is (by exactness of colimits) given by

$$\mathrm{colim}_B\, \Sigma^n \pi_n(E \otimes \Sigma^\infty_+ X_b) \simeq \Sigma^n\, \mathrm{colim}_B\, \pi_n(E \otimes \Sigma^\infty_+ X_b)$$

The functor on B given by $\pi_n(E \otimes \Sigma^\infty_+ X_b)$ agrees with the constant functor with value $\pi_n(E \otimes \Sigma^\infty_+ F) \cong E_n(F)$ since it takes values in the heart and B is simply-connected. Its colimit is given by

$$E_n(F) \otimes \Sigma^\infty_+(B).$$

Taking into account the shift Σ^n from above, this proves the first claim.

For the second claim, it is clear that the homotopy groups of the associated graded take the form $H_*(B; E_*(F))$. The reindexing is analogous to the (homological) Atiyah-Hirzebruch spectral sequence. □

Remark 5.16. Applying instead $\mathrm{map}(\Sigma^\infty_+ -, E)$ to $X \simeq \mathrm{colim}_B\, X_b$, we instead obtain

$$\mathrm{map}(\Sigma^\infty_+ X, E) \simeq \lim_B \mathrm{map}(\Sigma^\infty_+ X_b, E).$$

Again one can form a filtered spectrum whose n-th stage is

$$\lim_{B} \tau_{\geq n} \operatorname{map}(\Sigma_+^\infty X_b, E),$$

and prove that it is complete and its colimit is $\operatorname{map}(\Sigma_+^\infty X, E)$. One obtains a spectral sequence of the form

$$H^p(B; E^q(F)) \Rightarrow E^{p+q}(X),$$

with grading analogous to the cohomological Atiyah-Hirzebruch spectral sequence.

In the special case where E itself is an Eilenberg-MacLane spectrum, one obtains the *Serre spectral sequences*

$$H_p(B; H_q(F)) \Rightarrow H_{p+q}(X)$$
$$H^p(B; H^q(F)) \Rightarrow H^{p+q}(X).$$

Note that here, the filtrations are locally finite, so there are no convergence subtleties.

Among many of the powerful applications of these, one can use these (combining the unstable Postnikov tower and Hurewicz theorem) to prove that for simply-connected spaces, the homology groups are finitely generated if and only if the homotopy groups are finitely generated.

6 Cohomology of Eilenberg-MacLane spaces

In order to understand maps $[HA, \Sigma^n HB]$, i.e. cohomology of Eilenberg-MacLane spectra, we first study the cohomology of Eilenberg-MacLane spaces. Of special interest is the case $A = B = \mathbb{F}_p$.

As some preliminary observations, we can say that $H_*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ are finitely generated, and so the cohomology groups are also finitely generated and all possible forms of UCT and Künneth apply. We also have that $H^k(K(\mathbb{F}_p, n); \mathbb{F}_p) = 0$ for $k < n$ and $H^n(K(\mathbb{F}_p, n); \mathbb{F}_p) \cong \mathbb{F}_p$, for example by Hurewicz and the universal coefficient theorem. A choice of generator corresponds to a choice of natural equivalence $\widetilde{H}^n(X; \mathbb{F}_p) \simeq [X, K(\mathbb{F}_p, n)]$, let us fix such a generator and denote it by ι_n.

To analyze the relationship between $H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ for different n, we need an additional piece of information about the Serre spectral sequence:

Proposition 6.1 (Transgression)**.** *For $F \to X \to B$ a fiber sequence of pointed connected spaces with B simply-connected, consider the zig-zag*

$$H^n(B, \mathrm{pt}; A) \to H^n(X, F; A) \leftarrow H^{n-1}(F; A)$$

where the first map is induced by the map of pairs given by the vertical map in the pullback square for the fiber sequence, and the second map is a connecting homomorphism. Then, in the Serre spectral sequence for this fiber sequence with A-coefficients:

1. *Under the identification $\widetilde{H}^n(B, \mathrm{pt}; A) \cong E_2^{n,0}$, the kernel of the left map is $B_n^{n,0}$, and the image identifies with $E_n^{n,0}$.*
2. *The subgroup of $H^{n-1}(F; A) \cong E_2^{0,n-1}$ which, under the right map is taken into the image of the left map, is $Z_n^{0,n-1} = E_n^{0,n-1}$.*
3. *The map from that subgroup to the image corresponds, under the identifications, to $d_n : E_n^{0,n-1} \to E_n^{n,0}$.*

This is easiest to see in an alternative construction of the Serre spectral sequence, where instead one constructs a filtration from a skeletal filtration on B. The relationship between that construction and the one we used is that the latter is obtained from the former via *décalage* [4], in particular the spectral sequences agree up to a reindexing. $H^*(E, F; A)$ appears then explicitly as a certain quotient in that skeletal filtration, and the connecting homomorphism can directly be related to the differential. We skip the details.

Corollary 6.2. *Under $\widetilde{H}^n(X; A) \cong [X, K(A, n)]$ the map obtained by applying Ω on the right, together with a choice of equivalence $K(A, n) \simeq \Omega K(A, n+1)$, gives a map*

$$\sigma : \widetilde{H}^n(X; A) \to \widetilde{H}^{n-1}(\Omega X; A).$$

This is seen in the Serre spectral sequence of $\Omega X \to \mathrm{pt} \to X$ as follows: We have

$$\widetilde{H}^n(X; A) \cong E_2^{n,0}$$
$$\widetilde{H}^{n-1}(\Omega X; A) \cong E_2^{0,n-1},$$

and: The kernel of σ agrees with $B_n^{n,0}$, the image with $Z_n^{0,n-1}$, and the inverse of the resulting isomorphism $E_n^{n,0} = E_2^{n,0}/B_n^{n,0} \to Z_n^{0,n-1} = E_n^{0,n-1}$ with d_n.

Proof. For the fiber sequence $\Omega X \to \mathrm{pt} \to X$, the zigzag above becomes

$$H^n(X, \mathrm{pt}) \to H^n(\mathrm{pt}, \Omega X) \leftarrow H^{n-1}(\Omega X),$$

with the right map now an isomorphism (for $n \geq 2$). We can identify $H^n(\mathrm{pt}, \Omega X) \cong \widetilde{H}^n(\Sigma\Omega X)$, under that identification the first map is induced by the counit of the adjunction $\Sigma \dashv \Omega$, and the second map is the suspension isomorphism. This composite is just another description of σ. The statements of Proposition 6.1 then imply that the image of σ is given by $Z_n^{0,n-1}$, its kernel by $B_n^{n,0}$, and that the differential is its inverse. $\square$

So the last differential out of the edge term $E_n^{0,n-1}$ (to the other edge) can be identified with a kind of inverse to the loops map above.

Let us call an element of $H^{n-1}(\Omega X; A)$ *transgressive* if the corresponding element of $E^2_{0,n-1}$ in the Serre spectral sequence for $\Omega X \to \mathrm{pt} \to X$ survives to the E^n page. As seen above, these are precisely the elements that can be delooped to $H^n(X; A)$.

In the Serre spectral sequence for $K(\mathbb{F}_p, n) \to \mathrm{pt} \to K(\mathbb{F}_p, n+1)$, we may write

$$E_2^{p,q} = H^p(K(\mathbb{F}_p, n{+}1); H^q(K(\mathbb{F}_p, n); \mathbb{F}_p)) \cong H^p(K(\mathbb{F}_p, n{+}1); \mathbb{F}_p) \otimes_{\mathbb{F}_p} H^q(K(\mathbb{F}_p, n); \mathbb{F}_p).$$

The first nontrivial non-edge term (by total degree) is in bidegree $(p, q) = (n+1, n)$, a $\mathbb{F}_p$ generated by $\iota_n \iota_{n+1}$. This means that for degree reasons, every element of $H^{n+k}(K(\mathbb{F}_p, n); \mathbb{F}_p)$ for $k < n$ is transgressive, and no non-transgressive differentials hit $E_2^{n+k+1,0} = H^{n+k+1}(K(\mathbb{F}_p, n+1); \mathbb{F}_p)$ for $k \leq n$. This shows the first part of the following lemma:

Lemma 6.3. *1. The map $H^{n+k+1}(K(\mathbb{F}_p, n+1); \mathbb{F}_p) \to H^{n+k}(K(\mathbb{F}_p, n); \mathbb{F}_p)$ is an isomorphism for $k < n$ and injective for $k = n$.*

2. The map $H^k(H\mathbb{F}_p; \mathbb{F}_p) \to H^{n+k}(K(\mathbb{F}_p, n); \mathbb{F}_p)$ is an isomorphism for $k < n$ and injective for $k = n$.

Proof. We just saw the first part. For the second part, observe that for any spectra X, Y,

$$\begin{aligned}
\mathrm{Map}_{\mathrm{Sp}}(X, Y) &\simeq \lim_n \mathrm{Map}_{\mathcal{S}_*}(\Omega^{\infty-n} X, \Omega^{\infty-n} Y) \\
&\simeq \lim_n \mathrm{Map}_{\mathrm{Sp}}(\Sigma^{\infty-n} \Omega^{\infty-n} X, Y) \\
&\simeq \mathrm{Map}_{\mathrm{Sp}}(\mathrm{colim}_n\, \Sigma^{\infty-n} \Omega^{\infty-n} X, Y),
\end{aligned}$$

and hence an arbitrary spectrum as a colimit of shifts of suspension spectra,

$$X \simeq \mathrm{colim}_n\, \Sigma^{\infty-n} \Omega^{\infty-n} X.$$

In particular, this means that

$$\mathrm{map}(H\mathbb{F}_p, H\mathbb{F}_p) \simeq \lim_n \mathrm{map}(\Sigma^{\infty-n} K(\mathbb{F}_p, n), H\mathbb{F}_p).$$

Now with a similar disclaimer as when we discussed completeness of the Whitehead tower, in general homotopy groups don't commute with sequential limits, but in this case they do, since for large n the maps on the right become isomorphisms on homotopy groups for a bigger and bigger range of degrees. So we conclude that

$$[H\mathbb{F}_p, \Sigma^k H\mathbb{F}_p] \cong \widetilde{H}^{n+k}(K(\mathbb{F}_p, n); \mathbb{F}_p)$$

for any large enough n, and the precise claim follows from the first one. $\square$

6.1 $p = 2$

Let us write $\mathcal{A}_2 := \bigoplus_n [H\mathbb{F}_2, \Sigma^n H\mathbb{F}_2]$. This is a graded algebra, where the degree n part is given by maps $\mathbb{F}_2 \to \Sigma^n \mathbb{F}_2$. Of course, $\mathcal{A}_2$ agrees with $\pi_* \mathrm{map}(H\mathbb{F}_2, H\mathbb{F}_2)$ with the grading flipped by a minus sign. This $\mathcal{A}_2$ is called the ($p = 2$) *Steenrod algebra*.

By postcomposing, every element $\theta \in [H\mathbb{F}_2, \Sigma^n H\mathbb{F}_2]$ acts on $\widetilde{H}^*(X; \mathbb{F}_2) \cong [\Sigma^\infty X, \Sigma^* HF_2]$ by an additive degree $+n$ operation.

Lemma 6.4. *For every $\theta \in [H\mathbb{F}_2, \Sigma^n H\mathbb{F}_2]$, we get a commutative square*

$$\begin{array}{ccc} \widetilde{H}^m(X;\mathbb{F}_2) & \longrightarrow & \widetilde{H}^{m+n}(X;\mathbb{F}_2) \\ \downarrow{\scriptstyle\sigma} & & \downarrow{\scriptstyle\sigma} \\ \widetilde{H}^{m-1}(\Omega X;\mathbb{F}_2) & \longrightarrow & \widetilde{H}^{m+n-1}(\Omega X;\mathbb{F}_2), \end{array}$$

i.e. the operations from $\mathcal{A}_2$ commute with the "loop" operation.

Proof. Writing $\widetilde{H}^n(X;\mathbb{F}_2)$ unstably as $[X, K(\mathbb{F}_2, n)] = [X, \Omega^{\infty-n} H\mathbb{F}_2]$, we get to interpret the top map as composing with $\Omega^{\infty-n}(\theta)$ and the bottom as $\Omega^{\infty-n+1}(\theta)$. Since σ is itself defined in terms of applying Ω, the statement follows. □

Since the σ is inverse to the transgressive differential $d_n : E_n^{0,n-1} \to E_n^{n,0}$, this in particular tells us that θ takes transgressive elements to transgressive elements (and more precisely, commutes with the transgression). We now make the key observation strengthening Lemma 6.3

Lemma 6.5. *The map*

$$\sigma : H^{n+k+1}(K(\mathbb{F}_2, n+1);\mathbb{F}_2) \to H^{n+k}(K(\mathbb{F}_2, n);\mathbb{F}_2)$$

is an isomorphism for $k < n$. For $k = n$, it is injective and its image contains ι_n^2.

Proof. All of this was said in Lemma 6.3, except for the part about ι_n^2. In the $\mathbb{F}_2$-cohomology Serre spectral sequence for $K(\mathbb{F}_2, n) \to \mathrm{pt} \to K(\mathbb{F}_2, n+1)$, this is saying that ι_n^2 transgresses. For degree reasons, the only nontrivial thing it could hit by an earlier differential is $\iota_n \iota_{n+1}$, so we need to find $d_n(\iota_n^2)$. As ι_n also exists on the E_n page (i.e. it is in the kernel of the earlier differentials), ι_n^2 is still a product on that page, and we have

$$d_n(\iota_n^2) = 2\iota_n \iota_{n+1} = 0$$

since the differentials are derivations. □

It follows from the above that the element $\iota_n^2 \in H^{2n}(K(\mathbb{F}_2, n);\mathbb{F}_2)$ deloops to a unique element of $H^{2n+1}(K(\mathbb{F}_2, n+1);\mathbb{F}_2)$ and then to a unique element of $[H\mathbb{F}_2, \Sigma^n H\mathbb{F}_2]$, which is called Sq^n. For $n = 0$, we let $\mathrm{Sq}^0 = \mathrm{id}$.

Lemma 6.6. *For a space X, the operation $\mathrm{Sq}^n : H^k(X;\mathbb{F}_2) \to H^{k+n}(X;\mathbb{F}_2)$ satisfies $\mathrm{Sq}^n(\alpha) = \alpha^2$ if $k = n$ and $\mathrm{Sq}^n(\alpha) = 0$ if $k < n$.*

Proof. The first statement follows since $\Omega^{\infty-n}(\mathrm{Sq}^n) : K(\mathbb{F}_2, n) \to K(\mathbb{F}_2, 2n)$ is the map representing ι_n^2. For the second statement, one needs to check that the map $K(\mathbb{F}_2, n) \to K(\mathbb{F}_2, 2n)$ representing ι_n^2 loops to 0. This follows for example since in the Serre spectral sequence for $K(\mathbb{F}_2, n-1) \to \mathrm{pt} \to K(\mathbb{F}_2, n)$, there is a differential $d_n(\iota_{n-1}\iota_n) = \iota_n^2$, so $\sigma(\iota_n^2) = 0$. □

In the following theorem, call a sequence $I = (i_0, \ldots, i_k)$ of positive integers *admissible* if $i_k \geq 1$ and $i_r \geq 2i_{r+1}$ for all $0 \leq r \leq k-1$. For such a sequence, the *excess* is $e(I) = \sum_{r=0}^{k}(i_r - 2i_{r+1})$ (where $e_{k+1} = 0$ by convention). For such a sequence of integers, we write $\mathrm{Sq}^I = \mathrm{Sq}^{i_0} \cdots \mathrm{Sq}^{i_k}$.

Theorem 6.7 (Serre)**.** *The cohomology ring of $K(\mathbb{F}_2, n)$ with $\mathbb{F}_2$-coefficients is given by*

$$H^*(K(\mathbb{F}_2, n); \mathbb{F}_2) = \mathbb{F}_2[\mathrm{Sq}^I \iota_n \mid I \text{ admissible with } e(I) < n].$$

Proof sketch. For $n = 1$, the claim is $H^*(K(\mathbb{F}_2, 1); \mathbb{F}_2) = \mathbb{F}_2[\iota_1]$, which follows from $K(\mathbb{F}_2, 1) \simeq \mathbb{R}P^\infty$.

For higher n, it follows by inductively analyzing the Serre spectral sequence of $K(\mathbb{F}_2, n) \to \mathrm{pt} \to K(\mathbb{F}_2, n+1)$. If $I = (i_0, \ldots, i_k)$ is admissible of excess $e(I) = \sum_{r=0}^{k}(i_r - 2i_{r+1}) = i_0 - i_1 - \ldots - i_k$, then the total degree of $\mathrm{Sq}^I \iota_n$ can be written as

$$i_0 + \ldots + i_k + n = 2i_0 - e(I) + n.$$

So, for $i = 2i_0 - e(I) + n$,

$$\mathrm{Sq}^i \mathrm{Sq}^I \iota_n = (\mathrm{Sq}^I \iota_n)^2,$$

$i, i_0, \ldots, i_k$ is admissible iff $e(I) \leq n$, and the excess of $(i, i_0, \ldots, i_k)$ is

$$(i - 2i_0) + e(I) = n,$$

exactly n. Iterating this, this shows that elements of the form $(\mathrm{Sq}^I \iota_n)^{2^k}$ with I admissible of excess $< n$ are exactly given by $\mathrm{Sq}^J \iota_n$ with J admissible of excess $\leq n$, with $J = I$ if $k = 0$ (covering the $e(J) < n$ case), and otherwise J being obtained by applying the above extension procedure k times.

Inductively, $H^*(K(\mathbb{F}_2, n); \mathbb{F}_2)$ is polynomial on the $\mathrm{Sq}^I \iota_n$ with $e(I) < n$. So every monomial admits a unique description as square-free products of the $\mathrm{Sq}^J \iota_n$ with $e(J) \leq n$, or $e(J) < n+1$. These $\mathrm{Sq}^J \iota_n$ transgress to $\mathrm{Sq}^J \iota_{n+1}$. If we believe for a moment that $H^*(K(\mathbb{F}_2, n+1); \mathbb{F}_2)$ is what we claim it is, it will be a polynomial ring on those generators. The differentials are derivations, and so on each square-free product of the $\mathrm{Sq}^J \iota_n$ at least the first possibly nonzero differential is determined by knowing the differentials on the $\mathrm{Sq}^J \iota_n$. From this, one can give a combinatorially slightly involved description of what the spectral sequence would have to look like if the claim was true, and perform an inductive argument (based on the fact that we know the left edge $H^0(K(\mathbb{F}_2, n+1); H^*(K(\mathbb{F}_2, n); \mathbb{F}_2)) = H^*(K(\mathbb{F}_2, n); \mathbb{F}_2)$ of the E_2 page and the entire E_∞ page (it is 0 in positive degrees) to prove that it actually looks like that. □

In particular, in the "linear" range $H^{n+k}(K(\mathbb{F}_2, n); \mathbb{F}_2)$ for $k < n$ is free on $\mathrm{Sq}^I \iota_n$ for admissible sequences of degree $< k$ (which guarantees excess $< n$ automatically). Passing to the limit over n, we conclude:

Theorem 6.8. *$\mathcal{A}_2$ is free on the Sq^I where I ranges over all admissible sequences.*

Next, we prove a fact about the interaction between the operations Sq^n and products. We have on spaces X, Y the cross product

$$H^i(X;A) \times H^j(Y;B) \to H^{i+j}(X \times Y; A \otimes B).$$

This arises in terms of spectra as follows: Given a map $\Sigma_+^{\infty-i} X \to HA$ and $\Sigma_+^{\infty-j} Y \to HB$, we get a map

$$\Sigma_+^{\infty-i-j}(X \times Y) \simeq \Sigma_+^{\infty-i} X \otimes \Sigma_+^{\infty-j} Y \to HA \otimes HB \to H(A \otimes B),$$

where the last map is determined by its effect on π_0 since the target is coconnective and the source connective, and

$$\pi_0(HA \otimes HB) \cong A \otimes_{\mathbb{Z}} B,$$

for example using Hurewicz or by using free resolutions of A and B.

This perspective also allows us to get a more general cross product for cohomology of spectra instead of spaces (by replacing the suspension spectra with arbitrary spectra). We also have results like a Künneth theorem under suitable finiteness assumptions.

Lemma 6.9 (Cartan formula)**.** *For $x \in H^*(X;\mathbb{F}_2)$ and $y \in H^*(Y;\mathbb{F}_2)$, we have*

$$\mathrm{Sq}^n(x \times y) = \sum_{i+j=n} \mathrm{Sq}^i x \times \mathrm{Sq}^j y$$

in $H^*(X \otimes Y;\mathbb{F}_2)$, and similarly for cup products.

Proof. We may represent x by a map $\Sigma_+^{\infty-k} X \to H\mathbb{F}_2$ and y by a map $\Sigma_+^{\infty-l} Y \to H\mathbb{F}_2$. The cross product $x \times y$ is then represented by a map

$$\Sigma_+^{\infty-k-l}(X \times Y) \simeq \Sigma_+^{\infty-k} X \otimes \Sigma_+^{\infty-l} Y \to H\mathbb{F}_2 \otimes H\mathbb{F}_2 \to H\mathbb{F}_2,$$

where the last map is uniquely specified by its effect on π_0 by coconnectivity of the target, and there is given by the multiplication map $\mathbb{F}_2 \otimes_{\mathbb{Z}} \mathbb{F}_2 \to \mathbb{F}_2$.

We ask what that last map does on $[-, \Sigma^n H\mathbb{F}_2]$. It takes degree n elements of $\mathcal{A}_2$ to

$$H^n(H\mathbb{F}_2 \otimes H\mathbb{F}_2; \mathbb{F}_2) \cong \bigoplus_{i+j=n} H^i(H\mathbb{F}_2;\mathbb{F}_2) \otimes H^j(H\mathbb{F}_2;\mathbb{F}_2),$$

i.e. we get a map $\mathcal{A}_2 \to \mathcal{A}_2 \otimes \mathcal{A}_2$ (we are using here a version of Künneth which for example follows directly from the usual one by comparing to cohomology of Eilenberg-MacLane spaces instead of spectra). Writing the image of Sq^n as $\sum \theta' \otimes \theta''$ we get some version of the Cartan formula, that is some identity

$$\mathrm{Sq}^n(x \times y) = \sum \theta' x \times \theta'' y$$

for all x and y, we just have to identify θ' and θ''. We now use that for every $i+j=n$, we have a map

$$K(\mathbb{F}_2,i)\wedge K(\mathbb{F}_2,j)\to K(\mathbb{F}_2,n)$$

inducing $\iota_n\mapsto\iota_i\otimes\iota_j$ on cohomology. On one hand, this must take

$$\iota_n^2\mapsto\iota_i^2\otimes\iota_j^2=\operatorname{Sq}^i\iota_i\otimes\operatorname{Sq}^j\iota_j.$$

while on the other hand, $H^n(H\mathbb{F}_2;\mathbb{F}_2)\to H^{2n}(K(\mathbb{F}_2,n);\mathbb{F}_2)$ is injective.

So we learn that the map $H^i(H\mathbb{F}_2;\mathbb{F}_2)\otimes H^j(H\mathbb{F}_2;\mathbb{F}_2)\to H^{2i}(K(\mathbb{F}_2,i);\mathbb{F}_2)\otimes H^{2j}(K(\mathbb{F}_2,j);\mathbb{F}_2)$ is injective, and thus the degree (i,j) part of $\sum\theta'\otimes\theta''$ must be $\operatorname{Sq}^i\otimes\operatorname{Sq}^j$ as claimed. (For the slightly degenerate case $i=0$ we can do a similar but easier argument using $S^0\wedge K(\mathbb{F}_2,n)\to K(\mathbb{F}_2,n)$, analogously $j=0$.) □

Remark 6.10. Note that for this proof, we have only used that in the range $k\le n$, the map $H^k(H\mathbb{F}_2;\mathbb{F}_2)\to\widetilde{H}^{n+k}(K(\mathbb{F}_2,n);\mathbb{F}_2)$ is an isomorphism, and not the full computation of $\mathcal{A}_2$ in terms of admissible sequences.

Example 6.11. On $H^*(K(\mathbb{F}_2,1);\mathbb{F}_2)=\mathbb{F}_2[\iota_1]$, the action of $\mathcal{A}_2$ is given by

$$\operatorname{Sq}^k\iota_1^n=\binom{n}{k}\iota_1^{n+k}$$

Indeed, we can write $\operatorname{Sq}=\sum\operatorname{Sq}^i$ for the (inhomogeneous) "total Steenrod square", noting that for each x only finitely many of the $\operatorname{Sq}^i(x)$ are nonzero, so this is always a finite sum. The Cartan formula then just means

$$\operatorname{Sq}(xy)=\operatorname{Sq}(x)\operatorname{Sq}(y).$$

Now we have $\operatorname{Sq}(\iota_1)=\iota_1+\iota_1^2$ since Sq^0 is the identity and $\operatorname{Sq}^1(\iota_1)=\iota_1^2$. Now the claim follows from

$$\operatorname{Sq}(\iota_1^n)=\operatorname{Sq}(\iota_1)^n=\iota_1^n(1+\iota_1)^n$$

Lemma 6.12. *The map $H^*(H\mathbb{F}_2;\mathbb{F}_2)\to H^{*+1}(K(\mathbb{F}_2,1);\mathbb{F}_2)$ defined by acting on ι_1 takes*

$$\operatorname{Sq}^I\iota_1\mapsto\begin{cases}\iota_1^{2^k} & \text{if } I=(2^{k-1},2^{k-2},\dots,1)\\ 0 & \text{otherwise}\end{cases}$$

for admissible I

Proof. Clearly, $\operatorname{Sq}^{2^{k-1}}\operatorname{Sq}^{2^{k-2}}\cdots\operatorname{Sq}^1\iota_1=\iota_1^{2^k}$.

On the other hand, if $I=(i_0,\dots,i_n)$ is admissible of excess >1, then it vanishes on ι_1. □

Lemma 6.13. *For any n, the iterated product map*

$$\Sigma^{\infty-n}K(\mathbb{F}_2,1)^{\wedge n}\simeq(\Sigma^{\infty-1}K(\mathbb{F}_2,1))^{\otimes n}\to H\mathbb{F}_2$$

is injective on $H^k(-;\mathbb{F}_2)$ for $k\le n$.

Proof. This is saying equivalently that the map

$$\mathcal{A}_2 \to H^*(K(\mathbb{F}_2, 1)^{\wedge n}; \mathbb{F}_2)$$

obtained by acting on $\iota_1^{\otimes n}$, is injective up to degree $2n$ (in the target).

This is a completely algebraic statement: Using Künneth, we may write the cohomology ring on the right as a polynomial ring $\mathbb{F}_2[x_1, \ldots, x_n]$ on n degree 1 generators. The Steenrod action is again described in terms of the total square $\mathrm{Sq}(x_i) = x_i + x_i^2$. For example, this implies $\mathrm{Sq}(x_i^{2^k}) = (x_i + x_i^2)^{2^k} = x_i^{2^k} + x_i^{2^{k+1}}$, so

$$\mathrm{Sq}^n(x_i^{2^k}) = \begin{cases} x_i^{2^k} \text{ if } n = 0 \\ x_i^{2^{k+1}} \text{ if } n = 2^k \\ 0 \text{ otherwise} \end{cases}$$

Now let $I = (i_{k-1}, \ldots, i_0)$ be an admissible sequence of length k (we reverse the indexing from the previous convention for the following proof to be nicer, admissibility still means that the numbers have to at least double in each step from right to left), and for some $f \in \mathbb{F}_2[x_2, \ldots, x_n]$ consider $\mathrm{Sq}^I(x_1 \cdot f)$ as a polynomial in x_1 with coefficients in $\mathbb{F}_2[x_2, \ldots, x_n]$. As the above tells us, applying a single Sq^i operation can at most double the degree in x_1, so the resulting degree is at most 2^k. In fact, we claim that the coefficient in front of $x_1^{2^k}$ is exactly given by $\mathrm{Sq}^{I'}(f)$, where

$$I' = (i_{k-1} - 2^{k-1}, \ldots, i_0 - 1),$$

with trailing zeroes dropped if necessary. We may prove this inductively:

$$\mathrm{Sq}^{i_{k-2}} \cdots \mathrm{Sq}^{i_0}(x_1 \cdot f) = x_1^{2^{k-1}} \cdot (\mathrm{Sq}^{i_{k-2} - 2^{k-2}} \cdots \mathrm{Sq}^{i_0 - 1} f) + \text{ lower degree terms,}$$

and when we apply $\mathrm{Sq}^{i_{k-1}}$, the highest degree term is contributed to exactly by the $\mathrm{Sq}^{2^{k-1}} \otimes \mathrm{Sq}^{i_{k-1} - 2^{k-1}}$ part of the Cartan formula applied to the leading term.

Through the above construction, an admissible sequence I determines a k (its length) and another admissible sequence I' of length $\leq k$. Conversely, the pair of k and admissible I' determine I, by extending I' to a sequence of length k with zeroes on the right and then adding $(2^{k-1}, \ldots, 1)$. So passing from I to (k, I') is bijective. We also observe that the degree of I' is strictly smaller than the degree of I. Now suppose we have a nontrivial linear combination

$$\sum \mathrm{Sq}^{I_\alpha}(x_1 \cdots x_n) = 0,$$

where the I_α range through admissible sequences of degree n. Letting k be the maximal length among them, we can compare coefficients of $x_1^{2^k}$ and get

$$\sum \mathrm{Sq}^{I'_\alpha}(x_2 \cdots x_n) = 0,$$

where the sum is now only through those I'_α for which I_α has length exactly k. As all of those are different and of degree $\leq n-1$, we inductively know that the $\mathrm{Sq}^{I'_\alpha}(x_2 \cdots x_n)$ are linearly independent, finishing the proof. $\square$

Now, dually, these two statements imply the following:

Corollary 6.14. *1. The map $\Sigma^{\infty-1}K(\mathbb{F}_2,1) \to H\mathbb{F}_2$ induces on $H_*(-;\mathbb{F}_2)$ a map which is nonzero exactly in degrees 2^k-1. Denote by $\zeta_k \in H_{2^k-1}(H\mathbb{F}_2;\mathbb{F}_2)$ the image of the nonzero element of $H_{2^k}(K(\mathbb{F}_2,1);\mathbb{F}_2)$.*

2. The iterated product map

$$(\Sigma^{\infty-1}K(\mathbb{F}_2,1))^{\otimes n} \to H\mathbb{F}_2$$

is surjective on $H_k(-;\mathbb{F}_2)$ for $k \leq n$.

These two statements allow us to give a slick description of $H_*(H\mathbb{F}_2;\mathbb{F}_2)$ due to Milnor. Recall that we have a map

$$\mu : H\mathbb{F}_2 \otimes H\mathbb{F}_2 \to H\mathbb{F}_2$$

characterized by being the multiplication map $\mathbb{F}_2 \otimes \mathbb{F}_2 \to \mathbb{F}_2$ on π_0. Combined with the cross product in homology, this gives a multiplication map

$$H_*(H\mathbb{F}_2;\mathbb{F}_2) \otimes_{\mathbb{F}_2} H_*(H\mathbb{F}_2;\mathbb{F}_2) \to H_*(H\mathbb{F}_2;\mathbb{F}_2).$$

In fact, this gives an associative and (graded-)commutative ring structure, which one can check by observing that various maps

$$H\mathbb{F}_2 \otimes H\mathbb{F}_2 \otimes H\mathbb{F}_2 \to H\mathbb{F}_2, \quad H\mathbb{F}_2 \otimes H\mathbb{F}_2 \to H\mathbb{F}_2$$

are homotopic, since they are all characterized by their effect on π_0. Technically, this also requires further discussion about $\otimes$: We have described this as a functor $\mathrm{Sp}\otimes\mathrm{Sp} \to \mathrm{Sp}$ so far, but not said anything about associativity and symmetry of it. We will say more about this in section 7.1 below.

Theorem 6.15 (Milnor)**.** *The mod 2 homology of $H\mathbb{F}_2$, with product structure given by the multiplication map $H\mathbb{F}_2 \otimes H\mathbb{F}_2 \to H\mathbb{F}_2$, is given by*

$$H_*(H\mathbb{F}_2;\mathbb{F}_2) = \mathbb{F}_2[\zeta_1,\zeta_2,\ldots]$$

Proof. For every n, we have a unique map $(H\mathbb{F}_2)^{\otimes n} \to H\mathbb{F}_2$ characterized by its effect on π_0 given by the multiplication. Since the latter is commutative, these maps are invariant (up to homotopy) under permuting the factors, and the ring structure on $H_*(H\mathbb{F}_2;\mathbb{F}_2)$ is graded-commutative.

The n-th iterated product map is obtained from n copies of the map

$$\Sigma^{\infty-1}K(\mathbb{F}_2,1) \to H\mathbb{F}_2$$

by tensoring them together and composing with the multiplication $(H\mathbb{F}_2)^{\otimes n} \to H\mathbb{F}_2$. It factors through a map

$$\mathbb{F}_2\{\zeta_0,\zeta_1,\ldots\}^{\otimes n} \to H_*(H\mathbb{F}_2;\mathbb{F}_2)$$

surjective up to degree n. It is also not hard to see that ζ_0 goes to the unit, and so it follows that in $H_*(H\mathbb{F}_2;\mathbb{F}_2)$, every element up to degree n can be written as product of at most n factors from among the ζ_i with $i \geq 1$. Making n large, it follows that we have a surjective map

$$\mathbb{F}_2[\zeta_1, \ldots] \to H_*(H\mathbb{F}_2;\mathbb{F}_2).$$

Now, for an admissible sequence $(i_0, \ldots, i_n)$, write $e_k = i_{k-1} - 2i_k$. Then

$$i_{k-1} = e_k + 2i_k = e_k + 2e_{k+1} + 4i_{k+1} = \sum_{l\geq 0} 2^l e_{k+l},$$

and thus

$$|I| = \sum_{k\geq 1} i_{k-1} = \sum_{k\geq 1}\sum_{l\geq 0} 2^l e_{k+l} = \sum_{n\geq 1}(2^n - 1)e_n.$$

This means the admissible sequence I is exactly specified by the sequence of nonnegative integers e_n, and its degree agrees with the degree of

$$\prod \zeta_n^{e_n},$$

and so both sides of the surjective map above have the same (finite) dimension in each degree. Thus, it is also injective. □

6.2 Odd p

The case of odd p is similar, but more involved. Again, $\mathcal{A}_p := \pi_{-*}\operatorname{map}(H\mathbb{F}_p, H\mathbb{F}_p)$ acts on the $\mathbb{F}_p$-cohomology of each space, and preserves the transgressive elements in the spectral sequences $K(\mathbb{F}_p, n) \to \mathrm{pt} \to K(\mathbb{F}_p, n+1)$. The degree argument we used in the $p = 2$ case to prove that the ι_n^2 deloop uniquely to $\mathrm{Sq}^n : H\mathbb{F}_2 \to \Sigma^n H\mathbb{F}_2$ still only works through the same range of degrees (from H^n to H^{2n-1}), but now ι_n^2 does not deloop: In the corresponding spectral sequence, $d_{n+1}\iota_n = \iota_{n+1}$, but $d_{n+1}\iota_n^2 = 2\iota_n\iota_{n+1}$ if n is even, which is nonzero since we are no longer in characteristic 2. Instead, the thing that should have a chance to deloop is ι_n^p for n even, since here $d_{n+1}\iota_n^p = p\iota_n^{p-1}\iota_{n+1} = 0$, since we are working with mod p coefficients. But now a pure degree analysis will not help us determine the longer differentials on ι_n^p.

Definition 6.16. For an integer k and a spectrum X, we have a map

$$k : X \to X,$$

obtained as $n \cdot \mathrm{id}_X$ using the group structure on $[X, X]$. We write

$$[k] : \Omega^{\infty-n}X \to \Omega^{\infty-n}X$$

for the induced maps on underlying spaces.

In the case $X = H\mathbb{F}_p$, the multiplication-by-k map on $H\mathbb{F}_p$ only depends on the residue class of k mod p since $[H\mathbb{F}_p, H\mathbb{F}_p] \cong \mathrm{Hom}(\mathbb{F}_p, \mathbb{F}_p) \cong \mathbb{F}_p$. This means that the maps $[k]$ give an action of $\mathbb{F}_p^\times$ on $H^*(K(\mathbb{F}_p, n);\mathbb{F}_p)$.

Definition 6.17. For $i \in \mathbb{Z}/(p-1)$, let us say $\alpha \in H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ has *weight* i if

$$[k]^*\alpha = k^i \alpha$$

for all $k \in \mathbb{F}_p^\times$.

By a bit of representation theory, viewing $H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ as a $\mathbb{F}_p^\times$-representation, we get a decomposition

$$H^*(K(\mathbb{F}_p, n); \mathbb{F}_p) = \bigoplus_{i \in \mathbb{Z}/(p-1)} H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)^{\mathrm{wt}=i}$$

into the subgroups of weight i elements.

Clearly, the cup product of a weight i element and a weight j element has weight $i+j$. The element ι_n has weight 1, coming from the natural isomorphisms

$$H^n(K(\mathbb{F}_p, n); \mathbb{F}_p) \cong \mathrm{Hom}(H_n(K(\mathbb{F}_p, n)); \mathbb{F}_p) \cong \mathrm{Hom}(\pi_n K(\mathbb{F}_p, n); \mathbb{F}_p \cong \mathrm{Hom}(\pi_0 H\mathbb{F}_p; \mathbb{F}_p).$$

The action of $\mathcal{A}_p$ preserves weight, so all the $\theta\iota_n$ for $\theta \in \mathcal{A}_p$ also have weight 1. Finally, ι_n^p has weight $p = 1$ since weight is an element of $\mathbb{Z}/(p-1)$.

The fiber sequence $K(\mathbb{F}_p, n) \to \mathrm{pt} \to K(\mathbb{F}_p, n+1)$ is also acted upon by $\mathbb{F}_p^\times$ in this way, since it is natural in abelian groups. It follows that we also get a weight decomposition of the entire Serre spectral sequence. The transgressiveness of ι_n^p will arise from showing that the weight 1 part of this is again sparse in degrees roughly $\leq pn$. Since the weight 1 part of the spectral sequence is also contributed to by the weight i and weight j parts of $H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ and $H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ for $i + j = 1 \bmod p - 1$, we also need some control over those. The key result we will need is the following:

Lemma 6.18. *Let $1 \leq i \leq p-1$. Then the weight i part of $\widetilde{H}^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ vanishes in degrees $< ni$, that is*

$$\widetilde{H}^k(K(\mathbb{F}_p, n); \mathbb{F}_p)^{\mathrm{wt}=i} = 0$$

for $k < ni$. In degree $k = ni$, it is generated by ι_n^i (or 0).

Proof. For a pointed space X, we have an equivalence $\Sigma^\infty_+ X \to \mathbb{S} \oplus \Sigma^\infty X$, which is induced by the maps $X_+ \to X$ and $X_+ \to S^0$ of pointed spaces (where X_+ denotes X with a new basepoint added). That this is an equivalence can be seen directly in different ways, but also follows from Hurewicz and an easy consideration in homology. Since Σ^∞_+ takes products to $\otimes$, and Σ^∞ takes smash products to $\otimes$, this splitting also gives an equivalence

$$\Sigma^\infty_+ X^{\times m} \to \bigoplus_{S \subseteq \{1, \ldots, m\}} \Sigma^\infty X^{\wedge S}.$$

The map into the S-th factor is actually induced by a map of pointed spaces, namely the projection $p_S : X^{\times m} \to X^{\times S} \to X^{\wedge S}$, and it can again be directly checked by a homology argument that the resulting map of spectra is an equivalence.

Let us now write $\Delta_S : X \to X^{\times m}$ for the map which is the identity into each of the factors corresponding to elements of S, and the zero map into the others. Then $p_S \circ \Delta_T$ is the diagonal map $\overline{\Delta}_S : X \to X^{\wedge S}$ if $S \subseteq T$, and otherwise the zero map. We now claim the map

$$\Sigma^\infty_+ X \xrightarrow{\sum_{T \subseteq \{1,\dots,m\}} (-1)^{m-|T|} \Delta_T} \Sigma^\infty_+ X^{\times m}$$

factors through $\Sigma^\infty X^{\wedge m}$. Indeed, postcomposing with any p_S, we get

$$\sum_{T \subseteq \{1,\dots,m\}} (-1)^{m-|T|} p_S \circ \Delta_T = \sum_{S \subseteq T \subseteq \{1,\dots,m\}} (-1)^{m-|T|} \overline{\Delta}_S,$$

which is $\overline{\Delta}_S$ if $S = \{1, \dots, m\}$, and 0 if S is a proper subset, because then all terms in the alternating sum cancel. For example, if S does not contain t, then for each $S \subseteq T$ with $t \notin T$, the contributions of T and $T \setminus \{t\}$ cancel in this sum.

So under the splitting of $\Sigma^\infty X^{\times m}$, the map above lands only in the $\Sigma^\infty X^{\wedge m}$ summand. We get a commutative diagram

$$\begin{array}{ccc} \Sigma^\infty_+ X & \xrightarrow{\sum_{T \subseteq \{1,\dots,m\}} (-1)^{m-|T|} \Delta_T} & \Sigma^\infty_+ X^{\times m} \\ & \searrow^{\overline{\Delta}_S} & \uparrow \\ & & \Sigma^\infty X^{\wedge m} \end{array}$$

In spectra, the composite

$$H\mathbb{F}_p \xrightarrow{\Delta_T} H\mathbb{F}_p^{\oplus m} \xrightarrow{\text{codiag}} H\mathbb{F}_p$$

is multiplication with $|T|$. Applying $\Omega^{\infty-n}$, we learn that the composite

$$K(\mathbb{F}_p, n) \xrightarrow{\Delta_T} K(\mathbb{F}_p, n)^{\times m} \xrightarrow{\Omega^{\infty-n}\text{codiag}} K(\mathbb{F}_p, n)$$

is $[|T|]$. Applying Σ^∞_+ again and taking an alternating sum over subsets as above, we learn that

$$\Sigma^\infty_+ K(\mathbb{F}_p, n) \xrightarrow{\sum_{k=0}^m (-1)^{m-k} \binom{m}{k} [k]^*} \Sigma^\infty_+ K(\mathbb{F}_p, n)$$

factors through $\Sigma^\infty K(\mathbb{F}_p, n)^{\wedge m}$. On a weight i element in cohomology, the above map acts as

$$\sum_{k=0}^m (-1)^{m-k} \binom{m}{k} k^i.$$

This is seen to be the number of surjective maps $\{1, \dots, i\} \to \{1, \dots, m\}$: The term in the sum (without the sign) counts pairs of a k-element subset of $\{1, \dots, m\}$ and a function from $\{1, \dots, i\}$ with values in that subset. The alternating sum is then inclusion-exclusion counting functions whose image is not contained in any proper subset.

We now let $m = i$, so that the resulting number is $i!$. As we assumed $1 \le i \le p-1$, this is also coprime to p, so the above map acts in an invertible way on the weight i part of cohomology. Since $K(\mathbb{F}_p, n)^{\wedge i}$ is in-connective, it factors through 0 in degrees $< in$, proving that the weight i part vanishes there.

In degree in, this argument still shows that $H^{in}(K(\mathbb{F}_p, n)^{\wedge i}; \mathbb{F}_p)$ surjects onto the weight i part, via $\overline{\Delta}^*$. As this group (by Künneth) is generated by $\iota_n^{\otimes i}$, which under the diagonal map goes to ι_n^i, this shows that the weight i part in degree in is generated by ι_n^i. □

Note that $\iota_n^i = 0$ if $i \ge 2$ and n is odd, by graded-commutativity. On the other hand, if n is even, the fact that $H^*(\mathbb{C}P^\infty; \mathbb{F}_p)$ contains an element of degree n all of whose powers are nonzero shows that none of the ι_n^i is zero.

Lemma 6.19. *The map*

$$\sigma : H^{n+k+1}(K(\mathbb{F}_p, n+1); \mathbb{F}_p)^{\mathrm{wt}=1} \to H^{n+k}(K(\mathbb{F}_p, n); \mathbb{F}_p)^{\mathrm{wt}=1}$$

is an isomorphism for $k < (p-1)n$. For even n and $k = (p-1)n$, it is injective and its image contains ι_n^p.

Proof. We analyze the transgression in the weight 1 part of the $\mathbb{F}_p$-cohomology Serre spectral sequence for $K(\mathbb{F}_p, n) \to \mathrm{pt} \to K(\mathbb{F}_p, n+1)$.

Writing the E_2 page as

$$H^*(K(\mathbb{F}_p, n); \mathbb{F}_p) \otimes H^*(K(\mathbb{F}_p, n+1); \mathbb{F}_p),$$

we have a weight splitting on each of these factors. The weight 1 part is contributed to by

$$H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)^{\mathrm{wt}=1} \otimes H^0(K(\mathbb{F}_p, n+1)) H^0(K(\mathbb{F}_p, n)) \otimes H^*(K(\mathbb{F}_p, n+1); \mathbb{F}_p)^{\mathrm{wt}=1}$$

and

$$\widetilde{H}^*(K(\mathbb{F}_p, n); \mathbb{F}_p)^{\mathrm{wt}=i} \otimes \widetilde{H}^*(K(\mathbb{F}_p, n+1))^{\mathrm{wt}=j}$$

for $1 \le i \le p-1$, $1 \le j \le p-1$ and $i + j = p$. From the connectivity bounds on the weight i parts, these latter terms are trivial in total degrees below

$$in + j(n+1) = pn + j.$$

Furthermore, in total degree $pn + j$, they are generated by $\iota_n^i \iota_{n+1}^j$ (or zero).

It follows that there are no nontrivial non-edge terms in the spectral sequence of total degree $< pn+1$. It follows that every element of $E_*^{0,n+k}$ transgresses for $k < (p-1)n$, and no non-transgressive differential hits $E_*^{n+k+1,0}$ for $k \le (p-1)n$. Finally, as the only possibly non-zero non-edge term of total degree $pn + 1$ is generated by $\iota_n^{p-1} \iota_{n+1}$, the only possible nontransgressive differential on ι_n^p for even n is $d_{n+1} \iota_n^p$, which is zero again because d_{n+1} is a derivation. □

Again, it follows that for even $n = 2i$ the element $\iota_n^p \in H^{pn}(K(\mathbb{F}_p, n); \mathbb{F}_p)^{\text{wt}=1}$ deloops to a unique element in all the $H^{pn+k}(K(\mathbb{F}_p, n+k); \mathbb{F}_p)^{\text{wt}=1}$, which for $(p-1)n < k$ agrees with all of $H^{pn+k}(K(\mathbb{F}_p, n+k); \mathbb{F}_p)$ and hence $[H\mathbb{F}_p, \Sigma^{2i(p-1)} H\mathbb{F}_p]$. We denote the resulting element of $\mathcal{A}_p$ by P^i.

Lemma 6.20. *For a space X, the operation $P^i : H^k(X; \mathbb{F}_2) \to H^{k+2i(p-1)}(X; \mathbb{F}_2)$ satisfies $P^i(\alpha) = \alpha^p$ if $k = 2i$ and $P^i(\alpha) = 0$ if $k < 2i$.*

Proof. This follows as in the $p = 2$ case from the construction and the fact that in the $K(\mathbb{F}_p, 2i-1) \to \text{pt} \to K(\mathbb{F}_p, 2i)$ Serre spectral sequence there is a differential $d_{2i}\iota_{2i-1}\iota_{2i}^{p-1} = \iota_{2i}^p$. □

Contrary to the $p = 2$ case, not all differentials in the $K(\mathbb{F}_p, n) \to \text{pt} \to K(\mathbb{F}_p, n+1)$ Serre spectral sequence are determined completely by transgressive differentials. So the computation of $H^*(K(\mathbb{F}_p, n); \mathbb{F}_p)$ cannot proceed completely analogously. Instead, we first prove the Cartan formula:

Lemma 6.21 (Cartan formula)**.** *For $x \in H^*(X; \mathbb{F}_p)$ and $y \in H^*(Y; \mathbb{F}_p)$, we have*

$$P^n(x \times y) = \sum_{i+j=n} P^i x \times P^j y$$

in $H^*(X \otimes Y; \mathbb{F}_p)$, and similarly for cup products.

Proof. As in the $p = 2$ proof, one sees that there must be some Cartan formula

$$P^n(x \times y) = \sum \theta' x \times \theta'' y.$$

Now one has

$$K(\mathbb{F}_p, 2i) \wedge K(\mathbb{F}_p, 2j) \to K(\mathbb{F}_p, 2n)$$

inducing $\iota_{2n} \mapsto \iota_{2i} \otimes \iota_{2i}$ on cohomology. Again, this takes

$$\iota_{2n}^p \mapsto \iota_{2i}^p \otimes \iota_{2j}^p = P^i \iota_{2i} \otimes P^j \iota_{2j}.$$

On the other hand,

$$H^{2i(p-1)}(H\mathbb{F}_p; \mathbb{F}_p) \otimes H^{2j(p-1)}(H\mathbb{F}_p; \mathbb{F}_p) \to H^{2ip}(K(\mathbb{F}_p, 2i); \mathbb{F}_p) \otimes H^{2jp}(K(\mathbb{F}_p, 2j); \mathbb{F}_p)$$

is injective, and so this shows that the degree (i, j) part of $\sum \theta' \otimes \theta''$ is $P^i \otimes P^j$. □

In addition to the P^i, we also have a map $H\mathbb{F}_p \to \Sigma H\mathbb{F}_p$ obtained as the cofiber of $H\mathbb{Z}/p^2 \to H\mathbb{F}_p$. This is called the *Bockstein homomorphism* β. Starting with the top left square in the following diagram and taking horizontal and vertical cofibers, we obtain

$$\begin{array}{ccccc}
H\mathbb{Z}/p^3 & \longrightarrow & H\mathbb{Z}/p^2 & \longrightarrow & \Sigma\mathbb{F}_p \\
\downarrow & & \downarrow & & \downarrow \\
H\mathbb{F}_p & \longrightarrow & H\mathbb{F}_p & \longrightarrow & 0 \\
\downarrow & & \downarrow{\scriptstyle\beta} & & \downarrow \\
\Sigma H\mathbb{Z}/p^2 & \longrightarrow & \Sigma H\mathbb{F}_p & \xrightarrow{\beta} & \Sigma^2\mathbb{F}_p,
\end{array}$$

where the bottom right square witnesses that $\beta\beta = 0$.

Lemma 6.22 (Derivation property for the Bockstein)**.** *We have* $\beta(xy) = \beta(x)y + (-1)^{|x|}x\beta(y)$.

Proof. As in the other Cartan formula arguments, the existence of *some* formula for $\beta(xy)$ independent of x and y follows from expressing the image of β in

$$H^1(H\mathbb{F}_p \otimes H\mathbb{F}_p; \mathbb{F}_p)$$

via Künneth. Consider the commutative square

$$\begin{array}{ccc} H\mathbb{Z}/p^2 \otimes H\mathbb{Z}/p^2 & \longrightarrow & H\mathbb{Z}/p^2 \\ \downarrow & & \downarrow \\ H\mathbb{F}_p \otimes H\mathbb{F}_p & \longrightarrow & H\mathbb{F}_p \end{array}$$

Since the Bockstein is the cofiber of the right hand vertical map, the diagram tells us that the image of the Bockstein is in the kernel of

$$H^1(H\mathbb{F}_p \otimes H\mathbb{F}_p; \mathbb{F}_p) \to H^1(H\mathbb{Z}/p^2 \otimes H\mathbb{Z}/p^2; \mathbb{F}_p).$$

This is the composite of

$$H^1(H\mathbb{F}_p \otimes H\mathbb{F}_p; \mathbb{F}_p) \to H^1(H\mathbb{F}_p \otimes H\mathbb{Z}/p^2; \mathbb{F}_p) \to H^1(H\mathbb{Z}/p^2 \otimes H\mathbb{Z}/p^2; \mathbb{F}_p),$$

and from the long exact sequences we see that the kernel of the first sequence is generated by $\iota \otimes \beta$, the second by $\beta \otimes \iota$. It follows that the image of β is some linear combination of $\beta \otimes \iota$ and $\iota \otimes \beta$. By precomposing with $\mathbb{S} \otimes H\mathbb{F}_p \to H\mathbb{F}_p \otimes H\mathbb{F}_p$ and $H\mathbb{F}_p \otimes \mathbb{S} \to H\mathbb{F}_p \otimes H\mathbb{F}_p$, we identify both coefficients as 1.

Now the sign appears from translating for $x : X \to \Sigma^n H\mathbb{F}_p$ and $y : Y \to \Sigma^m H\mathbb{F}_p$ between the resulting rule for $x \times y$ viewed as map

$$\Sigma^{-n}X \otimes \Sigma^{-m}Y \to H\mathbb{F}_p$$

and the rule for $x \times y$ viewed as map

$$X \otimes Y \to \Sigma^{n+m}H\mathbb{F}_p,$$

there is a sign appearing here consistent with the Koszul sign rule. □

Example 6.23. We have $H^*(K(\mathbb{F}_p, 1); \mathbb{F}_p) = \Lambda_{\mathbb{F}_p}(\iota_1)[\beta\iota_1]$, with action by P^k and βP^k determined by the Cartan formula. Indeed, we can again write $P = \sum P^i$ for the total operation, and the Cartan formula then says $P(xy) = P(x)P(y)$. Applied here, we get

$$P(\iota_1) = \iota_1$$
$$P(\beta\iota_1) = \beta\iota_1 + (\beta\iota_1)^p,$$

and so

$$P((\beta\iota_1)^n) = (\beta\iota_1)^n(1+(\beta\iota_1)^{p-1})^n, P(\iota_1(\beta\iota_1)^n) = \iota_1(\beta\iota_1)^n(1+(\beta\iota_1)^{p-1})^n.$$

It follows that

$$P^k((\beta\iota_1)^n) = \binom{n}{k}(\beta\iota_1)^{n+(p-1)k}$$
$$P^k(\iota_1(\beta\iota_1)^n) = \binom{n}{k}\iota_1(\beta\iota_1)^{n+(p-1)k}.$$

We also learn that

$$\beta P^k((\beta\iota_1)^n) = 0$$
$$\beta P^k(\iota_1(\beta\iota_1)^n) = \binom{n}{k}(\beta\iota_1)^{n+(p-1)k+1},$$

using $\beta\beta = 0$ and the derivation property.

One interesting consequence is that βP^k is generally nontrivial. This means that $\beta P^k \in H^{2k(p-1)+1}(H\mathbb{F}_p;\mathbb{F}_p)$ is nonzero, and hence also $\beta P^k\iota_{2k+1} \in H^{2kp+2}(K(\mathbb{F}_p, 2k+1);\mathbb{F}_p)$ is nonzero. However, its image under σ, $\beta P^k\iota_{2k} = \beta(\iota_{2k}^p) = 0$ vanishes. This means that in the Serre spectral sequence for $K(\mathbb{F}_p, 2k) \to \mathrm{pt} \to K(\mathbb{F}_p, 2k+1)$, the element $\beta P^k\iota_{2k+1}$ in $E_*^{2kp+2,0}$ must be hit by a nontransgressive differential. As we saw in the proof of Lemma 6.19, the only nonzero non-edge term of weight 1 in total degree $2kp+1$ is generated by $\iota_{2k}^{p-1}\iota_{2k+1}$. Hence we must have that $d_{2k(p-1)+1}(\iota_{2k}^{p-1}\iota_{2k+1}) = u\beta P^k\iota_{2k+1}$ for $u \in \mathbb{F}_p^\times$.

Theorem 6.24 (Kudo transgression theorem)**.** *Assume for $F \to X \to B$ a fiber sequence of pointed connected spaces with B simply-connected, we have an element $x \in H^{2k}(F;\mathbb{F}_p)$ which in the Serre spectral sequence transgresses to $y = d_{2k+1}(x)$. Then x^p transgresses to $P^k y$, and $x^{p-1}y$ survives to the $E_{2k(p-1)+1}$-page and satisfies*

$$d_{2k(p-1)+1}(x^{p-1}y) = u\beta P^k y.$$

with some $u \in \mathbb{F}_p^\times$.

Proof. Pick a map $B \to K(\mathbb{F}_p, 2k+1)$ representing y, and form the diagram of fiber sequences

$$\begin{array}{ccc} F & \longrightarrow & K(\mathbb{F}_p, 2k) \\ \downarrow & & \downarrow \\ X & \longrightarrow & \mathrm{pt} \\ \downarrow & & \downarrow \\ B & \longrightarrow & K(\mathbb{F}_p, 2k+1), \end{array}$$

as follows: using that y is hit by a differential in the spectral sequence, it is in the kernel of $H^{2k+1}(B;\mathbb{F}_p) \to H^{2k+1}(X;\mathbb{F}_p)$ by an edge homomorphism consideration. This means we can find a nullhomotopy filling the bottom square, and pass to homotopy fibers.

The induced map between Serre spectral sequences takes ι_{2k} to an element of $H^{2k}(F;\mathbb{F}_p)$ which transgresses to y. It therefore differs from x by an element which is a permanent cycle. Realizing this by an element of $H^{2k}(X;\mathbb{F}_p)$, we may modify the homotopy of the bottom square by it (using that $[X, \Omega K(\mathbb{F}_p, 2k+1)]$ acts on the nullhomotopies) to achieve that the map $F \to K(\mathbb{F}_p, 2k)$ actually classifies x. Then the differentials in the latter spectral sequence prove the claim. □

Remark 6.25. The unit u can actually be seen to be 1, but that requires a more refined discussion of the Serre spectral sequence, for example in terms of a chain-complex model.

Theorem 6.26 (Cartan)**.** *For odd p, $H^*(K(\mathbb{F}_p, n);\mathbb{F}_p)$ is free as a graded-commutative $\mathbb{F}_p$-algebra on elements of the form*

$$\beta^{\varepsilon_0} P^{i_0} \cdots \beta^{\varepsilon_k} P^{i_k} \beta^{\varepsilon_{k+1}} \iota_n$$

with $\varepsilon_r \in \{0,1\}$, $2i_r \geq 2pi_{r+1} + \varepsilon_{r+1}$ (admissibility) for all r, and

$$\sum_{r\geq 0} (2i_r - 2pi_{r+1} - \varepsilon_{r+1}) < n$$

(excess).

Proof. Again, this is an inductive argument, with the structure of the spectral sequence determined by the Kudo transgression theorem. Writing $I = (\varepsilon_0, i_0, \ldots, i_k, \varepsilon_{k+1})$ for a sequence admissible in the above sense, $e(I)$ for its excess, and $\mathcal{P}^I = \beta^{\varepsilon_0} P^{i_0} \cdots P^{i_k} \beta^{\varepsilon_{k+1}}$, we have that the degree of $\mathcal{P}^I \iota_n$ is

$$|I| + n = \sum_{r\geq 0} \varepsilon_r + \sum_{r \geq 0} 2i_r(p-1) + n.$$

This element transgresses to $\mathcal{P}^I \iota_{n+1}$. If $|I| + n = 2i$ is even, then we have

$$P^i \mathcal{P}^I \iota_{n+1} = (\mathcal{P}^I \iota_{n+1})^p,$$

and more generally

$$P^{p^{k-1}i} \cdots P^i \mathcal{P}^I \iota_{n+1} = (\mathcal{P}^I \iota_{n+1})^{p^k}.$$

These sequences obtained by prepending I with $(0, p^{k-1}i, 0, \ldots, i)$ (the 0's referring to no further Bocksteins being involved), the resulting sequence is admissible and has excess

$$\begin{aligned} &(2p^{k-1}i - 2p(p^{k-2}i)) + \ldots + (2i - 2pi_0 - \varepsilon_0) + \ldots \\ =&2i - \sum_{r\geq 0} \varepsilon_r - \sum_{r\geq 0} 2i_r(p-1) = 2i - |I| = n \end{aligned}$$

Conversely, all admissible sequences of excess exactly n are of that form, possibly with an additional β in front.

Now, by the Kudo transgression theorem, $(\mathcal{P}^I \iota_{n+1})^{p^k}$ transgresses to $P^{p^{k-1}i} \cdots P^i \mathcal{P}^I \iota_{n+2}$, and the elements

$$(\mathcal{P}^I \iota_{n+1})^{p^{k-1}(p-1)}(P^{p^{k-2}i} \cdots P^i \mathcal{P}^I \iota_{n+2})$$

supports a differential which hits $u \cdot \beta P^{p^{k-1}i} \cdots P^i \mathcal{P}^I \iota_{n+2}$. This way, the (new) admissible sequences of excess exactly n arise. From this, one can fully describe what the spectral sequence has to look like, and prove by an inductive argument that this is the case. □

Theorem 6.27. *$\mathcal{A}_p$ is free on the $\mathcal{P}^I$ where I ranges over all admissible sequences in the above sense.*

Lemma 6.28. *1. The map $H^*(H\mathbb{F}_p; \mathbb{F}_p) \to H^{*+1}(K(\mathbb{F}_p, 1); \mathbb{F}_p)$ defined by acting on ι_1 takes*

$$\mathcal{P}^I \mapsto \begin{cases} (\beta\iota_1)^{p^k} & \text{if } \mathcal{P}^I = P^{p^{k-1}} \cdots P^1\beta \\ 0 & \text{otherwise} \end{cases}$$

for admissible I.

2. The map $H^*(H\mathbb{F}_p; \mathbb{F}_p) \to H^{*+2}(K(\mathbb{Z}, 2); \mathbb{F}_p)$ defined by acting on ι_2 takes

$$\mathcal{P}^I \mapsto \begin{cases} \iota_2^{p^k} & \text{if } \mathcal{P}^I = P^{p^{k-1}} \cdots P^1 \\ 0 & \text{otherwise} \end{cases}$$

for admissible I.

Proof. In the first case, if the sequence has excess > 1, the value is seen to be 0 from the fact that $P^i x = 0$ if $2i > |x|$. For it to have excess 1, it needs to be of the claimed form, and then one directly sees that the value is as claimed.

In the second case, because $H^*(K(\mathbb{Z}, 2); \mathbb{F}_p) = \mathbb{F}_p[\iota_2]$ is even, the only sequences that can act nontrivially are the ones without βs. Then, the excess needs to be exactly 2, leaving only sequences as claimed. □

We also have the following analogue of Lemma 6.13, in which the bound is not optimal, but sufficient for our purposes (the optimal bound may be found a posteriori in terms of Theorem 6.30 below, as the largest d such that every monomial in the τ_i and ξ_i of degree $\le d$ involves at most n factors of the τ_i and m factors of the ξ_i. This is slightly more complicated than the situation at $p = 2$ because of the fact that the τ_i square to 0.)

Lemma 6.29. *For any n, the iterated map*

$$\Sigma^{\infty - 3n} K(\mathbb{F}_p, 1)^{\wedge n} \wedge K(\mathbb{Z}, 2)^{\wedge m} \to H\mathbb{F}_p$$

is injective on H^k for $k \le n + m$.

Proof. This is saying equivalently that the map

$$\mathcal{A}_p \to H^*(K(\mathbb{F}_p,1)^{\wedge n} \wedge K(\mathbb{Z},2)^{\wedge m};\mathbb{F}_p)$$

obtained by acting on $\iota_1^{\otimes n} \otimes \iota_2^{\otimes m}$, is injective up to degree $n+2m+\min(n,m)$ (in the target). This is again a completely algebraic statement, where the cohomology ring here is given by the graded-commutative ring

$$\Lambda_{\mathbb{F}_p}(x_1,\dots,x_n)[\beta x_1,\dots,\beta x_n,y_1,\dots,y_m].$$

The proof now proceeds similarly to 6.13, based on the following two observations:

1. For sequences that end in β, we have that

$$\beta^{\varepsilon_{k-1}}P^{i_{k-1}}\cdots\beta^{\varepsilon_0}P^{i_0}\beta(x_1\cdot f)$$

   with f in the subring generated by $x_2,\dots,x_n$, $\beta x_2,\dots,\beta x_n$, and $y_1,\dots,y_m$, has largest power of βx_1 given by $(\beta x_1)^{p^k}$, with coefficient given by

$$\beta^{\varepsilon_k}P^{i_{k-1}-p^{k-1}}\cdots\beta^{\varepsilon_1}P^{i_1-1}f.$$

2. For sequences that do not end in β, we have that

$$\beta^{\varepsilon_{k-1}}P^{i_{k-1}}\cdots\beta^{\varepsilon_0}P^{i_0}(y_1\cdot f)$$

   with f in the subring generated by $x_1,\dots,x_n$, $\beta x_1,\dots,\beta x_n$, and $y_2,\dots,y_m$, has largest power of y_1 given by $y_1^{p^k}$, with coefficient given by

$$\beta^{\varepsilon_{k-1}}P^{i_{k-1}-p^{k-1}}\cdots\beta^{\varepsilon_0}P^{i_0-1}f.$$

Again, these allow us to turn a nontrivial linear relation $\sum\lambda_\alpha\mathcal{P}^{I_\alpha}(x_1\cdots x_n y_1\cdots y_m) = 0$ into a nontrivial linear relation of the form $\sum\lambda_\alpha\mathcal{P}^{I'_\alpha}(x_2\cdots x_n y_1\cdots y_n) = 0$ or $\sum\lambda_\alpha\mathcal{P}^{I'_\alpha}(x_1\cdots x_n y_2\cdots y_m) = 0$: We choose k based on the longest I_α, and then compare coefficients of $(\beta x_1)^{p^k}$ if any of them end in β, or $y_1^{p^k}$ otherwise. This leads to a linear relation involving sequences of strictly smaller degree, and reduces the total number of variables by 1. Inductively, we may assume that those elements are linearly independent, finishing the proof by contradiction. □

As in the $p=2$ case, by dualizing, we get maps $H_{*+1}(K(\mathbb{F}_p,1);\mathbb{F}_p) \to H_*(H\mathbb{F}_p;\mathbb{F}_p)$ and $H_{*+2}(K(\mathbb{Z},2);\mathbb{F}_p) \to H_*(H\mathbb{F}_p;\mathbb{F}_p)$ which are only nontrivial in degrees $* = 2p^k-1$ for $k\geq 0$ in the first case and $* = 2(p^k-1)$ for $k\geq 1$ in the second. In the first case, we denote the image of canonical generators τ_k, in the second ξ_k. The injectivity statement dualizes to a surjectivity statement on homology which proves that $H_*(H\mathbb{F}_p;\mathbb{F}_p)$ is generated as graded-commutative ring by those elements. As in the $p=2$ case, a counting argument then shows the following:

Theorem 6.30 (Milnor)**.** *For p odd, the mod p homology of $H\mathbb{F}_p$, with product structure given by the multiplication map $H\mathbb{F}_p \otimes H\mathbb{F}_p \to H\mathbb{F}_p$, is given by*

$$H_*(H\mathbb{F}_p; \mathbb{F}_p) = \Lambda_{\mathbb{F}_p}(\tau_0, \tau_1, \ldots)[\xi_1, \xi_2, \ldots]$$

Proof. By the dual of the injectivity statement, and graded commutativity, the map

$$\Sigma^{\infty-3n} K(\mathbb{F}_p, 1)^{\wedge n} \wedge K(\mathbb{Z}, 2)^{\wedge n} \to H\mathbb{F}_p$$

is surjective in degrees up to n. Since the image is spanned by monomials in the τ_i and ξ_i, we get that the map of graded-commutative rings

$$\Lambda_{\mathbb{F}_p}(\tau_0, \tau_1, \ldots)[\xi_1, \xi_2, \ldots] \to H_*(H\mathbb{F}_p; \mathbb{F}_p)$$

is surjective up to degree n for any n, so it is surjective. Now, a counting argument shows again that the dimensions on both sides agree, comparing the number of admissible sequences of a given degree with the number of monomials in the τ_i and ξ_i. □

7 The Steenrod algebra and its dual

7.1 Ring spectra and their modules

So far, we have used relatively little about the functor $\otimes : \mathrm{Sp} \times \mathrm{Sp} \to \mathrm{Sp}$ (maybe in some of the arguments using cup products, we used that it is associative and symmetric up to homotopy). Of course, it should be part of a symmetric monoidal structure on Sp. In 1-category theory, associativity and symmetry of a monoidal structure are witnessed by natural isomorphisms of functors, which are part of the structure and need to satisfy certain relations. For example, one has the braiding $X \otimes Y \to Y \otimes X$ (natural in X and Y), and composing this with itself should equal the identity.

In ∞-categories, these relations of course have to be witnessed by higher homotopies, and these again need to satisfy relations up to higher homotopies, and so on. This leads to an infinite amount of structure, the combinatorics of which is hard to make explicit. Instead, symmetric monoidal structures on ∞-categories are encoded using the language of (colored, ∞-)*operads*, which roughly encode all of the coherence homotopies of $\otimes$ in terms of the nerve of the category of finite sets. For details, this is fully developed in [13, Chapter 2]. We will not pursue these details here, instead trusting that most things work out in a similar way as in 1-categories (while of course pointing out the more subtle points as they arise).

The more coherent version of Lemma 3.16 is:

Proposition 7.1. Sp *carries a unique symmetric monoidal structure where the unit is $\mathbb{S}$ and $\otimes$ commutes with colimits in both variables. The functor*

$$\Sigma^\infty : \mathcal{S}_* \to \mathrm{Sp}$$

is symmetric monoidal with respect to the symmetric monoidal structure on $\mathcal{S}_*$ given by $\wedge$.

The functor $H : \mathrm{Ab} \to \mathrm{Sp}$ is not symmetric monoidal: $H\mathbb{F}_p \otimes H\mathbb{F}_p$ is different from $H(\mathbb{F}_p \otimes \mathbb{F}_p)$, since the former has homotopy $H_*(H\mathbb{F}_p; \mathbb{F}_p)$ which we saw is complicated, while the latter is an Eilenberg-Maclane spectrum. It is however *lax symmetric monoidal*. Informally, a lax symmetric-monoidal structure on a functor comes with natural transformations $F(X) \otimes F(Y) \to F(X \otimes Y)$, but of course there also need to be compatibility with the higher coherences in the monoidal structure. (A symmetric monoidal functor however is really just a lax symmetric monoidal one where the map $F(X) \otimes F(Y) \to F(X \otimes Y)$ happens to be an equivalence for each X, Y, no additional structure required.)

In the case of the functor H, informally the lax symmetric monoidal structure arises from the fact that maps $HA \otimes HB \to H(A \otimes B)$ are uniquely characterized by their effect on π_0, and both π_0 canonically identify with $A \otimes B$.

In a symmetric-monoidal category, we can talk about (associative and commutative) algebra objects. Informally, an associative algebra object consists of an object A with maps $\mathbb{1} \to A$ and $A \otimes A \to A$, homotopies witnessing unitality and associativity, and higher homotopies witnessing more subtle relations between these homotopies. A commutative algebra additionally comes with homotopies witnessing symmetry of the multiplication and higher coherences related to that. Both of these concepts also are made precise using the language of operads.

In Sp, we call associative and commutative algebras simply associative and commutative ring spectra. $\mathbb{S}$ as unit of the symmetric monoidal structure carries a unique commutative ring structure. Lax symmetric monoidal functors take algebra objects to algebra objects. In particular, if R is an (associative or commutative) ordinary ring, it is an (associative or commutative) algebra in Ab, and so HR inherits a canonical (associative or commutative) ring structure. (In this case, more is true, algebra structures on an Eilenberg-MacLane spectrum correspond uniquely to algebra structures on its π_0, viewed as object in Ab. This is because $\mathrm{Ab} \to \mathrm{Sp}$ is not just a full subcategory, it is in fact a *full suboperad*.)

If R is a ring spectrum, we get cross products on R-(co)homology: Given two maps

$$X \to \Sigma^n R, Y \to \Sigma^m R,$$

we may tensor them together and compose with the ring structure map to get

$$X \otimes Y \to \Sigma^{n+m} R \otimes R \to \Sigma^{n+m} R.$$

Similarly, given $\mathbb{S}^n \to R \otimes X$ and $\mathbb{S}^m \to R \otimes Y$, we may tensor them together, rearrange factors and compose to get

$$\mathbb{S}^{n+m} \to R \otimes R \otimes X \otimes Y \to R \otimes X \otimes Y.$$

This gives bilinear maps $R^n(X) \otimes R^m(Y) \to R^{n+m}(X \otimes Y)$ and $R_n(X) \otimes R_m(Y) \to R_{n+m}(X \otimes Y)$.

Given an algebra object R, we may also talk about R-modules. Informally, a (left) R-module consists of an object M together with a map $R \otimes M \to M$, homotopies witnessing associativity and unitality of the action, and again higher coherence information. Again, this is formalized using operads. Every object is uniquely a module over the monoidal unit, and lax symmetric monoidal functors preserve module structures. For example, for an ordinary ring R and an ordinary R-module M, HM inherits a HR-module structure.

For an algebra object R in a symmetric monoidal category $\mathcal{C}$, left modules form a category $\mathrm{LMod}_R(\mathcal{C})$. The forgetful functor $\mathrm{LMod}_R(\mathcal{C}) \to \mathcal{C}$ has a left adjoint, which takes $X \in \mathcal{C}$ to $R \otimes X$ with canonical R-module structure. This means

$$\mathrm{Map}_{\mathrm{LMod}_R(\mathcal{C})}(R \otimes X, M) \simeq \mathrm{Map}_{\mathcal{C}}(X, M).$$

If $\mathcal{C}$ is stable (and $\otimes$ exact in each variable), then $\mathrm{LMod}_R(\mathcal{C})$ is also stable, and the above equivalence of mapping spaces refines to one of mapping spectra. Analogously, we have a category of right modules $\mathrm{RMod}_R(\mathcal{C})$ and a corresponding adjunction.

Given a right R-module M and a left R-module N, assuming $\mathcal{C}$ has enough colimits, there exists a relative tensor product $M \otimes_R N$, which has no module structure left, it is just an object of $\mathcal{C}$. Informally is characterized by a universal property analogous to the relative tensor product of ordinary modules, where maps $M \otimes_R N \to X$ are the same as maps $M \otimes N \to X$ which "balance" the R-actions. Of course, this is not just a property, it also has to be expressed as homotopy coherent structure. A useful formula is that $M \otimes_R N$ is equivalent to a colimit of a simplicial diagram $\mathrm{Fun}(\Delta^{\mathrm{op}}, \mathcal{C})$ of the form

$$M \otimes N \leftleftarrows M \otimes R \otimes N \,\substack{\longleftarrow\\[-1ex]\longleftarrow\\[-1ex]\longleftarrow}\, \ldots.$$

The relative tensor product behaves in the usual way (for example, there are natural $M \otimes_R R \simeq M$). If R is commutative, one has an equivalence between $\mathrm{LMod}_R(\mathcal{C})$ and $\mathrm{RMod}_R(\mathcal{C})$, and we sometimes just write $\mathrm{Mod}_R(\mathcal{C})$ in that case. The resulting relative tensor product $\otimes_R$ on Mod_R can be upgraded to a symmetric-monoidal structure on $\mathrm{Mod}_R(\mathcal{C})$, that is $M \otimes_R N$ still carries a canonical R-module structure, and $\otimes_R$ has all the associativity and commutativity coherences necessary for a symmetric-monoidal structure.

In the case of ring spectra specifically, for a commutative ring spectrum R the mapping spectrum functor $\mathrm{map}_{\mathrm{Mod}_R} : \mathrm{Mod}_R^{\mathrm{op}} \times \mathrm{Mod}_R^{\mathrm{op}} \to \mathrm{Sp}$ can be upgraded to a functor to Mod_R (still denoted $\mathrm{map}_{\mathrm{Mod}_R}$), which is characterized by the adjunction

$$\mathrm{Map}_{\mathrm{Mod}_R}(M \otimes_R N, P) \simeq \mathrm{Map}_{\mathrm{Mod}_R}(M, \mathrm{map}_{\mathrm{Mod}_R}(N, P)),$$

that is, the symmetric-monoidal structure on Mod_R is *closed* (there is an "internal hom" right adjoint to $\otimes$ in the above sense).

For $R = \mathbb{S}$, every spectrum is uniquely a $\mathbb{S}$-module. So we have $\mathrm{Mod}_{\mathbb{S}}(\mathrm{Sp}) \simeq \mathrm{Sp}$. Under this equivalence, $\otimes_{\mathbb{S}}$ corresponds to the monoidal structure Sp we started with, and the internal hom $\mathrm{map}_{\mathrm{Mod}_{\mathbb{S}}}$ is just $\mathrm{map}_{\mathrm{Sp}}$. This allows us to think of spectra simply as module category. For emphasis, we will sometimes use the notation $\otimes_{\mathbb{S}}$ instead of just $\otimes$ for the monoidal structure on Sp, especially when there are other module structures around.

If R is connective, then $\mathrm{Mod}_R(\mathrm{Sp})$ has a t-structure where the connective and coconnective objects are simply the ones that go to connective and coconnective objects under the forgetful functor $\mathrm{Mod}_R(\mathrm{Sp}) \to \mathrm{Sp}$. Under the equivalence $\mathrm{Mod}_{HR}(\mathrm{Sp}) \simeq \mathcal{D}(R)$, this corresponds to the standard t-structure on the derived category (where connectivity and coconnectivity are defined in terms of homology).

Warning 7.2. We may view $\mathrm{Ab} \subseteq \mathrm{Sp}$ as full subcategory via H, and (by lax monoidality) also $\mathrm{Ring} \subseteq \mathrm{Alg}(\mathrm{Sp})$, $\mathrm{CRing} \subseteq \mathrm{CAlg}(\mathrm{Sp})$, and $\mathrm{Mod}_R(\mathrm{Ab}) \subseteq \mathrm{Mod}_{HR}(\mathrm{Sp})$ for an ordinary ring R. In that sense, ordinary algebra embeds fully faithfully into the world of spectra. Dropping the H from notation, we may for example think of ordinary rings as living under the sphere, as we have a ring map $\mathbb{S} \to \mathbb{Z}$. There are a number of notational challenges that arise here, though:

1. For an ordinary ring, $\mathrm{Mod}_R(\mathrm{Ab})$ and $\mathrm{Mod}_R(\mathrm{Sp})$ are not the same: The former is the usual abelian category of R-modules, while the latter is a stable ∞-category. It turns out that there is an equivalence $\mathrm{Mod}_R(\mathrm{Sp}) \simeq \mathcal{D}(R)$ to the derived category. This means that it can be ambiguous to simultaneously omit the H from notation and just write Mod_R instead of $\mathrm{Mod}_R(\mathrm{Sp})$. A solution employed by some people is to always write $\mathcal{D}(R) := \mathrm{Mod}_R(\mathrm{Sp})$ for modules over a ring spectrum, no matter if it is an Eilenberg-Maclane spectrum or not. This is elegant, but possibly confusing: While for general R, $\mathrm{Mod}_R(\mathrm{Sp})$ has a lot of similarities to the derived category of a ring, it is typically not literally the derived category of an abelian category. (We note however that the situation is similar in the algebraic geometry of stacks, where $\mathcal{D}(\mathcal{X})$ for a stack $\mathcal{X}$ is also not always the derived category of an abelian category, and the notation does not seem to cause offense there.)

2. The functor $\mathrm{Ab} \to \mathrm{Sp}$ is of course only lax symmetric monoidal, which should not surprise us any longer, since the monoidal structure on the right is $\otimes_{\mathbb{S}}$, while on the left it is $\otimes_{\mathbb{Z}}$. However, not even the functor $\mathrm{Ab} \to \mathrm{Mod}_{H\mathbb{Z}}(\mathrm{Sp})$ is symmetric monoidal, that is

$$HA \otimes_{H\mathbb{Z}} HB \not\simeq H(A \otimes_{\mathbb{Z}} B).$$

   What is true is that under the equivalence $\mathrm{Mod}_{H\mathbb{Z}}(\mathrm{Sp}) \simeq \mathcal{D}(\mathbb{Z})$, the monoidal structure $\otimes_{H\mathbb{Z}}$ corresponds to the *derived* tensor product in $\mathcal{D}(\mathbb{Z})$. In particular, $\pi_1(HA \otimes_{H\mathbb{Z}} HB) \cong \mathrm{Tor}_1^{\mathbb{Z}}(A, B)$. More generally, the monoidal structure in $\mathrm{Mod}_{HR}(\mathrm{Sp})$ corresponds to the derived tensor product in $\mathcal{D}(R)$. This means that if we drop all H's, something like $\mathbb{F}_p \otimes_{\mathbb{Z}} \mathbb{F}_p$ becomes quite ambiguous. In abstract arguments, this usually does not cause problems, but in any sort of computation we definitely need to involve both tensor products over ring spectra and underived tensor products of abelian groups at the same time, for example imagine expressing

$$\pi_0(HA \otimes_{H\mathbb{Z}} HB) \cong A \otimes_{\mathbb{Z}} B$$

without the H's (and the resulting convention that the $\pi_0(\ldots)$ on the left could also mean an Eilenberg-MacLane spectrum).

A solution analogous to the previous point would be to use some symbol like $\otimes^L$ for all tensor products over ring spectra, but that defeats the purpose of streamlining the notation. Some people have (somewhat jokingly) suggested instead introducing separate notation ($\otimes^U$ for "underived", or $\otimes^\heartsuit$ because we can think of the underived tensor product of R-modules as what we obtain by truncating the derived one back into the heart of $\mathcal{D}(R)$), but ultimately the best solution seems to be to be more explicit with H's whenever confusion approaches.

7.2 Künneth and UCT

The functor $\pi_* : \mathrm{Sp} \to \mathrm{gr\,Ab}$ to *graded abelian groups* is lax symmetric monoidal, where the symmetric monoidal structure on $\mathrm{gr\,Ab}$ is given by the usual tensor product of graded abelian groups (where $(A \otimes B)_n = \bigoplus_{i+j=n} A_i \otimes B_j$), and the braiding $A \otimes B \to B \otimes A$ is given by the Koszul sign rule, $a \otimes b \mapsto (-1)^{|a||b|} b \otimes a$. Informally, this comes from the maps $\pi_n(X) \otimes \pi_m(Y) \to \pi_{n+m}(X \otimes Y)$ which tensor together two maps $\mathbb{S}^n \to X$ and $\mathbb{S}^m \to Y$ to get a map $\mathbb{S}^{n+m} \to X \otimes Y$, with the Koszul sign forced upon us by the fact that the "flip" map $\mathbb{S}^n \otimes \mathbb{S}^m \to \mathbb{S}^m \otimes \mathbb{S}^n$ has degree $(-1)^{nm}$.

In particular, if R is a ring spectrum, then $\pi_*(R)$ is a graded ring (graded-commutative if R was commutative), and if M is an R-module, then $\pi_*(M)$ is a (graded) $\pi_*(R)$-module. For a graded abelian group A, we write $A(n)$ for the "degree shift by n", i.e. the graded abelian group with $A(n)_k = A_{k-n}$. Note that

$$\pi_*(\Sigma^n X) = \pi_*(X)(n).$$

Lemma 7.3. *Let R be an associative ring spectrum and M a left R-module. If $\pi_*(M)$ is a free graded $\pi_*(R)$-module with basis b_i (meaning that the b_i are homogeneous elements such that the map $\bigoplus_i \pi_*(R)(|b_i|) \to \pi_*(M)$ is an isomorphism), then we have an equivalence of R-modules*

$$\bigoplus_i \Sigma^{|b_i|} R \to M$$

Proof. The b_i are represented by maps $\mathbb{S}^{|b_i|} \to M$ of spectra, so by the adjunction, by R-module maps

$$\Sigma^{|b_i|} R \to M.$$

These induce the maps $\pi_*(R)(|b_i|) \to \pi_*(M)$ corresponding to b_i, and so their sum

$$\bigoplus_i \Sigma^{|b_i|} R \to M.$$

is an equivalence (by Whitehead). □

The tensor product of graded abelian groups is a closed monoidal structure, with the internal $\operatorname{Hom}(A_*, B_*)$ given by the graded abelian group whose degree n part is

$$\prod_k \operatorname{Hom}(A_k, B_{n+k}) = \operatorname{Hom}_{\mathrm{gr\,Ab}}(A_*(n), B_*),$$

i.e. homogeneous degree $+n$ maps $A_* \to B_*$.

On modules over a graded ring R_*, the tensor product of graded abelian groups gives rise to a relative tensor product, with

$$M_* \otimes_{R_*} N_*$$

a suitable quotient of $M_* \otimes N_*$. We also have a $\operatorname{Hom}_{R_*}(M_*, N_*)$ which is given by the graded subgroup of $\operatorname{Hom}(M_*, N_*)$ of those maps compatible with the module structure. For graded-commutative R_*, $\otimes_{R_*}$ upgrades to a closed symmetric-monoidal structure on graded R_*-modules, with $\operatorname{Hom}_{R_*}(M_*, N_*)$ its internal hom. Note that all of these formulas cannot be interpreted pointwise in $*$ (that is, in this notation we are not allowed to think of $*$ as representing an arbitrary integer anymore, it is just a symbol reminding us that all these are graded objects).

Corollary 7.4. *For a ring spectrum R and a left R-module M, if $\pi_*(M)$ is free as graded $\pi_*(R)$-module, then the canonical map*

$$\pi_* \operatorname{map}_{\mathrm{Mod}_R}(M, N) \to \operatorname{Hom}_{\pi_*(R)}(\pi_*(M), \pi_*(N))$$

is an isomorphism. Analogously, for a right R-module M with $\pi_*(M)$ free as graded $\pi_*(R)$-module, the canonical map

$$\pi_*(M) \otimes_{\pi_*(R)} \pi_*(N) \to \pi_*(M \otimes_R N)$$

is an isomorphism.

Proof. By Lemma 7.3, we may assume that M is a sum of shifts of R. In both cases, we may pull out coproducts and shifts, so we may assume that M is just R. In that case, it is clear, since $\operatorname{map}_{\mathrm{Mod}_R}(R, N) \simeq N$ and $R \otimes_R N \simeq N$. □

Corollary 7.5. *Let X be a spectrum, R a ring spectrum, and assume $R_*X = \pi_*(R \otimes X)$ is free as a graded $\pi_*(R)$-module. Then we have an isomorphism*

$$R^*X \to \operatorname{Hom}_{\pi_*(R)}(R_*X, \pi_*(R))$$

Proof. We may write

$$\operatorname{map}(X, R) \simeq \operatorname{map}_{\mathrm{Mod}_R}(R \otimes X, R),$$

and use Corollary 7.4 to get

$$R^*X = \pi_{-*} \operatorname{map}_{\mathrm{Mod}_R}(R \otimes X, R) = \operatorname{Hom}_{\pi_*(R)}(R_*X, \pi_*(R))$$

□

Corollary 7.6. *Let X, Y be spectra, R a commutative ring spectrum, and assume one of R_*X or R_*Y is free as graded $\pi_*(R)$-module. Then the canonical map*

$$R_*X \otimes_{\pi_*(R)} R_*Y \to R_*(X \otimes Y)$$

is an isomorphism.

Proof. We similarly write

$$R \otimes X \otimes Y \simeq (R \otimes X) \otimes_R (R \otimes Y),$$

and use Corollary 7.4. □

For Hk for k a field, the freeness conditions above are automatic. In those cases, we recover the familiar UCT and Künneth formulas

$$H^*(X;k) \simeq \mathrm{Hom}_k(H_*(X;k), k)$$
$$H_*(X \times Y;k) \simeq H_*(X;k) \otimes_k H_*(Y;k)$$

by applying the above to suspension spectra $\Sigma^\infty_+ X$, $\Sigma^\infty_+ Y$.

Now recall that these have versions over $\mathbb{Z}$, without any freeness assumption. Instead, the failure of the above comparison maps to be equivalences is measured by Ext and Tor terms. There is an analogous general statement about modules over a ring spectrum, with the short exact sequences being replaced by spectral sequences:

Theorem 7.7. *For a ring spectrum R and left R-modules M and N, there is a spectral sequence*

$$\mathrm{Ext}^{s,t}_{\pi_*(R)}(\pi_*(M), \pi_*(N)) \Rightarrow \pi_{t-s}\,\mathrm{map}_{\mathrm{Mod}_R}(M, N).$$

Here the Ext is a graded analogue of the usual Ext. It can for example be defined by resolving $\pi_*(M)$ by graded free $\pi_*(R)$-modules (in the abelian category of graded $\pi_*(R)$-modules), and applying the graded $\mathrm{Hom}_{\pi_*(R)}(-, \pi_*(N))$. The result is a cochain complex of graded abelian groups, whose s-th cohomology $\mathrm{Ext}^s_{\pi_*(R)}(\pi_*(M), \pi_*(N))$ is a graded abelian group, and $\mathrm{Ext}^{s,t}$ above refers to its degree t part.

We have analogously a bigraded Tor and a spectral sequence for relative tensor products:

Theorem 7.8. *For a ring spectrum R and right and left R-modules M and N, there is a spectral sequence*

$$\mathrm{Tor}^{\pi_*(R)}_{s,t}(\pi_*(M), \pi_*(N)) \Rightarrow \pi_{t+s}(M \otimes_R N).$$

As special cases, we have a *universal coefficient spectral sequence*

$$\mathrm{Ext}^{s,t}_{\pi_*(R)}(R_*X; \pi_*(R)) \Rightarrow R^{s-t}(X),$$

and a *Künneth spectral sequence*

$$\mathrm{Tor}_{s,t}^{\pi_*(R)}(R_*X, R_*Y) \Rightarrow R_{s+t}(X \otimes Y),$$

generalizing the above forms of the UCT and Künneth theorem. Their usefulness diminishes if $\pi_*(R)$ is very far from being a field and $\pi_*(M)$ very far from being a free module, but for example in the case of $R = \mathbb{Z}$, the fact that $\mathbb{Z}$ has global dimension 1 means that the Ext and Tor terms vanish for $s > 1$, in which case the spectral sequences degenerate to useful exact sequences analogous to the classical universal coefficient and Künneth theorems for $\mathbb{Z}$-homology.

Both spectral sequences can be constructed in different ways, the easiest being to use a free resolution of $\pi_*(M)$ to build a filtration on M by graded free R-modules (sums of shifts of R), obtain a filtration on $\mathrm{map}_{\mathrm{Mod}_R}(M, N)$ or $M \otimes_R N$, and use Corollary 7.4 to identify the homotopy groups of the associated graded with the complex computing Ext or Tor with the chosen free resolution. There is also a more canonical construction obtained by viewing the Whitehead tower $\tau_{\geq *}R$ as algebra in the category of filtered spectra $\mathrm{Fun}(\mathbb{Z}^{\mathrm{op}}, \mathrm{Sp})$ (with monoidal structure given by *Day convolution*), $\tau_{\geq *}M$ and $\tau_{\geq *}N$ as modules over it, and forming the spectral sequence of the resulting filtered spectra (taking a relative tensor product or internal hom). For now, we will not talk about this additional structure on filtered spectra, but we will revisit this point of view later.

7.3 The dual Steenrod algebra

Since we can write

$$\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} X \simeq (\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p) \otimes_{\mathbb{F}_p} (\mathbb{F}_p \otimes_{\mathbb{S}} X),$$

Künneth allows us to identify

$$\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} X) \simeq \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p) \otimes_{\mathbb{F}_p} \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} X).$$

Of course, the second factor is $H_*(X; \mathbb{F}_p)$. The first factor is $H_*(H\mathbb{F}_p; \mathbb{F}_p)$, whose dual is $H^*(H\mathbb{F}_p; \mathbb{F}_p)$, the Steenrod algebra $\mathcal{A}_p$ we identified earlier. We will fix p and (by a standard overloading of notation) simply write $\mathcal{A}$ for the Steenrod algebra. Since it is dual to the homology of $H\mathbb{F}_p$, one calls $H_*(H\mathbb{F}_p; \mathbb{F}_p)$ the *dual Steenrod algebra* and writes $\mathcal{A}_*$. Of course, the fact that degreewise both of these are finite-dimensional means that we can dualize back and forth.

The map

$$\mathbb{F}_p \otimes_{\mathbb{S}} X \to \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} X$$

obtained by using the unit map $\mathbb{S} \to \mathbb{F}_p$ in the middle, gives an interesting map

$$H_*(X; \mathbb{F}_p) \to \mathcal{A}_* \otimes_{\mathbb{F}_p} H_*(X; \mathbb{F}_p)$$

for any spectrum X. For $X = \mathbb{F}_p$ itself, we get a map

$$\mathcal{A}_* \to \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*.$$

This "comultiplication" is "coassociative", meaning the two maps $\mathcal{A}_* \to \mathcal{A}_*^{\otimes 3}$ we get by iterating this in two different ways agree, as one sees by writing both as induced by the map

$$\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \to \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p$$

which uses the unit map twice in the middle. Similarly, we get a "counit" $\mathcal{A}_* \to \mathbb{F}_p$ from the multiplication map $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \to \mathbb{F}_p$, and combining the comultiplication with the counit either way gives the identity on $\mathcal{A}_*$. In total, we learn that $\mathcal{A}_*$ is a *coalgebra*. The map we considered first exhibits every $H_*(X; \mathbb{F}_p)$ as $\mathcal{A}_*$-*comodule*.

Proposition 7.9. *The comultiplication map*

$$\Delta : \mathcal{A}_* \to \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*$$

is dual to the multiplication map

$$\mathcal{A} \otimes_{\mathbb{F}_p} \mathcal{A} \to \mathcal{A},$$

which we defined earlier in terms of composition, and the coaction map

$$\psi : H_*(X; \mathbb{F}_p) \to \mathcal{A}_* \otimes_{\mathbb{F}_p} H_*(X; \mathbb{F}_p)$$

is dual to the action map

$$\mathcal{A} \otimes_{\mathbb{F}_p} H^*(X; \mathbb{F}_p) \to H^*(X; \mathbb{F}_p)$$

Proof. The first claim is a special case of the second, for $X = \mathbb{F}_p$. So it suffices to prove the second. We need to prove the following: Given an element $\theta \in \mathcal{A}$, an element $\alpha \in H^*(X; \mathbb{F}_p)$, and an element $x \in H_*(X; \mathbb{F}_p)$, then

$$\langle \theta \otimes \alpha, \psi(x) \rangle = \langle \theta(\alpha), x \rangle,$$

where $\langle -, - \rangle$ denotes the duality pairings between $\mathcal{A} \otimes_{\mathbb{F}_p} H^*(X; \mathbb{F}_p)$ and $\mathcal{A}_* \otimes_{\mathbb{F}_p} H_*(X; \mathbb{F}_p)$ and $H^*(X; \mathbb{F}_p)$ and $H_*(X; \mathbb{F}_p)$ respectively.

By the adjunction for modules, $\theta : \mathbb{F}_p \to \Sigma^n \mathbb{F}_p$ also corresponds to a left $\mathbb{F}_p$ module map $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \to \Sigma^n \mathbb{F}_p$, which on homotopy is the map $\mathcal{A}_* \to \mathbb{F}_p(n)$ ((n) denoting the degree shift) given by $\langle \theta, - \rangle$. Similarly, α corresponds to a module map $\mathbb{F}_p \otimes_{\mathbb{S}} X \to \Sigma^m \mathbb{F}_p$. Now we get a commutative diagram

$$\begin{array}{ccccccc}
X & & \xrightarrow{\quad\quad\quad \mathrm{id} \quad\quad\quad} & & X & \xrightarrow{\mathrm{id}} & X \\
\downarrow & & & & \downarrow{\scriptstyle \mathrm{id} \otimes \alpha} & & \downarrow{\scriptstyle \theta \circ \alpha} \\
\mathbb{F}_p \otimes_{\mathbb{S}} X & \longrightarrow & \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} X & \xrightarrow{\mathrm{id} \otimes \alpha} & \mathbb{F}_p \otimes_{\mathbb{S}} \Sigma^m \mathbb{F}_p & \xrightarrow{\theta} & \Sigma^{n+m} \mathbb{F}_p
\end{array}$$

where the bottom composite is the left module map corresponding to $\langle \theta \otimes \alpha, \psi(-) \rangle$. The composite through the top however is $\theta \circ \alpha$, and so under the adjunction corresponds to the left $\mathbb{F}_p$-module map corresponding to $\langle \theta(\alpha), - \rangle$. This finishes the proof. □

Remark 7.10. There is a sign convention involved when defining the duality pairing $\mathcal{A} \otimes_{\mathbb{F}_p} H^*(X;\mathbb{F}_p) \otimes_{\mathbb{F}_p} \mathcal{A}_* \otimes_{\mathbb{F}_p} H_*(X;\mathbb{F}_p) \to \mathbb{F}_p$ of tensor products. The correct sign to choose here is informed by the Koszul sign rule, since in order to pair up $\mathcal{A}$ and $\mathcal{A}_*$, we need to commute the middle two terms. In the above proof, which sign rule appears is hidden in the somewhat lax treatment of the shifts that appear.

More formally, one can deal with this by instead producing a commutative square with

$$\mathrm{map}_{\mathrm{Mod}_{\mathbb{F}_p}}(\mathbb{F}_p \otimes \mathbb{F}_p, \mathbb{F}_p) \otimes \mathrm{map}_{\mathrm{Mod}_{\mathbb{F}_p}}(\mathbb{F}_p \otimes X, \mathbb{F}_p) \otimes \mathbb{F}_p \otimes X$$

in the top left corner, and $\mathbb{F}_p$ (without shift) in the bottom right. While the analogue of the upper composite only involves applying various evaluation pairings between these, the lower composite involes swapping factors $\mathrm{map}(X, \mathbb{F}_p)$ and $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p$. When analyzing the effect of this composite on homotopy groups on

$$\theta \otimes \alpha \otimes x \in \pi_* \mathrm{map}_{\mathrm{Mod}_{\mathbb{F}_p}}(\mathbb{F}_p \otimes \mathbb{F}_p, \mathbb{F}_p) \otimes \pi_* \mathrm{map}_{\mathrm{Mod}_{\mathbb{F}_p}}(\mathbb{F}_p \otimes X, \mathbb{F}_p) \otimes \pi_*(\mathbb{F}_p \otimes X),$$

the sign that appears is consistent with the Koszul sign convention.

So the coalgebra structure on $\mathcal{A}_*$ that appears naturally from thinking about these as homotopy of $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p$ actually is dual to the algebra structure on $\mathcal{A}$ that comes from the interpretation as cohomology operations. On $\mathcal{A}_*$, we also have an algebra structure, from the fact that $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p$ is a ring spectrum. It is with respect to that ring structure that we identified

$$\mathcal{A}_* = \mathbb{F}_2[\zeta_1, \zeta_2, \ldots]$$

at $p = 2$, and $\Lambda_{\mathbb{F}_p}(\tau_0, \tau_1, \ldots)[\xi_1, \xi_2, \ldots]$ at odd p. As the ring structure is induced on homology by Künneth and the multiplication map $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \to \mathbb{F}_p$, the same ingredients that dually on cohomology gave us the diagonal map $\mathcal{A} \to \mathcal{A} \otimes \mathcal{A}$ which gave rise to the Cartan formula (we proved that it takes $\mathrm{Sq}^n \mapsto \sum_{i+j=n} \mathrm{Sq}^i \otimes \mathrm{Sq}^j$, $P^n \mapsto \sum_{i+j=n} P^i \otimes P^j$, and $\beta \mapsto \beta \otimes 1 + 1 \otimes \beta$), the multiplication map on $\mathcal{A}_*$ dualizes to that comultiplication on $\mathcal{A}$. On the $\mathcal{A}_*$ side, it is also clear that the map inducing the comultiplication is a map of ring spectra, so it is compatible with the multiplication and both $\mathcal{A}_*$ and $\mathcal{A}$ have the structure of *bialgebras*. In fact, they are Hopf algebras, which are bialgebras with a certain type of involution map. This involution is seen easiest on the $\mathcal{A}_*$ side, where it is induced by the swap endomorphism of $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p$.

The coaction map $H_*(X;\mathbb{F}_p) \to \mathcal{A}_* \otimes_{\mathbb{F}_p} H_*(X;\mathbb{F}_p)$ is compatible with the cross product map on homology, in the sense that the following diagram commutes:

$$\begin{array}{ccc} H_*(X;\mathbb{F}_p) \otimes H_*(Y;\mathbb{F}_p) & \longrightarrow & \mathcal{A}_* \otimes H_*(X;\mathbb{F}_p) \otimes \mathcal{A}_* \otimes H_*(Y;\mathbb{F}_p) \\ \downarrow & & \downarrow \\ H_*(X \otimes Y;\mathbb{F}_p) & \longrightarrow & \mathcal{A}_* \otimes H_*(X \otimes Y;\mathbb{F}_p). \end{array}$$

Here the right vertical map is a combination of swapping the factors and multiplication map, and it is not hard to obtain this from a similar diagram of spectra. Dually, this implies that the action of $\mathcal{A}$ on cohomology relates cross products to the comultiplication on $\mathcal{A}$, which yields the general Cartan formula

$$\theta(x \times y) = \sum (-1)^{|\theta''|\cdot|x|} \theta'(x)\theta''(y)$$

if $\Delta(\theta) = \sum \theta' \otimes \theta''$. In the case where we can dualize back, both the homology coaction with its multiplicativity and the cohomology action with the Cartan formula encode the same information. This is the case if $\mathcal{A}_* \otimes H_*(X; \mathbb{F}_p)$ is finitely generated in each degree.

However, there is more: Every $\theta \in \mathcal{A}$ gives us a map $H^*(X; \mathbb{F}_p) \to H^*(X; \mathbb{F}_p)$, and by dualizing back, a dual map $\theta_* : H_*(X; \mathbb{F}_p) \to H_*(X; \mathbb{F}_p)$. Since dualizing is contravariant, $(\theta\rho)_* = \rho_*\theta_*$, and so this defines a *right action* of $\mathcal{A}$ on $H_*(X; \mathbb{F}_p)$, which decreases degrees by $|\theta|$. In fact, this can be done even if $H_*(X; \mathbb{F}_p)$ is not the dual of $H^*(X; \mathbb{F}_p)$, i.e. if the homology groups are not finitely generated: It is induced by the map $\mathbb{F}_p \otimes X \xrightarrow{\overline{\theta} \otimes \mathrm{id}_X} \Sigma^n \mathbb{F}_p \otimes X$, where $\overline{\theta}$ denotes the image of θ under Hopf algebra conjugation on $\mathcal{A}$.

Finally, if $H_*(X; \mathbb{F}_p) \otimes \mathcal{A}$ is finite-dimensional in each degree, which holds for example if the homology of $H_*(X; \mathbb{F}_p)$ is concentrated in finitely many degrees and finite-dimensional there, we may dualize again to obtain a right coaction

$$\lambda_* : H^*(X; \mathbb{F}_p) \to H^*(X; \mathbb{F}_p) \otimes \mathcal{A}_*$$

on cohomology. Even though out of the four structures we have considered, this is the most artificial one (for example, it is the only one that requires some finite-dimensionality assumption to exist, and so cannot really come from a general construction on spectra), it is also extremely useful: The Cartan formula translates back to this map just being multiplicative with respect to cross products, but since this is now about cohomology, it is also multiplicative with respect to cup products in cohomology of spaces. We can use this to identify the map $\mathcal{A}_* \to \mathcal{A}_* \otimes \mathcal{A}_*$.

Theorem 7.11 (Milnor)**.** *The comultiplication on $\mathcal{A}_*$ is given by*

$$\Delta(\zeta_n) = \sum_{i+j=n} \zeta_i^{2^j} \otimes \zeta_j.$$

for $p = 2$ and

$$\Delta(\xi_n) = \sum_{i+j=n} \xi_i^{p^j} \otimes \xi_j$$

$$\Delta(\tau_n) = \tau_n \otimes 1 + \sum_{i+j=n} \xi_j^{p^j} \otimes \tau_j$$

for odd p.

Proof. We have $K(\mathbb{F}_2, 1) \simeq \mathbb{R}P^\infty$, and so we may write $H^*(\mathbb{R}P^\infty; \mathbb{F}_2) = \mathbb{F}_2[\iota_1]$ and $H^*(\mathbb{R}P^n; \mathbb{F}_2) = \mathbb{F}_2[\iota_1]/\iota_1^{n+1}$.

Write $x_i \in H_i(\mathbb{R}P^n; \mathbb{F}_2)$ for the nonzero element. Of course, in the pairing between cohomology and homology, the ι_1^i form a dual basis to the x_i, that is $\langle \iota_1^i, x_j \rangle = \delta_{ij}$.

We defined the element $\zeta_k \in \mathcal{A}_*$ as image of x_{2^k} under the map

$$H_{2^k}(\mathbb{R}P^\infty; \mathbb{F}_2) \to H_{2^k-1}(H\mathbb{F}_2; \mathbb{F}_2).$$

This came from the fact that for I an admissible sequence, we have

$$\mathrm{Sq}^I \iota_1 = \begin{cases} \iota_1^{2^k} \text{ if } I = (2^{k-1}, \ldots, 1) \\ 0 \text{ otherwise.} \end{cases}$$

This means that under the pairing between $\mathcal{A}$ and $\mathcal{A}_*$, $\langle \mathrm{Sq}^I, \zeta_k \rangle$ is 1 if $I = (2^{k-1}, \ldots, 1)$, and 0 otherwise. So ζ_k is the dual basis element to $\mathrm{Sq}^{2^{k-1}} \cdots \mathrm{Sq}^1$ in the Serre-Cartan basis of $\mathcal{A}$ in terms of admissible sequences.

For the right homology action, we get

$$\langle \iota_1, x_n \cdot \mathrm{Sq}^I \rangle = \begin{cases} \iota_1^{2^k} \text{ if } n = 2^k \text{ and } I = (2^{k-1}, \ldots, 1) \\ 0 \text{ otherwise,} \end{cases}$$

and so expressing $\lambda_*(\iota_1)$ in the dual basis, we get

$$\lambda_*(\iota_1) = \sum_j \iota_1^{2^j} \otimes \zeta_j,$$

in $H^*(\mathbb{R}P^n; \mathbb{F}_2) \otimes \mathcal{A}_*$ for arbitrary n. Using the fact that λ_* is multiplicative, if we apply λ_* twice, it gives us

$$\sum_{k,l} \iota_1^{2^{i+j}} \otimes \zeta_i^{2^j} \otimes \zeta_j.$$

On the other hand, λ_* being a coaction means this expression must agree with

$$\sum_n \iota_1^{2^n} \otimes \Delta(\zeta_n).$$

Comparing coefficients, the formula for $\Delta(\zeta_n)$ follows.

For odd p, we proceed analogously, using instead skeleta of $K(\mathbb{F}_p, 1)$ and $K(\mathbb{Z}, 2) \simeq \mathbb{C}P^\infty$. The definition of ξ_i, together with the fact that the only $\mathcal{P}^I$ that act nontrivially on ι_2 in $H^*(K(\mathbb{Z}, 2); \mathbb{F}_p)$ are the $P^{p^{k-1}} \cdots P^1$, means that ξ_i is the dual basis element to that admissible sequence, and

$$\lambda_*(\iota_2) = \sum \iota_2^{p^k} \otimes \xi_k.$$

Similarly, the definition of τ_i, together with the fact that the only $\mathcal{P}^I$ that act nontrivially on $\iota_1 \in H^*(K(\mathbb{F}_p, 1); \mathbb{F}_p)$ are the identity and the $P^{p^{k-1}} \cdots P^1\beta$, means that τ_i is the dual basis element to that admissible sequence, and

$$\lambda_*(\iota_1) = \iota_1 \otimes 1 + \sum (\beta\iota_1)^{p^k} \otimes \tau_k.$$

Furthermore, the only $\mathcal{P}^I$ that act nontrivially on $\beta\iota_1$ are again the $P^{p^{k-1}} \cdots P^1$, so we also have

$$\lambda_*(\beta\iota_1) = \sum (\beta\iota_1)^{p^k} \otimes \xi_k.$$

Using multiplicativity of λ_* and the coaction property, we again get

$$\Delta(\xi_n) = \sum_{i+j=n} \xi_i^{p^j} \otimes \xi_j$$

$$\Delta(\tau_n) = \tau_n \otimes 1 + \sum_{i+j=n} \xi_j^{p^j} \otimes \tau_j$$

□

This completely describes $\mathcal{A}_*$ as a Hopf algebra.

Remark 7.12. By the universal coefficient theorem, $H^2(K(\mathbb{F}_2, 1); \mathbb{Z}) = \mathbb{Z}/2$, with generator mapping to ι_1^2 in mod 2 cohomology. It follows that we have a map $K(\mathbb{F}_2, 1) \to K(\mathbb{Z}, 2)$ which on cohomology takes $\iota_2 \mapsto \iota_1^2$. It is in particular injective, and since

$$\lambda_*(\iota_1^2) = (\sum \iota_1^{2^k} \otimes \zeta_k)^2 = \sum \iota_1^{2^{k+1}} \otimes \zeta_k^2$$

in any skeleton of $K(\mathbb{F}_2, 1)$, we have that $\lambda_*(\iota_2) = \sum \iota_2^{2^k} \otimes \zeta_k^2$ in any skeleton of $K(\mathbb{Z}, 2)$. So, the $p = 2$ analogue of the ξ_i are the ζ_i^2.

Note that we had not described $\mathcal{A}$ itself completely: We computed an additive basis, but gave no explicit formulas for either the multiplication or the comultiplication except for some special cases. In particular, for $a < 2b$, $\mathrm{Sq}^a \mathrm{Sq}^b$ is not an admissible sequence, so it must be expressible in terms of *some* admissible sequence. We could now dualize back from our knowledge of the comultiplication on $\mathcal{A}_*$ to determine these relations, known as *Adem relations.* This is a slightly involved but completely algebraic exercise. Inductively, the Adem relations then allow us to express every non-admissible sequence of Sq^i's in terms of admissible sequences. All of this works analogously at odd p.

Milnor's computation also tells us that maybe the basis for $\mathcal{A}$ in terms of admissible sequences was not the best choice: In the dual Steenrod algebra, the formulas seem to be much nicer in terms of monomials in the ζ_i. The *Milnor basis* on $\mathcal{A}$ is obtained as the dual basis to the monomial basis on $\mathcal{A}_*$. It is related to the original basis of admissible sequences in a subtle way, but behaves in a much nicer way. For example, there is an explicit combinatorial formula for the multiplication of $\mathcal{A}$ in the Milnor basis, which can be carried out by hand and requires no inductive simplification like the application of the Adem relations. We refer the interested reader to [16, Section 6].

8 The Adams spectral sequence

8.1 Descent of modules

For a spectrum X, $\pi_*(X)$ is hard to understand, but $\pi_*(\mathbb{F}_p \otimes X)$ is just homology. We want to explore how much information about $\pi_*(X)$ we can recover from "homology information".

This is formally similar to a phenomenon in algebraic geometry called *descent*. In its easiest incarnation, it is about recovering an R-module M from the S-module $S \otimes_R M$ and additional information (of course, this requires assumptions on $R \to S$). A classical example is to describe $\mathbb{R}$ vector spaces in terms of $\mathbb{C}$-vector spaces. One way to do this is the following: For an $\mathbb{R}$-vector space V, consider the cochain complex

$$\mathbb{C} \otimes_{\mathbb{R}} V \to \mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} \otimes_{\mathbb{R}} V \to \mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} \otimes_{\mathbb{R}} V \ldots$$

where the first map is $z \otimes v \mapsto 1 \otimes z \otimes v - z \otimes 1 \otimes v$, and the later ones are similarly given by alternating sums of all the maps which insert a 1. Then it is not hard to see that this complex is exact in positive degrees, and that the leftmost map has V as kernel. In other words, V can be recovered completely by a diagram involving only complex vector spaces (with some of the maps involved of course not $\mathbb{C}$-linear).

We can attempt something similar with spectra. Let us see what happens for the pair of rings $\mathbb{S} \to \mathbb{F}_p$. Instead of cochain complexes, we look at the augmented cosimplicial diagram

$$\mathbb{S} \longrightarrow \mathbb{F}_p \rightrightarrows \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \Rrightarrow \ldots .$$

This is indexed over the category Δ_+ of totally ordered finite sets. Informally, this diagram takes $\{0, \ldots, n\}$ to $\mathbb{F}_p^{\otimes_{\mathbb{S}} n+1}$, with maps induced by maps $f : S \to T$ given by totally ordered finite sets by using the unit maps $\mathbb{S} \to \mathbb{F}_p$ for each $t \in T$ with $f^{-1}(t) = \emptyset$, and more generally multiplication maps $\mathbb{F}_p^{\otimes_{\mathbb{S}} f^{-1}(t)} \to \mathbb{F}_p$ for each t. Formally, this also involved coherences, but is easy to write down correctly in the language of algebras over operads. In analogy with the situation of descent in algebraic geometry, this diagram is called *Cech nerve* of the map $\mathbb{S} \to \mathbb{F}_p$.

As $\emptyset$ is initial in Δ_+, we may think of the above diagram as a cone over its restriction to Δ, the full subcategory consisting of all *nonempty* totally ordered sets. It therefore gives a map

$$\mathbb{S} \to \lim_{\Delta} \Big(\mathbb{F}_p \rightrightarrows \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \Rrightarrow \ldots . \Big)$$

Let us abbreviate this limit with $\lim_\Delta(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet+1})$. We also have full subcategories $\Delta_{\leq n} \subseteq \Delta$ on all totally ordered sets with $\leq n$ many elements, and restriced limits $\lim_{\Delta_{\leq n}}(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet+1})$. More generally, we can form the Cech nerve (and the various limits considered here) for any map $S \to R$ of commutative ring spectra, or more generally a commutative ring spectrum S and an associative S-algebra R. For an S-module M, we more generally have a version of the form

$$M \to \lim_{\Delta} \Big(R \otimes_S M \rightrightarrows R \otimes_S R \otimes_S M \,\substack{\rightarrow\\\rightarrow\\\rightarrow}\, \dots \Big)$$

There are general formulas for limits over $\Delta_{\leq n}$ and the fiber of restriction maps $\lim_{\Delta_{\leq n}} \to \lim_{\Delta_{\leq n-1}}$ (for general cosimplicial objects) in terms of limits of cube-shaped diagrams, which can be proved by arguments with Kan extensions. In the present case, these imply the following:

Proposition 8.1. 1. *The fiber of* $M \to \lim_{\Delta_{\leq n}}(R^{\otimes_S \bullet+1} \otimes_S M)$ *is equivalent to*

$$\mathrm{fib}(S \to R)^{\otimes_S n+1} \otimes_S M$$

2. *The fiber of* $\lim_{\Delta_{\leq n}}(R^{\otimes_S \bullet+1} \otimes_S M) \to \lim_{\Delta_{\leq n-1}}(R^{\otimes_S \bullet+1} \otimes_S M)$ *is equivalent to*

$$R \otimes_S \mathrm{fib}(S \to R)^{\otimes_S n} \otimes_S M$$

Under certain connectivity assumptions, this will allow us to recover a module as limit over the Cech nerve:

Definition 8.2. For $S \to R$ a map of commutative ring spectra and M an S-module, we say M *descends along* $S \to R$ if

$$M \to \lim_{\Delta}(R^{\otimes_S \bullet+1} \otimes_S M)$$

is an equivalence.

Theorem 8.3. 1. *If the S-module M admits an R-module structure, M descends along $S \to R$.*

2. *Assume S and R are connective and $\pi_0(S) \to \pi_0(R)$ is an isomorphism. Then if M is a bounded below S-module, it descends along $S \to R$.*

3. *Assume S and R are connective, $\pi_0(S) \to \pi_0(R)$ is surjective, and $\pi_0(S)/p \to \pi_0(R)/p$ is an isomorphism. If M is a bounded below S-module all of whose homotopy groups are bounded p-torsion (meaning for each n there exists k such that every element of $\pi_n(M)$ is annihilated by p^k), then M descends along $S \to R$.*

Proof. The first part uses that if M admits an R-module structure, the cosimplicial diagram may be extended with "extra degeneracies" arising from the maps $R \otimes_S M \to M$, forcing the limit to be M (see [14, Lemma 6.1.3.16] for details).

Now for any S-module M, let us write $F(M)$ for the fiber of

$$M \to \lim_{\Delta}(R^{\otimes_S \bullet+1} \otimes_S M).$$

Evidently, M descends along $S \to R$ if and only if $F(M) = 0$. Since $\otimes$ and limits are exact, F is, too, and so the modules which descend are closed under fibers and cofibers (and hence also shifts).

Since S and R are connective, we have for M concentrated in degree 0 that S- and R-module structures on M are the same as a $\pi_0(S)$- and $\pi_0(R)$-module structures on $\pi_0(M)$, and so if $\pi_0(S) \to \pi_0(R)$ is an isomorphism, any S-module which is concentrated in degree 0 is automatically an R-module. With the first part, it follows by an induction over the Postnikov tower that any M which is bounded (bounded below and bounded above) descends along $S \to R$.

Now, for any M which is bounded below, and arbitrary n, consider the (co)fiber sequence

$$F(\tau_{\geq n} M) \to F(M) \to F(\tau_{\leq n-1} M).$$

As we just saw, the right hand term vanishes since $\tau_{\leq n-1} M$ is bounded below and above. So the left hand map is an equivalence. But from the interpretation of Proposition 8.1, if M is a-connective, all of the fibers of $M \to \lim_{\Delta_{\leq n}}(R^{\otimes_S \bullet+1} \otimes_S M)$ are also a-connective, since $\mathrm{fib}(S \to R)$ is connective by assumption. Their limit is $F(M)$, and from the way sequential limits interact with connectivity, $F(M)$ is then $(a-1)$-connective. Applying this to $\tau_{\geq n} M$, we see that $F(\tau_{\geq n} M)$ is $(n-1)$-connective. Altogether, it follows that $F(M)$ is $(n-1)$-connective for all n, so it is zero.

For the third part, one argues similarly, using that for M concentrated in degree 0 with $\pi_0(M)$ annihilated by p, S- and R-module structure on M are the same as $\pi_0(S)/p$ and $\pi_0(R)/p$-module structures on $\pi_0(M)$. So those M always descent along $S \to R$. If M more generally is an S-module concentrated in degree 0 and $\pi_0(M) =: A$ is bounded p-torsion, we have a finite filtration of ordinary $\pi_0(S)$-modules $p^i A$, where the associated graded terms are all annihilated by p. Inductively, this shows that all those descend along $S \to R$ as well. Then, one proceeds as in the second part to conclude that all M which are bounded below and have bounded p-torsion homotopy groups descend as well. □

Since the latter condition is what holds in the case $\mathbb{S} \to \mathbb{F}_p$, we are curious if we can push the result a bit further. What can we say about $\lim_{\Delta}(\mathbb{F}_p^{\otimes \bullet+1} \otimes M)$ if M is connective, but its homotopy groups are not bounded p-torsion? Clearly, we cannot fully recover M: For n coprime to p, the multiplication by n map on M induces equivalences on $\mathbb{F}_p \otimes M$ and all other terms in the Cech nerve, so also on the limit, but not necessarily on M itself. The right notion for what information we can recover is that of p-completeness and p-completion:

Definition 8.4. A spectrum X is *p-complete* if

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \xrightarrow{p} X \xrightarrow{p} X) \simeq 0.$$

For a spectrum X, its *p-completion* $X_p^\wedge$ is the limit

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \to X/p^2 \to X/p),$$

where X/p^n denotes the cofiber of $p^n : X \to X$.

Remark 8.5. There is a cofiber sequence

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \xrightarrow{p} X \xrightarrow{p} X) \to X \to \lim_{\mathbb{N}^{\mathrm{op}}}(\dots \to X/p^2 \to X/p),$$

which shows that a spectrum is p-complete if and only if the map $X \to X_p^\wedge$ is an equivalence.

Definition 8.6. We write $\mathrm{Sp}_p^\wedge$ for the full subcategory of Sp on p-complete spectra.

Proposition 8.7. *The inclusion $\mathrm{Sp}_p^\wedge \to \mathrm{Sp}$ admits a left adjoint which takes X to $X_p^\wedge$. In particular, p-complete spectra are closed under limits.*

Proof. Write $F(X)$ for the fiber of $X \to X_p^\wedge$. As discussed above, this is given by

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \xrightarrow{p} X \xrightarrow{p} X).$$

We have $F(X)/p \simeq F(X/p) \simeq 0$: The homotopy groups $\pi_* \operatorname{map}(X/p, X/p)$ sit in a long exact sequence with the multiplication-by-p map on $\pi_* \operatorname{map}(X/p, X)$. It follows that $\pi_* \operatorname{map}(X/p, X/p)$ is p^2-torsion, so in particular the multiplication-by-p^2 map $X/p \to X/p$ is nullhomotopic. (Somewhat unintuitively, the $(-)^2$ is essential, at least for $p = 2$.) So in the diagram for $F(X/p)$, the composite of any two successive maps is 0, a cofinality argument then shows that the limit is 0.

The vanishing of $F(X)/p$ means that multiplication by p is an equivalence on $F(X)$, and thus also p^n. So $F(X)/p^n \simeq F(X/p^n) \simeq 0$ as well. As limits commute with limits, $F(X_p^\wedge) \simeq \lim F(X/p^n) \simeq 0$, so $X_p^\wedge$ is in fact p-complete.

Next, since $\operatorname{map}(F(X), Y)/p^n$ can be written both as $\operatorname{map}(F(X), Y/p^n)$ and $\operatorname{map}(\Sigma^{-1}F(X)/p^n, Y) \simeq 0$, we learn that $\operatorname{map}(F(X), Y/p^n) = 0$ for all n. If Y is p-complete, that means

$$\operatorname{map}(F(X), Y) \simeq \lim_n \operatorname{map}(F(X), Y/p^n) \simeq 0,$$

and so

$$\operatorname{map}(X_p^\wedge, Y) \simeq \operatorname{map}(X, Y).$$

By the pointwise criterion for adjoints, it follows that we indeed have a left adjoint given by $X_p^\wedge$. □

Let us call a map $X \to Y$ of spectra a *p-complete equivalence* if it becomes an equivalence $X_p^\wedge \to Y_p^\wedge$ after p-completion. We have the following:

Lemma 8.8. *A map $X \to Y$ is a p-complete equivalence if and only if the induced map $X/p \to Y/p$ is an equivalence.*

Proof. By passing to the cofiber of $X \to Y$, we may reduce to the following claim: For a spectrum C, $C_p^\wedge \simeq 0$ if and only if $C/p \simeq 0$. In one direction, assume that $C_p^\wedge \simeq 0$. This means that

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \xrightarrow{p} C \xrightarrow{p} C) \simeq C.$$

As we saw in the previous proof, p acts as equivalence on this limit, hence also on C, so $C/p \simeq 0$.

Conversely, if $C/p \simeq 0$, p acts as equivalence on C, so

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \xrightarrow{p} C \xrightarrow{p} C) \simeq C,$$

and as the completion is the cofiber of that map, $C_p^{\wedge} \simeq 0$. □

Lemma 8.9. *If R is a ring spectrum and $\pi_0(R)$ is bounded p-torsion, any R-module is p-complete.*

Proof. Choose n such that $p^n = 0$ in $\pi_0(R)$. Let M be an R-module. In the limit diagram

$$\lim_{\mathbb{N}^{\mathrm{op}}}(\dots \xrightarrow{p} M \xrightarrow{p} M),$$

the composite of any n successive maps is homotopic to the zero map. From this, a cofinality argument shows that the limit is 0. □

Theorem 8.10. *Let $S \to R$ be a map of connective commutative ring spectra, and assume that $\pi_0(R)$ is bounded p-torsion, $\pi_0(S) \to \pi_0(R)$ is surjective and $\pi_0(S)/p \to \pi_0(R)/p$ is an isomorphism. Then for a connective S-module M,*

$$M \to \lim_{\Delta}(R^{\otimes_S \bullet+1} \otimes_S M)$$

is a p-completion, meaning that the target is p-complete and the map is a p-complete equivalence.

Proof. As each of the terms in the limit is an R-module, they and their limit are p-complete. It remains to check that the map is a p-complete equivalence. We may check this mod p. As limits are exact, this gives

$$M/p \to \lim_{\Delta}(R^{\otimes_S \bullet+1} \otimes_S M/p).$$

But M/p satisfies the assumptions of Theorem 8.3, since it is connective and all its homotopy groups are p^2-torsion (by the long exact sequence). □

So the general statement we get for Cech nerves of maps such as $\mathbb{S} \to \mathbb{F}_p$ is that for connective M, the limit recovers exactly the p-completion $M_p^{\wedge}$. Fortunately, $\pi_*(M_p^{\wedge})$ is not too far off from $\pi_*(M)$ itself. For example, if all homotopy groups of M are finitely generated, $\pi_n(M_p^{\wedge})$ is just the usual algebraic p-completion of $\pi_n(M)$. In general, $\pi_n(M_p^{\wedge})$ is influenced by both $\pi_n(M)$ and $\pi_{n-1}(M)$, but nothing else. We refrain from making a precise statement here.

Proposition 8.11. *For a map $X \to Y$ of* bounded below *spectra, it is also true that it is a p-complete equivalence if and only if $\mathbb{F}_p \otimes X \to \mathbb{F}_p \otimes Y$ is an equivalence.*

Proof. To see this, we can pass to cofibers and reduce to showing that for a bounded below spectrum Z we have $Z/p = 0$ iff $\mathbb{F}_p \otimes Z = 0$. Observe that the left map in the cofiber sequence

$$Z \xrightarrow{p} Z \to Z/p$$

becomes 0 after tensoring with $\mathbb{F}_p$. This means that $\mathbb{F}_p \otimes (Z/p) \simeq \mathbb{F}_p \otimes Z \oplus \Sigma(\mathbb{F}_p \otimes Z)$.

Now if $Z/p \simeq 0$, $\mathbb{F}_p \otimes Z$ is 0, since it is a retract of this sum. Conversely, assume $\mathbb{F}_p \otimes Z = 0$, then also $\mathbb{F}_p \otimes (Z/p) = 0$. But by the above, Z/p descends along $\mathbb{S} \to \mathbb{F}_p$, so from this it also follows that $Z/p \simeq 0$. □

Remark 8.12. This is wrong without the connectivity assumption: There exist non-bounded below spectra which are already p-complete and nontrivial but satisfy $\mathbb{F}_p \otimes X = 0$. We will see later (when we discuss nilpotence) how these can be constructed by inverting non-nilpotent maps $\Sigma^n Y \to Y$ which are not seen by $\mathbb{F}_p$-homology.

8.2 The classical Adams spectral sequence

For a p-complete connective spectrum X, we obtain a limit diagram

$$X \simeq \lim_{\Delta} \mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet + 1} \otimes_{\mathbb{S}} X.$$

Taking the Whitehead filtration pointwise, we obtain a filtered spectrum

$$F_{\geq *}^{\mathrm{Adams}} X \simeq \lim_{\Delta} \tau_{\geq *}(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet + 1} \otimes_{\mathbb{S}} X).$$

Since limits commute with limits, it is complete. Its colimit is also clearly X: If X is n-connective, all $\mathbb{F}_p^{\otimes_{\mathbb{S}} k+1} \otimes_{\mathbb{S}} X$ are n-connective, and so the whole filtration becomes constant from $F_{\geq n}^{\mathrm{Adams}} X = X$ on.

What can we say about its associated graded? It is given by

$$\Sigma^{-i} \operatorname{gr}_i^{\mathrm{Adams}} X \simeq \lim_{\Delta} \pi_i(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet + 1} \otimes_{\mathbb{S}} X).$$

The homotopy groups of such a cosimplicial diagram of spectra which are concentrated in degree 0 agree with the cohomology groups of the corresponding cochain complex which has $\pi_i(\mathbb{F}_p^{\otimes_{\mathbb{S}} k+1} \otimes_{\mathbb{S}} X)$ in degree k, and differential given by the alternating sum of the "coface maps", the maps induced by the inclusions of totally ordered sets $\{0, \dots, k\} \to \{0, \dots, k+1\}$.

Here of course, $\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} X) = H_*(X; \mathbb{F}_p)$. Using Künneth, we may also describe

$$\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} X) \cong \mathcal{A}_* \otimes_{\mathbb{F}_p} H_*(X; \mathbb{F}_p),$$

as we saw in the section about the dual Steenrod algebra. More generally,

$$\pi_*(\mathbb{F}_p^{\otimes_{\mathbb{S}} k+1} \otimes_{\mathbb{S}} X) \cong \mathcal{A}_*^{\otimes_{\mathbb{F}_p} i} \otimes_{\mathbb{F}_p} H_*(X; \mathbb{F}_p).$$

The coface maps are given by inserting a 1 (in the case of the leftmost one), using the comultiplication on $\mathcal{A}_*$ (in case of the inner ones), and the coaction on $H_*(X;\mathbb{F}_p)$ (in case of the rightmost one). The resulting cochain complex is called the *cobar complex* of the comodule $H_*(X;\mathbb{F}_p)$, and its cohomology groups have an interpretation as Ext groups

$$\mathrm{Ext}^*_{\mathcal{A}_*}(\mathbb{F}_p, H_*(X;\mathbb{F}_p))$$

Note that the cobar complex is a cochain complex of graded abelian groups, so each Ext^s, the s-th cohomology group, is a graded abelian group. We write $\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, H_*(X;\mathbb{F}_p))$ for the homogeneous degree t part of the s-th Ext group.

Theorem 8.13 (The Adams spectral sequence)**.** *For connective p-complete X, there is a spectral sequence*

$$E^2_{n,s} = \mathrm{Ext}^{s,n+s}_{\mathcal{A}_*}(\mathbb{F}_p, H_*(X;\mathbb{F}_p)) \Rightarrow \pi_n(X).$$

This is usually displayed in "Adams grading", with d_r of bidegree $(-1, r)$, and the associated graded of the abutment filtration on $\pi_n(X)$ showing up in the n-th column.

As we will see in examples below, the filtration is no longer locally finite. This is not surprising, since all of the terms on the E^2 page are $\mathbb{F}_p$-vector spaces, so something like $\pi_0(\mathbb{S}^\wedge_p) \cong \mathbb{Z}_p$ cannot show up through a finite filtration. This is where some of the aforementioned convergence subtleties come in, but for our purposes we will just note that if all mod p homology groups of X are finitely generated, then for each (n, s) each $E^r_{n,s}$ stabilizes for some r, the resulting E^∞ page is isomorphic to the associated graded of the abutment filtration, and the abutment filtration is complete in the sense that

$$\pi_n(X) \cong \lim_s \pi_n(X)/F^s\pi_n(X).$$

8.3 Derived comodules

In order to work with the Adams E^2 page, we will want a better interpretation of what the cobar complex of the $\mathcal{A}_*$-comodule $H_*(X;\mathbb{F}_p)$ computes. The classical approach to this is to describe an abelian category of $\mathcal{A}_*$-comodules, derive it, and identify the cobar complex with the complex computing Ext in terms of a canonical injective resolution.

In this section, we want to take a slightly different perspective, and describe the stable ∞-category $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ in a way which is more directly connected to the cobar complex (and less obviously a derived category of an abelian category). While these perspectives are completely interchangeable in the case of the ordinary Adams spectral sequence, it will allow us to circumvent technical issues regarding the Adams-Novikov spectral sequence later, and also be very useful in the context of *synthetic spectra*.

If we start with the Cech nerve of $\mathbb{S} \to \mathbb{F}_p$ and apply π_*, we obtain a cosimplicial diagram of graded-commutative rings

$$\mathbb{F}_p \rightrightarrows \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p) \substack{\longrightarrow\\\longrightarrow\\\longrightarrow} \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p) \dots .$$

As discussed before, Künneth allows us to rewrite the entire diagram in terms of $\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p) \simeq \mathcal{A}_*$:

$$\mathbb{F}_p \rightrightarrows \mathcal{A}_* \mathrel{\substack{\rightarrow\\\rightarrow\\\rightarrow}} \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_* \ldots,$$

with the maps now encoding the coalgebra structure on $\mathcal{A}_*$.

If we apply $\mathrm{Mod}_{(-)}(\mathrm{gr\,Ab})$ to this diagram of graded rings, we get a cosimplicial diagram of categories

$$\mathrm{Mod}_{\mathbb{F}_p}(\mathrm{gr\,Ab}) \rightrightarrows \mathrm{Mod}_{\mathcal{A}_*}(\mathrm{gr\,Ab}) \mathrel{\substack{\rightarrow\\\rightarrow\\\rightarrow}} \mathrm{Mod}_{\mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*}(\mathrm{gr\,Ab}) \ldots,$$

with functors given by basechange along our maps of graded rings. What is its limit? We may informally think of an object in a limit of categories like that as a family of objects $X_0, X_1, \ldots$ in each of the individual categories, together with isomorphisms $F_{ij}(X_i) \cong X_j$ for every morphism in the diagram, a homotopy between the two resulting isomorphisms $F_{jk}F_{ij}(X_i) \cong X_k$ for each pair of composable morphisms, etc., but since these are 1-categories, the homotopy just becomes an equality and there are no higher conditions. In our case, this means an object in the limit consists of a graded $\mathbb{F}_p$-module M_0, a graded $\mathcal{A}_*$-module M_1, etc., together with a lot of isomorphisms: For example, the two maps $\mathbb{F}_p \to \mathcal{A}_*$ in our diagram (even though they are the same) give two isomorphisms $\mathcal{A}_* \otimes_{\mathbb{F}_p} M_0 \cong M_1$, etc. Identifying M_1 with $\mathcal{A}_* \otimes_{\mathbb{F}_p} M_0$ along, say, the map coming from the 0-th coface map $\{0\} \to \{0,1\}$, the other such isomorphism provides an interesting map $M_0 \to \mathcal{A}_* \otimes_{\mathbb{F}_p} M_0$. Similarly, we may identify $M_2 \cong \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_* \otimes_{\mathbb{F}_p} M_0$ along one of the maps in our diagram, all the other maps $M_1 \to M_2$ are then determined under these identifications in terms of the map $M_0 \to \mathcal{A}_* \otimes_{\mathbb{F}_p} M_0$ and $\mathcal{A}_* \to \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*$. Finally, the identities for compositions amount exactly to $M_0 \to \mathcal{A}_* \otimes_{\mathbb{F}_p} M_0$ giving M_0 the structure of a $\mathcal{A}_*$-comodule. That is, we have (very sketchily) verified that

$$\mathrm{Comod}_{\mathcal{A}_*} \simeq \lim\left(\mathrm{Mod}_{\mathbb{F}_p}(\mathrm{gr\,Ab}) \rightrightarrows \mathrm{Mod}_{\mathcal{A}_*}(\mathrm{gr\,Ab}) \mathrel{\substack{\rightarrow\\\rightarrow\\\rightarrow}} \mathrm{Mod}_{\mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*}(\mathrm{gr\,Ab}) \ldots \right).$$

For a graded ring R_*, we also have a derived category of modules, $\mathcal{D}^{\mathrm{gr}}(R_*)$, where objects are represented by chain complexes of objects of $\mathrm{Mod}_{R_*}(\mathrm{gr\,Ab})$. This is a stable ∞-category with a t-structure, whose heart is $\mathrm{Mod}_{R_*}(\mathrm{gr\,Ab})$. To not confuse ourselves with the word "degree" (which we for example sometimes also use for the index of homotopy groups), let us call the additional grading here "formal degree", and refer to things such as the index of homotopy groups and the chain complex degree "homotopical degree". One object of $\mathcal{D}^{\mathrm{gr}}(R_*)$ is the chain complex which is a copy of R_* in homotopical degree 0, let us more generally write $\Sigma^n R_*(m)$ for the object represented by the chain complex which is a free R_*-module on a generator in homotopical degree n and formal degree m. We have that $[\Sigma^n R_*(m), -]$ takes an object of $\mathcal{D}^{\mathrm{gr}}(R_*)$ to the formal degree m part of its n-th homology group, more generally $\mathrm{map}_{\mathcal{D}^{\mathrm{gr}}(R_*)}(R_*(m), -)$ takes an object of $\mathcal{D}^{\mathrm{gr}}(R_*)$ to a spectrum whose homotopy groups are the formal degree m part of its homology groups.

Definition 8.14. We let

$$\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*}) := \lim\Big(\mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p) \rightrightarrows \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*) \Rrightarrow \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*) \dots \Big).$$

An object in here a priori corresponds to a sequence of objects $X_0 \in \mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$, $X_1 \in \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*)$, etc., identifications between the various base-changes, and higher coherences (which now cannot be ignored). Given two such objects, $\mathrm{Map}(X, Y) \simeq \lim \mathrm{Map}(X_i, Y_i)$, i.e. mapping spaces in a limit of categories are the limit of the mapping spaces. Since universal properties of both limits and colimits can be expressed in terms of limits on the outside mapping spaces, limits and colimits in a limit of categories are computed "pointwise" as long as they are preserved by the functors in the diagram, and so the fact that all of these derived categories are stable and the functors exact implies that the limit is also stable.

Each of these derived categories comes with a t-structure, where the connectives are characterized by having trivial homology in negative (homotopical) degrees, and the coconnectives by trivial homology in positive (homotopical) degrees. Now crucially, in our case the maps from $\mathbb{F}_p \to \mathcal{A}_*^{\otimes_{\mathbb{F}_p} k}$ are faithfully flat. This means that an object $X_0 \in \mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$ is (co)connective if and only if its base-change in $\mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*^{\otimes_{\mathbb{F}_p} k})$ is. So, an object (X_i) in the limit has X_0 (co)connective if and only if all X_i are (co)connective. This implies that we have a t-structure on $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ where connectives and coconnectives are characterized "pointwise".

The heart of this t-structure is now clearly the full subcategory of all (X_i) where each X_i is in the heart, which means

$$\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})^{\heartsuit} := \lim\Big(\mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)^{\heartsuit} \rightrightarrows \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*)^{\heartsuit} \Rrightarrow \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*)^{\heartsuit} \dots \Big).$$

But this is precisely the previous limit description of $\mathrm{Comod}_{\mathcal{A}_*}$. In particular, this exhibits $\mathrm{Comod}_{\mathcal{A}_*}$ as a full subcategory of $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$, as it should be.

Let $\mathbb{F}_p$ be endowed with the trivial comodule structure, and M an arbitrary comodule. We may view them as objects of $\mathrm{Comod}_{\mathcal{A}_*}$. What is $\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(\mathbb{F}_p, M)$? Since $\mathbb{F}_p$ maps to the free module $\mathcal{A}_*^{\otimes k}(0)$ in $\mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*^{\otimes k})$, and M to $\mathcal{A}_*^{\otimes k} \otimes M$, we have

$$\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(\mathbb{F}_p, M) \simeq \lim_{\Delta} \mathrm{map}_{\mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*^{\otimes \bullet})}(\mathcal{A}_*^{\otimes \bullet}(0), \mathcal{A}_*^{\otimes \bullet} \otimes M),$$

with k-th term just being the homogeneous (formal) degree 0 part of $\mathcal{A}_*^{\otimes k} \otimes M$, in homotopical degree 0. the homotopy groups of the limit are therefore in nonpositive degrees, and computed by the formal degree 0 part of a cochain complex

$$M \to \mathcal{A}_* \otimes M \to \dots,$$

the cobar complex! More generally, $\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(\mathbb{F}_p(t), M)$ is related in the same way to the formal degree t part of the cobar complex. We write

$$\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(N, M) := \pi_{-s} \mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(N(t), M)$$

for ordinary comodules N and M.

To summarize, we have constructed a stable ∞-category $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ where ordinary comodules form a full subcategory, and the homotopy groups of map are computed by the cobar complex. Finally, we want to relate this to a more classical perspective on Ext as derived functor. To do so, consider the diagram

$$\begin{array}{ccccccc} & & \mathbb{F}_p & \rightrightarrows & \mathcal{A}_* & \rightrightarrows & \cdots \\ & & \downarrow & & \downarrow & & \\ \mathbb{F}_p & \longrightarrow & \mathcal{A}_* & \rightrightarrows & \mathcal{A}_* \otimes \mathcal{A}_* & \rightrightarrows & \cdots \end{array}$$

where the top row is the cosimplicial diagram considered before, and the bottom row is the coaugmented cosimplicial diagram where the cosimplicial part is obtained from the top diagram by inserting another $\mathcal{A}_*$ factor on the right of each term. The bottom diagram extends to a split cosimplicial diagram ("extra codegeneracy").

If we apply $\mathcal{D}^{\mathrm{gr}}$, the bottom diagram therefore has limit $\mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$. The vertical maps induce base-change functors on $\mathcal{D}^{\mathrm{gr}}$, which have right adjoints given by restriction functors. Crucially, these restrictions still commute with the basechange functors in the cosimplicial direction. So, on limits, we get an adjoint pair of functors $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*}) \rightleftarrows \mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$.

On the lower split cosimplicial diagram, the equivalence between the limit and $\mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$ is also given by the forgetful map to $\mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*)$ followed by the splitting map $\mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*) \to \mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$. The composite in the above diagram of rings $\mathbb{F}_p \to \mathcal{A}_* \to \mathbb{F}_p$ from the top left $\mathbb{F}_p$ to the bottom left $\mathbb{F}_p$ is the identity, and so the left adjoint of the adjunction we just constructed is simply the forgetful $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*}) \to \mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$. Its right adjoint takes a graded derived $\mathbb{F}_p$-module X to a comodule whose underlying $\mathbb{F}_p$-module is $\mathcal{A}_* \otimes_{\mathbb{F}_p} X$ by a similar diagram chase. We think of this right adjoint as forming "coinduced comodules", and will write it as $\mathcal{A}_* \otimes_{\mathbb{F}_p} -$.

The adjunction implies that $\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(X, \mathcal{A}_* \otimes Y) \simeq \mathrm{map}_{\mathcal{D}(\mathbb{F}_p)}(X, Y)$. If N and M are concentrated in (homotopical) degree 0, i.e. ordinary comodules, we can derive from this a formula for $\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(N, M)$. Indeed, if F_0 is an ordinary graded $\mathbb{F}_p$ vector space and $\mathcal{A}_* \otimes_{\mathbb{F}_p} F_0$ the corresponding coinduced comodule, and we choose $M \to \mathcal{A}_* \otimes_{\mathbb{F}_p} F_0$ a map of comodules which is injective, then the cofiber M_1 is again a comodule concentrated in (homotopical) degree 0. We may iterate this and obtain a sequence of cofiber sequences of ordinary comodules (i.e. short exact sequences) of the form

$$M_k \to \mathcal{A}_* \otimes F_k \to M_{k+1}.$$

Applying $\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(N(t), -)$ and taking homotopy groups, we get long exact sequences. As

$$\mathrm{map}_{\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})}(N(t), \mathcal{A}_* \otimes F_k) \simeq \mathrm{map}_{\mathcal{D}(\mathbb{F}_p)}(N(t), F_k)$$

is concentrated in homotopical degree 0 (and given by $\mathrm{Hom}_{\mathrm{Comod}_{\mathcal{A}_*}}(N(t), \mathcal{A}_* \otimes F_k) \cong \mathrm{Hom}_{\mathbb{F}_p}(N(t), F_k)$ there), these may be spliced together to prove the following:

Proposition 8.15. $\mathrm{Ext}^{s,t}(N, M)$ *may be computed as s-th cohomology group of the cochain complex obtained by applying* $\mathrm{Hom}_{\mathrm{Comod}_{\mathcal{A}_*}}(N(t), -)$ *to a resolution of* M *by coinduced comodules* $\mathcal{A}_* \otimes F_k$.

Note that contrary to the classical approach, where one would *define* Ext by a formula like above, here it arises just as a formula for computing the Ext which is defined a priori in terms of mapping spectra in the stable ∞-category $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$. In our later discussion of the Adams-Novikov spectral sequence, this perspective will pay off: There, $\mathbb{F}_p$ and $\mathcal{A}_*$ will be replaced by more complicated rings, and some of the homological algebra here becomes more subtle (in particular, the description of $\mathrm{Ext}(N, M)$ in terms of a resolution will require conditions on N), whereas the definition of the category of derived comodules as a limit of categories goes through unchanged. We will also see later that this limit of categories perspective leads us quite naturally to *synthetic spectra*.

9 Examples of Adams spectral sequences

9.1 $H\mathbb{F}_p$

As a very simple reality check, how does the Adams spectral sequence see $\pi_*(H\mathbb{F}_p)$?

For $X = H\mathbb{F}_p$, $\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} X) = \mathcal{A}_*$ is itself a coinduced comodule. It follows that

$$\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, H_*(X; \mathbb{F}_p)) = \begin{cases} \mathbb{F}_p \text{ if } (s,t) = (0,0) \\ 0 \text{ otherwise.} \end{cases}$$

The resulting Adams spectral sequence degenerates, and this of course just reflects the fact that $\pi_*(H\mathbb{F}_p)$ is $\mathbb{F}_p$ in degree 0 (picked up by the $(s,t) = (0,0)$ term on the E^∞ page).

9.2 $H\mathbb{Z}$

Next, we want to see a slightly more nontrivial example. Let $X = H\mathbb{Z}$. Then, the limit of its Cech nerve

$$\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z} \rightrightarrows \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z} \Rrightarrow \dots.$$

is, by the results in the previous section, the p-completion $X_p^\wedge \simeq \lim_n H\mathbb{Z}/p^n \simeq H\mathbb{Z}_p^\wedge$.

In order to see how the Adams spectral sequence sees $\pi_*(H\mathbb{Z}_p)$ (which consists of a single $\mathbb{Z}_p$ in degree 0), we need to understand $\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z}) \cong \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z}_p)$ as $\mathcal{A}_*$-comodule.

Proposition 9.1. *The canonical map*

$$\pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z}) \to \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p) = \mathcal{A}_*$$

is injective, with image the subring generated by ζ_1^2 and the elements $\zeta_i - \zeta_1\zeta_i^2$ with $i \geq 2$ for $p = 2$, and the subring generated by elements $\tau_i - \tau_0\xi_i$ for $i \geq 1$ and ξ_i for $i \geq 1$ for odd p.

Proof. The left map in the cofiber sequence

$$\mathbb{Z} \xrightarrow{p} \mathbb{Z} \to \mathbb{F}_p$$

becomes 0 after tensoring with $\mathbb{F}_p$. This means that the connecting homomorphisms in the long exact sequence on homotopy groups in the rotated cofiber sequence

$$\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z} \to \mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p \to \Sigma\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{Z}$$

are 0. We thus get a short exact sequence of comodules

$$0 \to H_*(H\mathbb{Z}_p; \mathbb{F}_p) \to H_*(H\mathbb{F}_p; \mathbb{F}_p) \to H_{*-1}(H\mathbb{Z}_p; \mathbb{F}_p). \to 0,$$

which already contains the injectivity statement we claimed. To identify the image, we can compose these maps the other way to get a degree -1 map of comodules

$$H_{*+1}(H\mathbb{F}_p; \mathbb{F}_p) \to H_*(H\mathbb{F}_p; \mathbb{F}_p)$$

with the homology $H_*(H\mathbb{Z}_p; \mathbb{F}_p)$ given as its image. As this is a comodule map $\mathcal{A}_*(-1) \to \mathcal{A}_*$ with coinduced target, the adjunction identifies it with a map $\mathcal{A}_*(-1) \to \mathbb{F}_p$ of graded $\mathbb{F}_p$ vector spaces. It cannot be the trivial map, as then the map $\mathcal{A}_*(-1) \to \mathcal{A}_*$ would be trivial too, and we would get $H_*(H\mathbb{Z}_p; \mathbb{F}_p) = 0$, contradicting for example Hurewicz. So it is a unit multiple of the map taking $\tau_0 \mapsto 1$ (or $\zeta_1 \mapsto 1$).

Under the adjunction, the corresponding comodule map can be described as composite

$$\mathcal{A}_*(-1) \to \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathcal{A}_*(-1) \to \mathcal{A}_* \otimes_{\mathbb{F}_p} \mathbb{F}_p \cong \mathcal{A}_*.$$

where the first map is the comultiplication, and the second uses $\tau_0 \mapsto 1$ (or $\zeta_1 \mapsto 1$).

It takes $\tau_i \mapsto \xi_i$ (with $\xi_0 = 1$ as usual), as we see from the formula for the comultiplication on $\mathcal{A}_*$, and similarly $\xi_i \mapsto 0$. At $p = 2$, it takes $\zeta_i \mapsto \zeta_{i-1}^2$. In fact, as $\mathcal{A}_*(-1) \to \mathbb{F}_p$ corresponds to the degree 1 element of the (non-dual) Steenrod algebra β (Sq^1 at $p = 2$), this map is the right action of β on $\mathcal{A}_*$. As β satisfies the derivation property, β is characterized as the unique derivation on $\mathcal{A}_*$ that takes $\xi_i \mapsto 0$ and $\tau_i \mapsto \xi_i$. At $p = 2$, it similarly is given by the derivation taking $\zeta_i \mapsto \zeta_{i-1}^2$.

For p odd, its image is the subring of $\mathcal{A}_*$ generated by all ξ_i and $\tau_i - \tau_0\xi_i$ for $i \geq 1$. Indeed, these elements are evidently in the kernel of the right action by β on $\mathcal{A}_*$, and the whole $\mathcal{A}_*$ is a rank 2 free module over this subring, on 1 and τ_0, because we can also write $\mathcal{A}_*$ as free graded-commutative algebra on τ_0, the

$\tau_i - \tau_0\xi_i$ and ξ_i. Because of the derivation property, β is an endomorphism of this free rank 2 module, and takes $1 \mapsto 0$ and $\tau_0 \mapsto 1$, so the kernel is exactly that subring and also coincides with the image of β.

At $p = 2$, we analogously get the subring of $\mathcal{A}_*$ generated by ζ_1^2 and $\zeta_i - \zeta_1\zeta_{i-1}^2$ for $i \geq 2$. In both cases, it is easy to check this: All these elements are in the kernel of β, and the whole $\mathcal{A}_*$ is a rank 2 free module over this subring, on 1 and τ_0 (resp. ζ_1). Because of the derivation property, β is a module map and takes $1 \mapsto 0$ and $\tau_0 \mapsto 1$, so the kernel is exactly that subring and coincides with the image. □

Remark 9.2. In terms of the antipode of the Hopf algebra structure on $\mathcal{A}_*$ (arising from the flip $\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{F}_p$), the above description becomes a bit easier: The antipode interchanges the right β-action with the left β-action, and the left β-action on $\mathcal{A}_*$ is very simple: It takes $\tau_0 \mapsto 1$, and all other generators τ_i, ξ_i to 0 (resp. $\zeta_1 \mapsto 1$ and all other $\zeta_i \mapsto 0$ at $p = 2$). By a similar argument as the above, its image is seen to be simply the subring generated by all τ_i with $i \geq 1$ and all ξ_i (ζ_1^2 and all ζ_i for $i \geq 2$ at $p = 2$). Under the antipode, this describes $\pi_*(\mathbb{F}_p \otimes \mathbb{Z}) \subseteq \mathcal{A}_*$ now as the subring generated by $\overline{\tau}_i$ and $\overline{\xi}_i$ for $i \geq 1$ (resp. $\overline{\zeta}_1^2$ and all $\overline{\zeta}_i$ for $i \geq 2$ for $p = 2$), where $\overline{(-)}$ denotes the antipode. These are not exactly the same generators as the ones found above, but they of course generate the same subring.

In order to compute $\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, H_*(H\mathbb{Z}; \mathbb{F}_p))$, we can splice copies of the short exact sequence

$$0 \to H_*(H\mathbb{Z}; \mathbb{F}_p) \to \mathcal{A}_* \to H_*(H\mathbb{Z}; \mathbb{F}_p)(1) \to 0$$

to get a coinduced resolution

$$0 \to H_*(H\mathbb{Z}; \mathbb{F}_p) \to \mathcal{A}_* \to \mathcal{A}_*(1) \to \mathcal{A}_*(2) \to \dots,$$

where the maps are given by the right action with β (Sq^1 at $p = 2$). This means that

$$\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, H_*(H\mathbb{Z}; \mathbb{F}_p)) = \begin{cases} \mathbb{F}_p \text{ if } s = t \text{ and } s \geq 0 \\ 0 \text{ otherwise.} \end{cases}$$

A different perspective which also explains some additional structure is obtained as follows. In $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$, we can view the exact sequences as cofiber sequences, and rotate them to give

$$\Sigma^{-1}M(1) \to M \to \mathcal{A}_*.$$

where we wrote M for $H_*(H\mathbb{Z}; \mathbb{F}_p)$, viewed as object of $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$. If we denote the map $\Sigma^{-1}M(1) \to M$ by h_0, then this represents an element of $\mathrm{Ext}^{1,1}_{\mathcal{A}_*}(M, M)$, and the above cofiber sequence shows $\mathrm{cofib}(h_0 : M \to M) \simeq \mathcal{A}_*$.

The resulting long exact sequence, together with the fact that $\mathrm{Ext}^{s,t}(\mathbb{F}_p, \mathcal{A}_*)$ is $\mathbb{F}_p$ in degree 0, shows the following:

Proposition 9.3.

$$\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, M) \cong \mathbb{F}_p[h_0],$$

compatible with the action of $h_0 \in \mathrm{Ext}^{*,*}_{\mathcal{A}_*}(M, M)$ *by composition.*

In fact, $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ has a symmetric-monoidal structure, coming from its description as limit of symmetric-monoidal categories, and the functors $\mathrm{Sp} \to \mathrm{Comod}_{\mathcal{A}_*}$, $\mathrm{Comod}_{\mathcal{A}_*} \to \mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ and $\mathrm{Ext}^{s,t}(\mathbb{F}_p, -)$ from $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ are lax symmetric monoidal. This also gives rise to a bigraded ring structure on the above Ext, which by an Eckmann-Hilton style argument has to be compatible with composition, so $\mathbb{F}_p[h_0]$ also describes the ring structure.

In the associated Adams spectral sequence, we still have no room for differentials (everything is concentrated in the column $t - s = 0$, but differentials decrease $t-s$ by 1). So the spectral sequence degenerates, and we must have an abutment filtration on $\pi_*(H\mathbb{Z}_p)$ with associated graded $\mathbb{F}_p[h_0]$. This can only be the p-adic filtration, where $p^s\mathbb{Z}_p$ is the filtration $\geq s$ part. In other words, the element $p \in \pi_0(H\mathbb{Z}_p)$ is detected in the $s = 1$ line of the Adams spectral sequence by h_0. (A priori, a unit multiple, but if one is a bit more precise about conventions, the coefficient comes out as 1).

In both these examples, we have combined some a priori knowledge of the homotopy groups to conclude a full description of the Adams spectral sequence. There are some additional interesting examples of this flavor, where one still can fully describe the entire Adams spectral sequence. These include for example ku and ko, the connective versions of (complex and real) topological K-theory.

We skip a discussion of these examples, and go straight to the most complicated example:

9.3 The sphere spectrum

Here, the homology $H_*(\mathbb{S}; \mathbb{F}_p) = \pi_*(\mathbb{F}_p \otimes_{\mathbb{S}} \mathbb{S})$ is very simple, it is given by a single $\mathbb{F}_p$ concentrated in degree 0. So the Adams E^2 page is given by $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, \mathbb{F}_p)$.

This can be computed by resolving $\mathbb{F}_p$ by coinduced comodules. In fact, there is a minimal way of doing this, and the resulting "minimal resolution" has generators exactly corresponding to a basis of $\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, \mathbb{F}_p)$. Dually, one can resolve $\mathbb{F}_p$ as a $\mathcal{A}$-module by free $\mathcal{A}$-modules in a minimal way, in general one has an isomorphism $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, M) \cong \mathrm{Ext}^{*,*}_{\mathcal{A}}(M^\vee, \mathbb{F}_p)$ if M is bounded below and degreewise finite-dimensional.

This gives an efficient way to compute the E^2 page of the Adams spectral sequence for the sphere (and other bounded below spectra) through a range of degrees $t \leq N$ for any N, and with some additional modifications (which amount to sometimes using $H_*(H\mathbb{Z}; \mathbb{F}_p)$ instead of only shifts of $\mathcal{A}_*$ in resolutions) even through a range of the form $t - s \leq N$, i.e. a range of finitely many columns. Notably, the resulting $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, \mathbb{F}_p)$ is quite complicated, and does not admit a global description (for example in the sense of generators and relations).

At $p = 2$, in the range $t - s \leq 10$, the Adams E^2 page looks as follows.

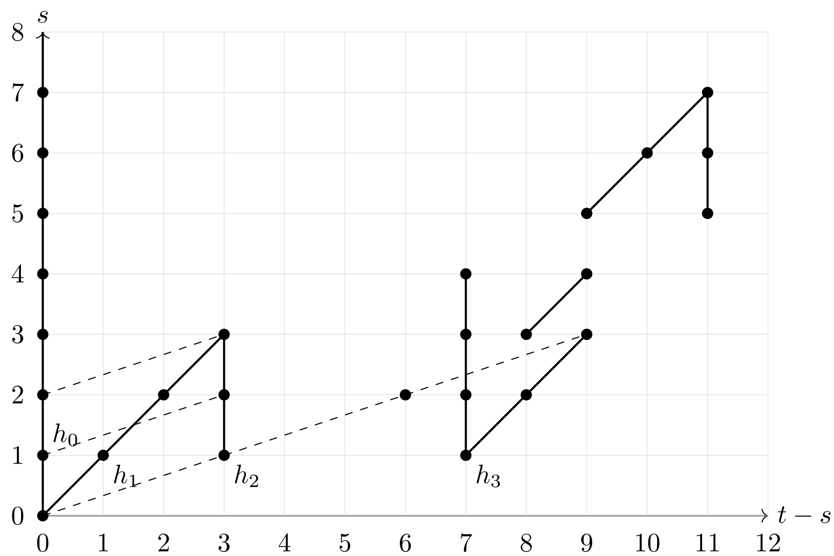

Figure 1: $\mathbb{F}_2$-based Adams spectral sequence for $\mathbb{S}$

Here, each dot represents a single $\mathbb{F}_2$ in the corresponding bidegree. Of course, there is no reason to expect $\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_2, \mathbb{F}_2)$ to be 1-dimensional over $\mathbb{F}_2$ for each s, t, and this is not the case in general, but happens to be so in the indicated range.

The elements h_i here are maps $h_i : \Sigma^{-1}\mathbb{F}_2(2^i) \to \mathbb{F}_2$ obtained by rotating the cofiber sequence obtained from the short exact sequence of comodules

$$0 \to \mathbb{F}_2 \to M_i \to \mathbb{F}_2(2^i) \to 0$$

where M_i is a comodule which is free as graded $\mathbb{F}_2$ vector space on elements a_0 and a_{2^i} with coaction given by $\psi(a_{2^i}) = 1 \otimes a_{2^i} + \zeta_1^{2^i} \otimes a_0$. It can be constructed for example as a subcomodule of $\mathcal{A}_*$ spanned by $1, \zeta_1^{2^i}$, since the comultiplication takes $\zeta_1^{2^k}$ to

$$(\zeta_1 \otimes 1 + 1 \otimes \zeta_1)^{2^i} = \zeta_1^{2^i} \otimes 1 + 1 \otimes \zeta_1^{2^i}.$$

This is dual to the $\mathcal{A}$-module on two basis elements connected by a nontrivial action by Sq^{2^i}. On the module side, the existence of this module is related to the fact that the Sq^{2^i} are indecomposable elements in $\mathcal{A}$.

The vertical, diagonal and dashed lines in the above chart indicate multiplication by h_0, h_1 and h_2 under the composition product on $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_2, \mathbb{F}_2)$, or the product structure arising from lax symmetric monoidal structures and the fact that $\mathbb{S}$ is an algebra. In the range indicated, that plus the fact that the differentials are derivations, actually rules out all possible differentials. For example, $d_2(h_1)$ must be zero, since

$$0 = d_2(h_0 h_1) = h_0 d_2(h_1)$$

by the derivation rule and $d_2(h_0) = 0$, but from degree $(t-s, s) = (0, 3)$ to $(0, 4)$ multiplication by h_0 is bijective, so also $d_2(h_1) = 0$.

Except for the h_0-tower in the $t - s = 0$ column, each column is finite. This tells us that each of the homotopy groups $\pi_n(\mathbb{S}_2^\wedge)$ for $n > 0$ are finite. Analogous phenomena happen at all odd primes as well. This does not directly tell us that $\pi_n(\mathbb{S})$ is finite for $n > 0$, since a priori there could be huge rational information that gets killed by the p-completion (for example, if $\pi_1(\mathbb{S})$ would be $\mathbb{Z}/2 \oplus \mathbb{Q}^\infty$, we would not see the huge rational summand in any of the p-completions). But an inductive argument with the Serre spectral sequence can be used to prove that all homotopy groups of spheres, and hence also all stable homotopy groups of spheres, are finitely generated. Then, finiteness of the homotopy groups of the p-completions really implies that $\pi_n(\mathbb{S}_p^\wedge)$ is a finite abelian group for each $n > 0$, and $\pi_n(\mathbb{S}_p^\wedge)$ is just the subgroup of $\pi_n(\mathbb{S})$ consisting of the p-power torsion elements. In particular, the Adams spectral sequence above tells us about the 2-power torsion part of $\pi_*(\mathbb{S})$.

Of course, $\pi_0(\mathbb{S}_2^\wedge) \cong \mathbb{Z}_2$, the 2-completion of $\pi_0(\mathbb{S}) \cong \mathbb{Z}$, which we know from Hurewicz. In the Adams spectral sequence, this is seen in the same way as $\pi_0(H\mathbb{Z}_2^\wedge)$, as associated graded of the 2-adic filtration, with 2 detected by h_0.

The fact that there is just a single $\mathbb{F}_2$ in the first column, generated by h_1 in degree $(t - s, s) = (1, 1)$, tells us that $\pi_1(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/2$. The unique nontrivial element in homotopy that is detected by h_1 is called η and is the stable version of the Hopf map $S^3 \to S^2$. Similarly, $\pi_2(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/2$, with the unique nontrivial element detected by h_1^2 in the Adams E^∞ page. Of course, η^2 is detected by h_1^2, and so η^2 is nonzero and generates $\pi_2(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/2$. Here, there are two ways of thinking about the multiplicative structure on $\pi_*(\mathbb{S}_2^\wedge)$: Either in terms of the symmetric-monoidal structure, using that the unit $\mathbb{S}$ is a ring spectrum and completion is lax symmetric monoidal, or as composition product on $[\Sigma^*\mathbb{S}, \mathbb{S}_2^\wedge] \cong [\Sigma^*\mathbb{S}_2^\wedge, \mathbb{S}_2^\wedge]$. It turns out both lead to the same ring structure on $\pi_*(\mathbb{S}_2^\wedge)$ using an Eckmann-Hilton argument.

In the third column, we see three $\mathbb{F}_2$'s on the E^∞ page, which implies that $\pi_3(\mathbb{S}_2^\wedge)$ has order 8. Let us write $\nu \in \pi_3(\mathbb{S}_2^\wedge)$ for an element detected by h_2, there are of course 4 different choices to make (a coset of the subgroup of index 2 consisting of all elements detected in $s \geq 2$), but one canonical choice comes from the quaternionic Hopf map $\nu : S^7 \to S^4$. As 2 is detected by h_0, we get that 2ν is detected by h_0h_2, and 4ν by $h_0^2h_2$, which are nonzero. So, $4\nu \neq 0$, and therefore we must have $\pi_3(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/8$, generated by ν. We also see in the E_2 page that $h_1^3 = h_0^2h_2$, and since there is nothing above those, i.e. the $E_{3,3}^\infty$ term (in $(t-s, s)$ grading) is really a subgroup of $\pi_3(\mathbb{S}_2^\wedge)$ instead of just a subquotient, this implies $\eta^3 = 4\nu$ in $\pi_3(\mathbb{S}_2^\wedge)$. It turns out that in this degree there is also a contribution from $p = 3$ integrally, $\pi_3(\mathbb{S}_3^\wedge) \cong \mathbb{Z}/3$, so $\pi_3(\mathbb{S}) \cong \mathbb{Z}/24$, and integrally one has $\eta^3 = 12\nu$.

Continuing similarly, one reads off $\pi_4(\mathbb{S}_2^\wedge) = 0$, $\pi_5(\mathbb{S}_2^\wedge) = 0$, and $\pi_6(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/2$. The octonionic Hopf map $\sigma : S^{15} \to S^8$ generates $\pi_7(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/16$.

In degree 8, the nonzero element in $E_{8,3}^\infty$ is represented by a unique $\varepsilon \in \pi_8(\mathbb{S}_2^\wedge)$, not decomposable as a product of earlier terms. We also see that the nonzero element of $E_{8,2}^\infty$ is h_1h_3 and detects $\eta\sigma$, so because $2\eta\sigma = 0$, we have that $\pi_8(\mathbb{S}_2^\wedge) \cong \mathbb{Z}/2 \oplus \mathbb{Z}/2$ generated by ε and $\eta\sigma$. The next group, $\pi_9(\mathbb{S}_2^\wedge)$, is similarly seen to be $\mathbb{Z}/2 \oplus \mathbb{Z}/2 \oplus \mathbb{Z}/2$, generated by $\eta^2\sigma$, $\eta\varepsilon$, and one other element μ_9.

As $h_2^3 = h_1^2 h_3$, ν^3 and $\eta^2\sigma$ have to agree modulo the subgroup of index 2 generated by $\eta\varepsilon$ and μ_1. The spectral sequence alone is incapable of telling us whether $\nu^3 - \eta^2\sigma$ is 0 or nontrivial, and in fact it is known that

$$\nu^3 - \eta^2\sigma = \eta\varepsilon.$$

Note that even though we had some freedom in picking the representatives ν and σ of h_2 and h_3, the value of $\nu^3 - \eta^2\sigma$ does not change under changing the representatives. So, this relation cannot be "defined away", and is a reminder that even in this range where all differentials vanish and the "additive extensions" are all decided by the fact that h_0 detects 2, the Adams spectral sequence still reflects imperfect information about $\pi_*(\mathbb{S}_2^\wedge)$.

One can carry on this analysis further, but starting with $\pi_{14}(\mathbb{S}_2^\wedge)$, nontrivial differentials appear, and the spectral sequence quickly becomes quite messy. Still, the Adams spectral sequence remains one of the main tools to compute stable homotopy groups of spheres to this day, and its structure has been determined roughly up to degree 90 at $p = 2$, and even further at bigger primes (where it becomes sparser, and various "small prime exceptions" go away).

10 Thom spectra and MU

So far, our examples of spectra were the sphere spectrum $\mathbb{S}$, and Eilenberg-MacLane spectra. These essentially live at different extreme ends of the world of stable homotopy theory: Spheres have simple homology, but extremely complicated homotopy, while Eilenberg-Maclane spectra have simple homotopy and pretty complicated (but understandable) homology. The Adams spectral sequence can be thought of as a tool to leverage our full understanding on one end (the fact that we *can* describe the homology of Eilenberg-MacLane spectra as well) to gain understanding on the other end.

One of the miracles of stable homotopy theory is the existence of a spectrum MU which lands somewhere in the middle: Both its homotopy and homology groups are quite nontrivial, but fully understandable and very regular. Since it is in a sense closer to $\mathbb{S}$ than $H\mathbb{F}_p$ was, the MU-based Adams spectral sequence, i.e. the analogue of the Adams spectral sequence obtained from $\mathbb{S} \to MU$ instead $\mathbb{S} \to H\mathbb{F}_p$, behaves better than the classical Adams spectral sequence in many ways. This is called the Adams-Novikov spectral sequence, and it is one of the main tools in our later discussion of nilpotence and related phenomena.

In order to discuss MU, we will first need to discuss Thom spectra.

10.1 Thom spectra

Given a functor $B \to \mathcal{S}$, its colimit gives us the total space of the associated unstraightening $X \to B$. If we start instead with a functor $B \to \mathcal{S}_* = \mathcal{S}_{\mathrm{pt}\,/}$, the colimit (in $\mathcal{S}$) comes with a map $B \to X$ over B, i.e. a section to $X \to B$. From general descriptions of colimits in slice categories, the colimit of $B \to \mathcal{S}_*$ in $\mathcal{S}_*$ is given by X/B, i.e. the total space with its 0-section collapsed.

Definition 10.1. Let $F : B \to \mathcal{S}_*$ be a functor which takes values in the full subcategory of pointed spaces equivalent to some S^n for $n \in \mathbb{N}$. We call

$$\operatorname{colim}_B^{\mathcal{S}_*} F =: \operatorname{Th}(B, F)$$

the associated *Thom space.*

Example 10.2. For the functor $\operatorname{const}_{S^n} : B \to \mathcal{S}_*$, we get

$$\operatorname{Th}(B, \operatorname{const}_{S^n}) = \operatorname{colim}_B \operatorname{const}_{S^n} \simeq S^n \times B/B \simeq S^n \wedge B_+ \simeq \Sigma^n B_+.$$

So, we may think of Thom spaces as "twisted suspensions". We can also talk about a stable version:

Definition 10.3. Let $F : B \to \mathrm{Sp}$ be a functor which takes values in the full subcategory of spectra equivalent to $\mathbb{S}^n$'s for $n \in \mathbb{Z}$. We call

$$\operatorname{colim}_B^{\mathrm{Sp}} F =: \operatorname{th}(B, F)$$

the associated Thom spectrum.

Example 10.4. For the constant functor $\operatorname{const}_{\mathbb{S}}$, we get

$$\operatorname{th}(B, \operatorname{const}_{\mathbb{S}}) \simeq \Sigma^\infty_+ B.$$

Example 10.5. For a functor $B \to \mathcal{S}_*$ which takes values in spheres, $\Sigma^\infty \circ F : B \to \mathrm{Sp}$ takes values in spheres. Since Σ^∞ preserves colimits,

$$\operatorname{th}(B, \Sigma^\infty \circ F) \simeq \Sigma^\infty \operatorname{Th}(B, F).$$

Where do we get such functors from? We have a topologically enriched groupoid $\mathrm{Vect}_{\mathbb{R}}^{\mathrm{iso}}$ whose objects are finite-dimensional $\mathbb{R}$-vector spaces, and where the mapping space between V and W is given by $\operatorname{Iso}(V, W)$, endowed with the subspace topology from the finite-dimensional $\mathbb{R}$ vector space $\operatorname{Hom}_{\mathbb{R}}(V, W)$. Each object is isomorphic to $\mathbb{R}^n$, with automorphisms $\mathrm{GL}_n(\mathbb{R})$. In algebraic topology, it is customary to replace $\mathrm{GL}_n(\mathbb{R})$ by the homotopy equivalent space $O(n)$.

The homotopy coherent nerve of $\mathrm{Vect}_{\mathbb{R}}^{\simeq}$ is an ∞-groupoid that can be written as $\coprod_n BO(n)$, where $BO(n)$ comes from the full subgroupoid on $\mathbb{R}^n$. The notation is justified by the fact that $\operatorname{Map}(\mathbb{R}^n, \mathbb{R}^n) \simeq O(n)$, so $BO(n)$ is a pointed space with $\Omega BO(n) \simeq O(n)$.

One-point compactification assigns to each finite-dimensional real vector space V a sphere $S^V = V \cup \{\infty\}$. This defines a functor between topologically enriched categories, and hence after homotopy coherent nerves, a functor

$$\coprod_n BO(n) \to \mathcal{S}_*.$$

taking values in spheres.

Any vector bundle on a space B is classified (by a construction similar to straightening) by a functor $B \to \coprod_n BO(n)$, and so in general vector bundles give rise to sphere-valued functors $B \to \mathcal{S}_*$, and hence to Thom spaces.

The functor $\oplus\mathbb{R}$ on vector spaces induces a map $BO(n) \to BO(n+1)$ which fits into a diagram

$$\begin{array}{ccc} BO(n) & \longrightarrow & \mathcal{S}_* \\ \downarrow{\scriptstyle \oplus\mathbb{R}} & & \downarrow{\scriptstyle \wedge S^1} \\ BO(n+1) & \longrightarrow & \mathcal{S}_*. \end{array}$$

After postcomposing with Σ^∞ and shifting, this means that the composites

$$BO(n) \to \mathcal{S}_* \xrightarrow{\Sigma^{\infty-n}} \mathrm{Sp}$$

are compatible with each other. So, they factor over the colimit $BO := \mathrm{colim}_n\, BO(n)$, and we have a functor $BO \to \mathrm{Sp}$ which takes values in the full subcategory on $\mathbb{S}$. We can think of a map $B \to BO$ as a *stable vector bundle*, and so stable vector bundles give rise to Thom spectra. More generally, we can build a commutative diagram

$$\begin{array}{ccc} \coprod_n BO(n) & \longrightarrow & \mathcal{S}_* \\ \downarrow & & \downarrow{\scriptstyle \Sigma^\infty} \\ \mathbb{Z} \times BO & \longrightarrow & \mathrm{Sp} \end{array}$$

where the restriction of the bottom horizontal functor on $n \times BO$ is just built from the functor $BO \to \mathrm{Sp}$ by composing with Σ^n.

Remark 10.6. Direct sum of vector spaces endows $\coprod_n BO(n)$ with the structure of a commutative monoid in $\mathcal{S}$, or a symmetric-monoidal ∞-groupoid, and the functors to $\mathcal{S}_*$ and Sp with a symmetric-monoidal structure. In Sp, it takes values in invertible objects. It turns out there is a universal way to invert the elements of a commutative monoid in $\mathcal{S}$, "group completion", and it can be proved that this turns $\coprod_n BO(n)$ into a space equivalent to $\mathbb{Z} \times BO$. This gives $\mathbb{Z} \times BO$ (and BO) a commutative group structure, and the functor to Sp a symmetric-monoidal structure.

In particular, for two vector bundles $B \to \coprod_n BO(n)$, we can form their difference as maps $B \to \mathbb{Z} \times BO$, and so we think of $\mathbb{Z} \times BO$ classifying space for *virtual vector bundles*, formal differences of vector bundles. Stable vector bundles can be thought of as 0-dimensional virtual vector bundles. In that point of view, the stabilisation of an n-dimensional vector bundle ξ is the virtual vector bundle $\xi - \underline{R}^n$.

Theorem 10.7 (Thom isomorphism)**.** *Let $F : B \to \mathrm{Sp}$ be a functor with values in the full subcategory on $\mathbb{S}$. Let A be an abelian group, and if A is not a $\mathbb{F}_2$ vector space, assume additionally that the composite $\pi_0 \circ F : B \to \mathrm{Ab}$ is constant. Then,*

$$H_*(\mathrm{th}(B, F); A) \cong H_*(B; A).$$

Proof. The functor $\pi_0 \circ F : B \to \mathrm{Ab}$ takes all points to $\mathbb{Z}$, and is hence determined on the connected component of a point $b \in B$ by a homomorphism $\pi_1(B, b) \to \mathrm{Aut}(\mathbb{Z}) = \{\pm 1\}$. The functor $B \to \mathrm{Ab}$ corresponding to $HA \otimes F(-) : B \to \mathrm{Sp}^\heartsuit$ is therefore determined on the connected component of a point $b \in B$ by letting this homomorphism act on A, hence it is automatically constant if A is a $\mathbb{F}_2$ vector space. Otherwise, it is constant by assumption.

We obtain:

$$HA \otimes \mathrm{th}(B, F) \simeq \mathrm{colim}_B\, HA \otimes F \simeq \mathrm{colim}_B\, \mathrm{const}_{HA} \simeq HA \otimes \Sigma^\infty_+ B,$$

which on homotopy gives the desired isomorphism. □

One analogously proves a version for cohomology. If F comes from a vector bundle, the constancy condition on $\pi_0 \circ F$ is equivalent to classical orientability of vector bundles. So, while we may think of a Thom spectrum as kind of a "twisted suspension spectrum", the Thom isomorphism says that for ordinary (co)homology, the only relevant part of the twist is orientability.

10.2 MU

Repeating the constructions of $\coprod_n BO(n)$ and $\mathbb{Z} \times BO$ for complex vector bundles, we get $\coprod_n BU(n)$ and $\mathbb{Z} \times BU$, with "forgetful maps" $BU(n) \to BO(2n)$ and $BU \to BO$.

Definition 10.8. MU is the Thom spectrum associated to the composite $BU \to BO \to \mathbb{Z} \times BO \to \mathrm{Sp}$. Analogously, we define $MU(n)$ as Thom spectrum of the map $BU(n) \to BO(2n) \to \mathrm{Sp}$.

Note that we have canonical maps $\Sigma^{-2n} MU(n) \to MU$ since the map $BU(n) \to BU \to BO \to \mathrm{Sp}$ differs from the map $BU(n) \to BO(2n) \to \mathrm{Sp}$ by a shift by $2n$. The fact that $BU \simeq \mathrm{colim}\, BU(n)$ implies $MU \simeq \mathrm{colim}_n\, \Sigma^{-2n} MU(n)$.

Based on Remark 10.6, BU carries the structure of a commutative group in $\mathcal{S}$, the functor $BU \to \mathrm{Sp}$ a symmetric-monoidal structure. One has that $\mathrm{th} : \mathcal{S}_{/BO} \to \mathrm{Sp}$ is symmetric-monoidal from an argument with Day convolution and Kan extensions. In particular, one gets a commutative algebra structure on the colimit MU. Also, the "direct sum" maps $BU(n) \times BU(m) \to BU(n+m)$ induce maps $MU(n) \otimes MU(m) \to MU(n+m)$, which under the maps to MU are compatible with the algebra structure.

In order to compute its homology and homotopy, we first need some facts about the homology of $BU(n)$:

Theorem 10.9. *1. $H^*(BU(1)) \cong \mathbb{Z}[c_1]$ with c_1 in degree 2.*

2. $H^*(BU(n)) \cong \mathbb{Z}[c_1, \ldots, c_n]$, with c_i taken by the map

$$H^*(BU(n)) \to H^*(BU(1)^\times n) \cong \bigotimes_{j=1}^{n} H^*(BU(1)) \cong \mathbb{Z}[x_1, \ldots, x_n]$$

to the i-th elementary symmetric polynomial in the n degree 2 generators x_j coming from c_1 under Künneth.

3. *Writing β_n for the generator of $H_{2n}(BU(1))$ dual to c_1^n, we have that*

$$H_*(BU) \cong \mathbb{Z}[\beta_1, \beta_2, \ldots],$$

under the ring structure coming from the commutative group structure on BU (β_0 becomes the unit). More generally, $H_*(BU(n))$ identifies along the map $H_*(BU(n)) \to H_*(BU)$ with the subgroup spanned by monomials in $\leq n$ factors β_i.

Proof. We have $BU(1) \simeq \mathbb{C}P^\infty \simeq K(\mathbb{Z}, 2)$ (the latter also follows from the fact that $BU(1)$ is connected and $\Omega BU(1) \simeq U(1) \simeq S^1$), so this is just $H^*(K(\mathbb{Z}, 2)) \cong \mathbb{Z}[\iota_2]$, only that we write c_1 for ι_2.

We do not prove the second statement, but refer the reader to a course on characteristic classes, for example [17, Theorem 14.5].

For the third statement, observe that $H^*(BU(n)) \to H^*(BU(1)^{\times n})$ is a split injection (since $\mathbb{Z}[x_1, \ldots, x_n]$ is a free module over the ring of symmetric polynomials). So the dual $H_*(BU(1)^{\times n}) \to H_*(BU(n))$ is a split surjection. Different permutations of the map $BU(1)^{\times n} \to BU(n)$ are homotopic (as different permutations of sums $L_1 \oplus \cdots \oplus L_n$ of 1-dimensional vector spaces are naturally isomorphic), so the map here is symmetric and factors through a surjective map $\mathrm{Sym}^n(\mathbb{Z}\{a_0, a_1, \ldots\}) \to H_*(BU(n))$. Here, $\mathrm{Sym}^n(\mathbb{Z}\{a_0, a_1, \ldots\})$ can be thought of as the subgroup of $\mathbb{Z}[a_1, \ldots]$ spanned by monomials with $\leq$ many a_i factors (taking $a_0 \mapsto 1$). A counting argument shows that there are as many of those monomials as monomials in $\mathbb{Z}[c_1, \ldots, c_n]$, so the above surjection is an isomorphism. In the colimit for $n \to \infty$, this proves the last statement. □

Remark 10.10. Similar to the situation with the Steenrod algebra and its dual, the pairing between the monomial bases in the β_i on $H_*(BU)$ and the c_i on $H^*(BU)$ is very nontrivial.

Under the Thom isomorphism (using the fact that the algebra structure on MU comes from the group structure of BU under th), we get

Corollary 10.11. *Writing b_i for the element of $H_*(\Sigma^{-2}MU(1))$ corresponding to $\beta_i \in H_*(BU(1))$ under the Thom isomorphism, we have that*

$$H_*(MU) = \mathbb{Z}[b_1, b_2, \ldots],$$

and $H_*(\Sigma^{-2n}MU(n))$ identifies with the subgroup spanned by monomials of $\leq n$ factors b_i.

We next want to identify the structure of $H_*(MU; \mathbb{F}_p)$ as $\mathcal{A}_*$-comodule. The Thom isomorphism does not preserve the Steenrod algebra action. While one can generaly describe how the action gets twisted under the Thom isomorphism (in terms of characteristic classes), here there is an easier solution:

Lemma 10.12. *There is an equivalence $MU(1) \simeq \Sigma^\infty BU(1)$.*

Proof. Note that this is a highly nonformal statement: The description of $MU(1) = \mathrm{th}(BU(1))$ expresses $MU(1)$ as a twisted version of $\Sigma^{\infty+2}_{+} BU(1)$, while the statement here describes $MU(1)$ instead as untwisted, reduced suspension of $BU(1)$.

To see this, we use a geometric argument. We have the functor $S(-) : BU(1) \to \mathcal{S}$ that takes a 1-dimensional complex vector space to $V \setminus \{0\}$ (which is homotopy equivalent to an S^1). We may describe the functor $S^- : BU(1) \to \mathcal{S}$ that takes V instead to the 2-sphere S^V as pushout

$$S^V \simeq \mathrm{pt} \amalg_{S(V)} \mathrm{pt}$$

(arising from the pushout of $S^V \setminus \{0\}$ and $S^V \setminus \{\infty\}$ over their intersection). Taking colimits in unpointed spaces, we get a pushout diagram

$$\begin{array}{ccc} \mathrm{colim}^{\mathcal{S}}_{BU(1)} S(V) & \longrightarrow & BU(1) \\ \downarrow & & \downarrow \\ BU(1) & \longrightarrow & \mathrm{colim}^{\mathcal{S}}_{BU(1)} S^V. \end{array}$$

As the colimit in pointed spaces is instead obtained by collapsing the section given by the points $\infty \in S^V$, it sits in a pushout diagram

$$\begin{array}{ccc} \mathrm{colim}^{\mathcal{S}}_{BU(1)} S(V) & \longrightarrow & BU(1) \\ \downarrow & & \downarrow \\ \mathrm{pt} & \longrightarrow & \mathrm{colim}^{\mathcal{S}_*}_{BU(1)} S^V. \end{array}$$

Now, crucially, $\mathrm{colim}^{\mathcal{S}}_{BU(1)} S(V)$ is contractible. Indeed, straightening the map $\mathrm{pt} \to BU(1)$ given by the basepoint, we get a functor $BU(1) \to \mathcal{S}$ which takes every point to a space homotopy equivalent to $U(1) \simeq S^1$, and this turns out to be $S(V)$ (morally, because a point in the fiber over some V is an element of the pullback $\{V\} \times_{BU(1)} \{\mathbb{C}\}$, since the basepoint comes from the 1-dimensional vector space $\mathbb{C}$. But this is equivalent to the mapping space $\mathrm{Iso}(\mathbb{C}, V) \simeq S(V)$. To make this fully precise, one has to be more explicit about the construction of $BU(1)$ from the topological groupoid of 1-dimensional vector spaces than we are here). So, the colimit recovers pt. It follows that the right hand map in the above pushout gives an equivalence between $BU(1)$ and $\mathrm{Th}(BU(1), S^V)$ as pointed spaces. As

$$\mathrm{th}(BU(1)) = \Sigma^\infty \, \mathrm{Th}(BU(1), S^V),$$

the claim follows. □

This gives another identification between $H_*(\Sigma^{-2} MU(1))$ and $H_*(BU(1))$, this time with $H_n(\Sigma^{-2} MU(1)) \cong \widetilde{H}_{n+2}(BU(1))$, with a shift by 2. Here

$b_i \in H_*(\Sigma^{-2}MU(1))$ corresponds to $\beta_{i+1} \in H_*(BU(1))$. (To see that the coefficient is actually 1 here, one writes the composite isomorphism $\widetilde{H}_{*+2}(BU(1)) \to H_*(BU(1))$ as dual to multiplication with c_1, by relating the Thom isomorphism to characteristic classes.)

In what follows, we will write $\mathcal{P}_* \subseteq \mathcal{A}_*$ for the subalgebra generated by the ξ_i (resp. the ζ_i^2 at $p=2$). This is a sub Hopf algebra, since the comultiplication on those generators takes values in $\mathcal{P}_* \otimes \mathcal{P}_*$. In particular, we can view $\mathcal{P}_*$ as a left $\mathcal{A}_*$ comodule (by viewing the comultiplication with values in $\mathcal{A}_* \otimes \mathcal{P}_*$.)

Theorem 10.13. *$H_*(MU;\mathbb{F}_p)$, as a left $\mathcal{A}_*$ comodule, is isomorphic to*

$$\mathcal{P}_* \otimes \mathbb{F}_p[b_k \mid k+1 \textit{ not a power of } p]$$

Proof. We do the proof for odd p, but the $p=2$ case follows analogously by replacing ξ_i by ζ_i^2, compare Remark 7.12.

Recall that $\lambda_*(\iota_2) = \sum \iota_2^{p^i} \otimes \xi_i$ in skeleta of $K(\mathbb{Z},2) \simeq BU(1)$. All these values lie in $\mathcal{P}_*$, and so the same holds for $\lambda_*(\iota_2^k)$. Dualizing, we learn that the coaction on $H_*(BU(1);\mathbb{F}_p)$ takes values in $\mathcal{P}_* \otimes H_*(BU(1);\mathbb{F}_p)$, and furthermore in $\psi(\beta_k)$ the coefficient on β_1 is ξ_i if $k=p^i$, and 0 otherwise. Under the identification with $H_*(\Sigma^{-2}MU(1))$ coming from $\Sigma^{-2}MU(1) \simeq \Sigma^{\infty-2}BU(1)$, β_i turns into b_{i-1}, and so the coefficient of b_0 in $\psi(b_k)$ is ξ_i if $k+1=p^i$, and 0 otherwise. (This isomorphism didn't come from a Thom isomorphism, just from a suspension isomorphism, so it is compatible with the coaction). Now in $H_*(MU;\mathbb{F}_p)$, this means that the coaction map $H_*(MU;\mathbb{F}_p) \to \mathcal{P}_* \otimes H_*(MU;\mathbb{F}_p)$ satisfies

$$\psi(b_k) = \begin{cases} \xi_i \otimes 1 + 1 \otimes b_k + \dots & \text{if } k+1 = p^i \\ 1 \otimes b_k + \dots & \text{otherwise.} \end{cases}$$

Here, the remainder terms are *decomposables*: They are in I^2 if I denotes the ideal of positive-degree elements of $\mathcal{P}_* \otimes H_*(MU;\mathbb{F}_p)$. We can postcompose the coaction map with the quotient map taking all $b_{p^i-1} \mapsto 0$ to obtain a comodule map

$$H_*(MU_*;\mathbb{F}_p) \to \mathcal{P}_* \otimes \mathbb{F}_p[b_k \mid k+1 \text{ not a power of } p]$$

(where the comodule structure in the target comes exclusively from the left factor). Indeed, both sides are polynomial rings on one generator of each even degree (with ξ_i replacing b_{p^i-1} on the right), and the map takes b_k to $\zeta_i \otimes 1$ plus decomposable elements if $k+1=p^i$, and to $1 \otimes b_k$ plus decomposable elements otherwise. Now it is an easy inductive argument that a map between graded polynomial rings on positive-degree generators, which modulo decomposables takes generators to generators (bijectively), is an isomorphism. ☐

As a consequence, we have a splitting

$$\mathrm{Ext}^{*,*}_{\mathrm{Comod}_{\mathcal{A}_*}}(\mathbb{F}_p, H_*(MU;\mathbb{F}_p)) \cong \mathrm{Ext}^{*,*}_{\mathrm{Comod}_{\mathcal{A}_*}}(\mathbb{F}_p, \mathcal{P}_*)[b_k \mid k+1 \text{ not a power of } p]$$

We note that $H_*(MU;\mathbb{F}_p)$ has the structure of an algebra in $\mathrm{Comod}_{\mathcal{A}_*}$, and so does $\mathcal{P}_*$ (as sub Hopf algebra of $\mathcal{A}_*$). Furthermore, the maps above are all ring maps, and so this is in fact an isomorphism compatible with the graded ring structure on $\mathrm{Ext}^{*,*}_{\mathrm{Comod}_{\mathcal{A}_*}}$. Next, we identify $\mathrm{Ext}^{*,*}_{\mathrm{Comod}_{\mathcal{A}_*}}(\mathbb{F}_p, \mathcal{P}_*)$.

Recall that $\mathcal{A}_*$ is a Hopf algebra, and hence has an antipode $\mathcal{A}_* \to \mathcal{A}_*$ which is an involution which flips the comultiplication. In particular, one has

$$\Delta(\overline{\xi}_n) = \sum \overline{\xi}_i \otimes \overline{\xi}_j^{p^i}$$
$$\Delta(\overline{\tau}_n) = 1 \otimes \overline{\tau}_n + \sum \overline{\tau}_i \otimes \overline{\xi}_j^{p^i}$$

and similarly at $p = 2$. One of the identities of the antipode, comultiplication and multiplication implies

$$\sum_{i+j=n} \xi_i^{p^j} \overline{\xi}_j = 0$$

for $n \geq 1$, and from this one inductively checks that $\mathcal{P}_*$ can also be described as subring generated by the $\overline{\xi}_n$ instead.

Definition 10.14. For $J \subseteq \mathbb{Z}_{\geq 0}$, we let C_J denote the subalgebra of $\mathcal{A}_*$ generated by all $\overline{\xi}_n$ and all $\overline{\tau}_j$ with $j \in J$. For $p = 2$, we analogously write C_J for the subalgebra generated by all $\overline{\zeta}_n^2$ and all $\overline{\zeta}_{j+1}$ for $j \in J$.

One observes directly that C_J is a left sub-$\mathcal{A}_*$-comodule of $\mathcal{A}_*$, since the comultiplication takes each of the generators into $\mathcal{A}_* \otimes C_J$.

Now, for any J and any $n \notin J$, we have a short exact sequence of comodules

$$0 \to C_J \to C_{J \cup \{n\}} \to C_J(2p^n - 1) \to 0.$$

Rotating, this defines a map

$$a_n : \Sigma^{-1} C_J(2p^n - 1) \to C_J$$

in $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$. In particular, for $J = 0$, this is a derived comodule map $\Sigma^{-1}\mathcal{P}_* \to \mathcal{P}_*$. All C_J are comodule algebras, and the extension used to define $h_{n+1,0}$ is one of C_J-modules. By restricting along $\Sigma^{-1}\mathbb{F}_p(2p^n - 1) \to \Sigma^{-1} C_{J(2p^n-1)}$, the map a_n corresponds to an element in $\mathrm{Ext}^{1,2p^j-1}_{\mathcal{A}_*}(\mathbb{F}_p, C_J)$, and the map a_n above is obtained by acting with it (using the algebra structure of C_J). Furthermore, $J = \emptyset$ yields $C_\emptyset = \mathcal{P}_*$, and all other C_J are algebras under $\mathcal{P}_*$. All $h_{n+1,0}$ can therefore be viewed as elements in $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, \mathcal{P}_*)$, which acts on all $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, C_J)$, and the sequences above show that $C_{J\cup\{n\}}$ is the cofiber of a_n on C_J.

We note that a_0 even lifts to an element of $\mathrm{Ext}^{1,1}_{\mathcal{A}_*}(\mathbb{F}_p, \mathbb{F}_p)$, which we met earlier under the name h_0. This was what detected p in $\pi_0(\mathbb{S}_p^\wedge)$. One way to see this is by observing that $C_\{n \mid n \geq 1\}$ is equal to our earlier description of $H_*(H\mathbb{Z};\mathbb{F}_p)$, compare Remark 9.2, and the maps induced on $\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, -)$ by

$$\mathbb{F}_p \to \mathcal{P}_* \to C_{\{n \mid n \geq 1\}}$$

are all isomorphisms on the $t - s = 0$ part.

Proposition 10.15.

$$\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, C_J) \cong \mathbb{F}_p[a_n \mid n \notin J].$$

In particular,

$$\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, \mathcal{P}_*) \cong \mathbb{F}_p[a_n \mid n \geq 0].$$

Proof. We first prove this for cofinite J, by an induction on the size of $\mathbb{Z}_{\geq 0} \setminus J$. If $J = \mathbb{Z}_{\geq 0}$, then $C_J = \mathcal{A}_*$, and

$$\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, \mathcal{A}_*) \cong \mathbb{F}_p$$

as claimed. Inductively, assume the claim is known for $C_{J\cup\{n\}}$. We have a cofiber sequence

$$\Sigma^{-1} C_J(2p^n - 1) \to C_J \to C_{J\cup\{n\}}.$$

As $\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, C_{J\cup\{n\}})$ is generated by the a_m for $m \notin J \cup \{n\}$, all of which lift to C_J, and the right-hand map is an algebra map, it induces a surjective map on Ext. So, the long exact sequence on Ext degenerates to a short exact sequence

$$0 \to \mathrm{Ext}^{*-1,*-(2p^n-1)}_{\mathcal{A}_*}(\mathbb{F}_p, C_J) \xrightarrow{a_n} \mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, C_J) \to \mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, C_{J\cup\{n\}}) \to 0$$

which implies the claim.

Finally, assume J is not necessarily cofinite, and let $J' = J \cup \{n \mid n \geq N\}$. Then J' is cofinite, and the map $C_J \to C_{J'}$ is an isomorphism in formal degrees $< 2p^N - 1$. This also shows that

$$\mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, C_J) \to \mathrm{Ext}^{s,t}_{\mathcal{A}_*}(\mathbb{F}_p, C_{J'})$$

is an isomorphism for $t < 2p^N - 1$ (in fact, $t - s < 2p^N - 1$, using minimal resolutions). In that range, this implies the claim. Since we can choose N arbitrarily large (and J' accordingly), the claim follows in general. □

Remark 10.16. The usual way this computation is done is by expressing $\mathcal{P}_*$ as image of the right adjoint to a "forgetful functor" $\mathrm{Comod}_{\mathcal{A}_*} \to \mathrm{Comod}_{E_*}$ for E_* a certain quotient Hopf algebra of $\mathcal{A}_*$, and saying some words about how this adjunction derives. We will return to this viewpoint later.

Theorem 10.17. *We have that* $\pi_*(MU) = \mathbb{Z}[x_1, x_2, \ldots]$ *is a polynomial algebra on generators of degree* $|x_n| = 2n$.

Proof. We first prove this after p-completion. To compute $\pi_*(MU^\wedge_p)$, we can use an Adams spectral sequence whose E_2 page is given by

$$\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_p, \mathcal{P}_*) \otimes \mathbb{F}_p[b_k \mid k+1 \text{ not a power of } p].$$

This is itself polynomial, in the generators b_k of bidegree $(t-s, s) = (2k, 0)$, and the generators a_n of bidegree $(t-s, s) = (2p^n - 2, 1)$. Notably, everything lives in degrees where $t - s$ is even, and so the spectral sequence degenerates. Then, for $k \geq 1$, we may choose $x_k \in \pi_{2k}(MU^\wedge_p)$ detected by b_k if $k+1$ is not a power of p, and by a_i if $k+1=p^i$. Also, p is detected by a_0. Together, these define a map of graded rings

$$\mathbb{Z}_p^\wedge[x_1, x_2, \dots] \to \pi_*(MU_p^\wedge).$$

If we endow the first ring with a multiplicative filtration where p is in filtration 1, x_k is in filtration 1 if $k+1$ is a power of p, and 0 otherwise (extended to all monomials by adding up the individual filtrations), this filtration is compatible with the abument filtration. The associated graded is a polynomial algebra over $\mathbb{F}_p$ on the elements represented by p and all the x_i, and so the map becomes an isomorphism on associated gradeds. From completeness of the filtrations, it follows that the map is an isomorphism.

To obtain an integral statement, observe that because MU is connective and the homology groups $H_*(MU)$ are finitely generated in each degree, also the $\pi_*(MU)$ are finitely generated. Their p-completion is then just given by tensoring with $\mathbb{Z}_p$. They also have to be p-torsion free, since if there was a $\mathbb{Z}/p^k$ summand, it would also be present after p-completion. Similarly, they have to be free of the same rank as we see after p-completion, in particular even. Suppose I denotes the ideal of positive-degree elements in $\pi_*(MU)$, and look at the short exact sequence

$$0 \to (I^2)_{2n} \to I_{2n} \to (I/I^2)_{2n} \to 0.$$

We know that after tensoring with the (flat!) $\mathbb{Z}_p$, the right-hand term is $\mathbb{Z}_p$, generated by x_n. But since it base-changes to $\mathbb{Z}_p$ at each p, it must also be free of rank 1. Letting $x_n \in \pi_{2n}(MU)$ denote any representative of a generator of $(I/I^2)_{2n}$, the claim follows, since these x_i define a map

$$\mathbb{Z}[x_1, x_2, \dots] \to \pi_*(MU)$$

which becomes an isomorphism after base-change to $\mathbb{Z}_p^\wedge$ for each p, and both sides are degreewise finitely-generated. □

So, both homology and homotopy of MU are simply polynomial rings of the same shape! We can say a bit more. Consider the sequential diagram

$$\mathbb{S} \xrightarrow{2} \mathbb{S} \xrightarrow{3} \mathbb{S} \xrightarrow{4} \dots$$

where the maps are multiplication with all positive integers. Since homotopy groups commute with filtered colimits, the homotopy groups of its colimit are obtained by an analogous colimit of $\pi_*(\mathbb{S})$. Since $\pi_*(\mathbb{S})$ is torsion in positive degrees, the colimit is equivalent to $H\mathbb{Q}$. Commuting this colimit with $-\otimes X$, we also learn that

$$\mathbb{Q} \otimes_{\mathbb{S}} X \simeq \operatorname{colim}(X \xrightarrow{2} X \xrightarrow{3} X \xrightarrow{4} \dots),$$

and hence that $\pi_*(\mathbb{Q} \otimes_{\mathbb{S}} X) \simeq \mathbb{Q} \otimes \pi_*(X)$. Looking at homotopy groups, the canonical map $\mathbb{Q} \otimes_{\mathbb{S}} \mathbb{Z} \to \mathbb{Q}$ is an equivalence. In particular,

$$X \to \mathbb{Z} \otimes_{\mathbb{S}} X$$

becomes an equivalence after tensoring with $\mathbb{Q}$, and so $\mathbb{Q} \otimes \pi_*(X) \to \mathbb{Q} \otimes H_*(X)$ is an isomorphism.

In the case of MU, we also know that $\pi_*(MU)$ and $H_*(MU)$ are torsion free, hence map injectively into their rationalisations. It follows that the map

$$\pi_*(MU) \to H_*(MU)$$

is injective. It is of course not an isomorphism. For example, a generator $x_1 \in \pi_2(MU)$ maps to $\pm 2b_1$ in $H_2(MU)$, here the 2 is related to $\pi_1(\mathbb{S}) \cong \mathbb{Z}/2$ through the Atiyah-Hirzebruch spectral sequence $H_*(MU; \pi_*(\mathbb{S})) \Rightarrow \pi_*(MU)$.

A more elaborate comparison of the Adams spectral sequences for MU and $H\mathbb{Z} \otimes MU$ can be used to show

$$x_i \mapsto \begin{cases} \pm p b_i + \dots & \text{if } i+1 \text{ is a power of } p \\ \pm b_i + \dots & \text{otherwise.} \end{cases}$$

Here, the remainder terms are again decomposables. This at least determines the index of $\pi_*(MU)$ in $H_*(MU)$. Using this, Quillen was able to determine the subring $\pi_*(MU) \subseteq H_*(MU)$ explicitly:

Theorem 10.18 (Quillen). *Let $f \in H_*(MU)[\![x]\!]$ denote the power series*

$$f(x) = x + b_1 x^2 + b_2 x^3 + \dots.$$

Then $\pi_*(MU) \subseteq H_*(MU)$ is the subring generated by the coefficients of the power series in two variables given by

$$f^{-1}(f(x) + f(y)),$$

where f^{-1} denotes the unique power series with vanishing constant term and $f^{-1}(f(x)) = f(f^{-1}(x)) = x$.

The proof proceeds by showing that $E^*(\mathbb{C}P^\infty)$ is a power series ring over $\pi_*(E)$ for E any of MU, $H\mathbb{Z}$ or $H\mathbb{Z} \otimes_{\mathbb{S}} MU$, and that f appears when one expresses the two generators of $(H\mathbb{Z} \otimes_{\mathbb{S}} MU)^*(\mathbb{C}P^\infty)$ coming from MU and $H\mathbb{Z}$ in terms of each other. The power series $f^{-1}(f(x)+f(y)$ can then be interpreted in terms of the map $MU^*(\mathbb{C}P^\infty) \to MU^*(\mathbb{C}P^\infty \times \mathbb{C}P^\infty)$ induced by the map $\mathbb{C}P^\infty \times \mathbb{C}P^\infty \to \mathbb{C}P^\infty$ which corresponds to addition of cohomology classes under $K(\mathbb{Z}, 2) \simeq \mathbb{C}P^\infty$. This shows that $\pi_*(MU)$ contains the coefficients of $f^{-1}(f(x) + f(y))$, and the proof is then finished by proving (algebraically) that both of these subrings of $H_*(MU)$ have the same index. (In fact, a lower bound for the index of $\pi_*(MU) \subseteq H_*(MU)$ suffices, which is somewhat easier to see.)

For example, as $f^{-1}(x) = x - b_1 x^2 + \dots$, we have

$$f^{-1}(f(x) + f(y)) = x + y - 2b_1 xy + \dots,$$

and so the image of $\pi_2(MU) \to H_2(MU)$ is spanned by $2b_1$. This description of $\pi_*(MU)$, and a related description of $\pi_*(MU \otimes \ldots \otimes MU) \subseteq H_*(MU \otimes \ldots \otimes MU)$, has a much more illuminating interpretation in terms of the theory of *formal groups*. We will skip this for now, but refer the reader for example to [2] or [12] for a more in-depth discussion.

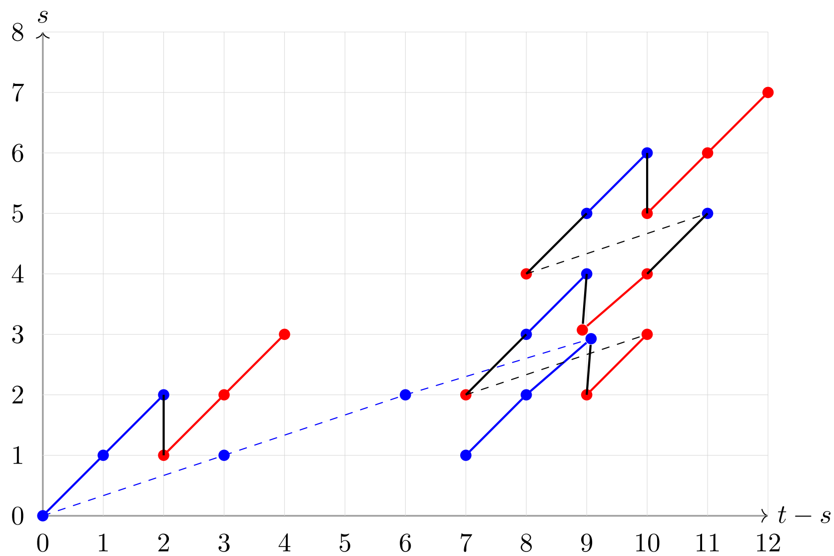

Figure 2: $\mathbb{F}_2$-based Adams spectral sequence for $\mathbb{S}/2$

11 Nilpotence and periodicity: First observations

11.1 The Adams spectral sequence for $\mathbb{S}/2$

We can consider the Adams spectral sequence for $\mathbb{S}/2$, the cofiber of the degree 2 map $\mathbb{S} \to \mathbb{S}$. Here, $H_*(\mathbb{S}/2; \mathbb{F}_2)$ can be identified (along the map $\mathbb{S}/2 \to \mathbb{F}_2$) with the sub left $\mathcal{A}_*$ comodule of $\mathcal{A}_*$ given by the linear span of $1, \zeta_1$. Note that this is not a subalgebra, in particular *there cannot be an algebra structure on* $\mathbb{S}/2$. In $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$, this comodule is in fact exactly the cofiber of $h_0 : \Sigma^{-1}\mathbb{F}_2(1) \to \mathbb{F}_2$, since h_0 can be described in terms of the extension

$$0 \to \mathbb{F}_2 \to \langle 1, \zeta_1 \rangle \to \mathbb{F}_2(1) \to 0.$$

In particular, $\mathrm{Ext}^{*;*}_{\mathcal{A}_*}(\mathbb{F}_2, H_*(\mathbb{S}/2; \mathbb{F}_2))$ sits in a long exact sequence with the multiplication by h_0 map on $\mathrm{Ext}^{*;*}_{\mathcal{A}_*}(\mathbb{F}_2, \mathbb{F}_2)$. Since in graded $\mathbb{F}_2$ vector spaces, every short exact sequence splits, this exhibits the Ext in question as split extension of $\mathrm{coker}(h_0)$ and a shift of $\ker(h_0)$, seen as the blue and red part in Figure 2. Of course, this does not determine the full structure as $\mathrm{Ext}^{*;*}_{\mathcal{A}_*}(\mathbb{F}_2, \mathbb{F}_2)$-module, this additional information (multiplication with h_0, h_1 and h_2) is indicated by the additional black lines. These are easiest determined by computing Ext with minimal resolutions.

Again, in the present range all possible differentials can be ruled out purely by multiplicativity of the spectral sequence (this time, with respect to the

$\mathrm{Ext}_{\mathcal{A}_*}(\mathbb{F}_2,\mathbb{F}_2)$-module structure). Curiously, we see that $\pi_2(\mathbb{S}/2) \cong \mathbb{Z}/4$ (which also implies that $\mathbb{S}/2$ cannot have the structure of a ring spectrum, as $\pi_0(\mathbb{S}/2) \cong \mathbb{Z}/2$).

We also see some interesting repetition of patterns, in the form of different "lightning flashes". Focusing on the top right one, let $x \in \pi_8(\mathbb{S}/2)$ be the unique element detected by the nonzero class in degree $(8,4)$. Since it is 2-torsion, the composite

$$\mathbb{S}^8 \xrightarrow{2} \mathbb{S}^8 \xrightarrow{x} \mathbb{S}/2$$

is nullhomotopic, and x extends to a self-map $\theta : \Sigma^8\mathbb{S}/2 \to \mathbb{S}/2$. In fact, the map

$$\mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_2, H_*(\mathbb{S}/2;\mathbb{F}_2)) \to \mathrm{Ext}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_2, H_*(MU\otimes\mathbb{S}/2;\mathbb{F}_2))$$

can be seen to take this class to $h_{2,0}^4$, and from this that the image of x in $\pi_*(MU\otimes\mathbb{S}/2)$ is x_1^4. It follows from this that θ acts on $\pi_*(MU\otimes\mathbb{S}/2)$ by multiplication with x_1^4.

In particular, MU detects that θ cannot be nilpotent! For example, θ^k takes $1 \in MU_*(\mathbb{S}/2) \cong \pi_*(MU)/2$ to $x_1^{4k} \in MU_*(\mathbb{S}/2) \cong \pi_*(MU)/2$.

Adams observes that this self-map in fact exhibits a certain 8-periodic pattern along a line of slope $\frac{1}{2}$ in the Adams spectral sequence for $\mathbb{S}/2$, and a related (but a little bit harder to describe) 8-periodicity phenomenon in the Adams spectral sequence for $\mathbb{S}_2^\wedge$ (referred to as "Adams periodicity"). In fact, if we write $\mathbb{F}_2/h_0$ for the comodule $H_*(\mathbb{S}/2;\mathbb{F}_2)$, then since the class $\Sigma^{-4}\mathbb{F}_2(12) \to \mathbb{F}_2/h_0$ representing the nonzero element in bidegree $(8,4)$ is h_0-torsion, we can (as algebraic analogue of the construction of θ) form an extension

$$\theta_{\mathrm{alg}} : \Sigma^{-4}\mathbb{F}_2/h_0(12) \to \mathbb{F}_2/h_0,$$

which is a $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ version of the self-map θ. Now Adams proves that the cofiber of that map in $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ has $\mathrm{map}(\mathbb{F}_2,-)$ ("$\mathrm{Ext}_{\mathcal{A}_*}$", but of course this cofiber is not an ordinary comodule anymore) concentrated below a line of slope strictly smaller than $\frac{1}{2}$. (Adams proves $\frac{1}{3}$, but it turns out the optimal bound is $\frac{1}{5}$.) In other words, θ_{alg} induces an isomorphism

$$\mathrm{Ext}^{s,t}(\mathbb{F}_2,\mathbb{F}_2/h_0) \to \mathrm{Ext}^{s+4,t+12}(\mathbb{F}_2,\mathbb{F}_2/h_0)$$

for $s \geq \frac{1}{5}(t-s) + C$ for some constant C.

This can be used to understand a big part of the Adams spectral sequence of $\mathbb{S}/2$ (and less directly, $\mathbb{S}_2^\wedge$). An analogous (but in fact easier) periodicity self-map exists on $\mathbb{S}/p$ as well. We can ask if this generalizes:

Question 11.1. 1. Are there more non-nilpotent self-maps $\Sigma^k X \to X$ for other spectra X?

2. Are they always detected by MU_*?

As it turns out, these questions have very satisfactory (positive) answers, in Hopkins-Smith's Periodicity Theorem and Devinatz-Hopkins-Smith's Nilpotence Theorem. We will state (and prove) those later.

Remark 11.2. In Remark 8.12 we pointed out that there exist nonconnective spectra X with $X/p \neq 0$ and $\mathbb{F}_p \otimes X = 0$. Indeed,

$$X = \operatorname{colim}(\mathbb{S}/2 \xrightarrow{\theta} \Sigma^{-8}\mathbb{S}/2 \xrightarrow{\theta} \Sigma^{-16}\mathbb{S}/2 \xrightarrow{\theta} \ldots)$$

is one of those: The map θ is clearly zero on $\mathbb{F}_2$-homology, and so the colimit here has vanishing $\mathbb{F}_2$-homology. Also, X is nonzero, because its MU-homology is $\pi_*(MU)/2[x_1^{-1}]$. The MU homology of $X/2$ is similarly nonzero (since 2 acts by 0 on $\pi_*(MU)/2[x_1^{-1}]$).

11.2 Nishida's nilpotence theorem

After constructing the non-nilpotent self-map of $\mathbb{S}/2$, another immediate question is why we did have to look at $\mathbb{S}/2$ at all: Are there maybe already non-nilpotent self-maps of $\mathbb{S}$?

In this section, we prove the following:

Theorem 11.3 (Nishida)**.** *Let $\alpha \in \pi_n(\mathbb{S})$ with $n > 0$. Then $\alpha^k = 0$ for some k.*

The proof is based on analyzing power operations in homotopy groups. Given $\alpha : \mathbb{S}^n \to \mathbb{S}$, its k-th power can be obtained as composition

$$(\mathbb{S}^n)^{\otimes k} \xrightarrow{\alpha^{\otimes k}} \mathbb{S}^{\otimes k} \xrightarrow{\simeq} \mathbb{S}$$

where we use the functor $(-)^{\otimes k}$ followed by a canonical equivalence arising from the fact that $\mathbb{S}$ is the monoidal unit.

One can enhance this by so-called *extended powers*. Let us write BG for (the nerve of) the category with one object with endomorphisms given by a group G, and think of a functor $BG \to \mathrm{Sp}$ as a "spectrum with G-action". If $X : BG \to \mathrm{Sp}$ is such a functor, we write $X_{hG} := \operatorname{colim}_{BG} X$, and $X^{hG} := \lim_{BG} X$. Part of symmetric-monoidal structures is that for $X \in \mathrm{Sp}$, $X^{\otimes k}$ has a canonical Σ_k-action, so refines to a functor $B\Sigma_k \to \mathrm{Sp}$. Writing $D_k(X) = (X^{\otimes k})_{h\Sigma_k}$, the above diagram fits into

$$\begin{array}{ccccc} (\mathbb{S}^n)^{\otimes k} & \xrightarrow{\alpha^{\otimes k}} & \mathbb{S}^{\otimes k} & \xrightarrow{\simeq} & \mathbb{S} \\ \downarrow & & \downarrow & \nearrow_{\mu_k} & \\ D_k(\mathbb{S}^n) & \xrightarrow{D_k(\alpha)} & D_k(\mathbb{S}), & & \end{array}$$

where the μ_k map exists because the Σ_k action on $\mathbb{S}^{\otimes k}$ is trivial (a stronger way of saying this is that $\mathbb{S}$ has a canonical E_∞ ring structure since it is the unit of the symmetric-monoidal structure on Sp).

So α^k extends over a map $D_k(\mathbb{S}^n) \to \mathbb{S}$. More generally, this works for $D_k^G(X) = (X^{\otimes k})_{hG}$ for G any group with homomorphism $G \to \Sigma_k$. We will refer to the resulting extensions $D_k^G(\mathbb{S}^n) \to \mathbb{S}$ as $\widetilde{\alpha^k}$.

An element β of $\pi_m(D_k^G(\mathbb{S}^n))$ allows us to form a composite

$$\mathbb{S}^m \xrightarrow{\beta} D_k^G(\mathbb{S}^n) \xrightarrow{\widetilde{\alpha^k}} \mathbb{S},$$

producing from $\alpha \in \pi_n(\mathbb{S})$ a new element of $\pi_m(\mathbb{S})$ we write $P_\beta(\alpha)$. P_β is an example of so-called *power operations*.

The nilpotence theorem is based on the fact that P_β is not quite additive, but we can express its failure to being additive explicitly in terms of *transfers*.

Definition 11.4. For a finite group G, and $X \in \mathrm{Fun}(BG, \mathrm{Sp})$, the composition $\mathrm{Fun}(BG, \mathrm{Sp}) \to \mathrm{Sp} \to \mathrm{Fun}(BG, \mathrm{Sp})$ of the forgetful functor followed by its right adjoint takes X to $\prod_{g \in G} X$ with the action permuting the factors. The canonical map $X \to \prod_{g \in G} X$ induces on homotopy orbits a map

$$X_{hG} \to X,$$

which is a natural transformation of functors $\mathrm{Fun}(BG, \mathrm{Sp}) \to \mathrm{Sp}$. We call this the *transfer* tr_G.

Lemma 11.5 (Nishida). *Let $\beta \in \pi_m(((\mathbb{S}^n)^{\otimes p})_{hC_p})$ and $\alpha \in \pi_n(\mathbb{S})$. Then*

$$P_\beta(k\alpha) = kP_\beta(\alpha) + \frac{k^p - k}{p}\,\mathrm{tr}_{C_p}(\beta)\alpha^p$$

for any $k \geq 0$.

Proof. Since we can write $k\alpha$ as composite

$$\mathbb{S}^n \xrightarrow{\mathrm{diag}} (\mathbb{S}^n)^{\bigoplus k} \xrightarrow{\mathrm{codiag}} \mathbb{S}^n \xrightarrow{\alpha} \mathbb{S},$$

we can write $\widetilde{(k\alpha)^p} : D_p^{C_p}(\mathbb{S}^n) \to \mathbb{S}$ as composite

$$D_p^{C_p}(\mathbb{S}^n) \xrightarrow{D_p^{C_p}(\mathrm{diag})} D_p^{C_p}((\mathbb{S}^n)^{\bigoplus k}) \xrightarrow{D_p^{C_p}(\mathrm{codiag})} D_p^{C_p}(\mathbb{S}^n) \xrightarrow{\widetilde{\alpha^p}} \mathbb{S},$$

and hence $P_\beta(k\alpha)$ as a composite

$$\mathbb{S}^m \xrightarrow{\beta} D_p^{C_p}(\mathbb{S}^n) \xrightarrow{D_p^{C_p}(\mathrm{diag})} D_p^{C_p}((\mathbb{S}^n)^{\bigoplus k}) \xrightarrow{D_p^{C_p}(\mathrm{codiag})} D_p^{C_p}(\mathbb{S}^n) \xrightarrow{\widetilde{\alpha^p}} \mathbb{S}.$$

We analyze the middle part. We have

$$((\mathbb{S}^n)^{\bigoplus k})^{\otimes p} \simeq \bigoplus_{k^p} \mathbb{S}^{pn},$$

with sum indexed by the functions $p \to k$ (writing p for the p-element set $\{1, \ldots, p\}$, and k analogously). This set k^p splits C_p-equivariantly into k trivial orbits (the constant functions $p \to k$) and $\frac{k^p - k}{p}$ free orbits: Since C_p has only the trivial and the full subgroup, every orbit is either trivial or free. The trivial orbits are clearly the constant maps, and so there are $k^p - k$ elements of free orbits, hence $\frac{k^p-k}{p}$ such orbits.

Compatibly with the C_p-action, $((\mathbb{S}^n)^{\bigoplus k})^{\otimes p}$ hence splits into k summands of the form $(\mathbb{S}^n)^{\otimes p}$ with action permuting the tensor factors, and $\frac{k^p-k}{p}$ summands of the form $\prod_{g\in C_p} \mathbb{S}^{pn}$, with action permuting the factors of the product. After homotopy orbits, this means that $D_p^{C_p}((\mathbb{S}^n)^{\bigoplus k})$ splits into p summands of the form $D_p^{C_p}(\mathbb{S}^n)$, and $\frac{k^p-k}{p}$ summands of the form $\mathbb{S}^{pn}$. The map $D_p^{C_p}(\text{diag})$ is given by the identity into each $D_p^{C_p}(\mathbb{S}^n)$-summand, and the transfer tr_{C_p} : $D_p^{C_p}(\mathbb{S}^n) = ((\mathbb{S}^n)^{\otimes p})_{hC_p} \to \mathbb{S}^{pn}$ into each $\mathbb{S}^{pn}$ summand. The map $D_p^{C_p}(\text{codiag})$ is given by the identity on each $D_p^{C_p}(\mathbb{S}^n)$-summand, and the canonical map $\iota : \mathbb{S}^{pn} \to D_p^{C_p}(\mathbb{S}^n)$ on each $\mathbb{S}^{pn}$-summand. We hence obtain that the composite $D_p^{C_p}(\text{codiag}) \circ D_p^{C_p}(\text{diag})$ is given by $k \cdot \mathrm{id}_{D_p^{C_p}(\mathbb{S}^n)} + \frac{k^p-k}{p}\iota \circ \mathrm{tr}_{C_p}$. Precomposing with β, and postcomposing with $\widetilde{\alpha^p}$, we get

$$P_\beta(k\alpha) = kP_\beta(\alpha) + \frac{k^p - k}{p}\,\mathrm{tr}_{C_p}(\beta)\alpha^p,$$

since $\widetilde{\alpha^p} \circ \iota = \alpha^p$ by construction of $\widetilde{\alpha^p}$. ☐

Corollary 11.6.

$$P_\beta(p^e\alpha) = p^e P_\beta(\alpha) + p^{e-1}(p^{(p-1)e} - 1)\,\mathrm{tr}(\beta)\alpha^p$$

Corollary 11.7. *If $p^e\alpha = 0$, then $p^{e-1}\,\mathrm{tr}(\beta)\alpha^{p+1} = 0$ for any $\beta \in \pi_*(D_p(\mathbb{S}^n))$.*

Proof. From the previous corollary, we see

$$0 = P_\beta(p^e\alpha) = p^e P_\beta(\alpha) + (p-1)p^{e-1}\frac{p^{(p-1)e} - 1}{p^{p-1} - 1}\,\mathrm{tr}(\beta)\alpha^p$$

Multiplying with α, the first term becomes 0, and the second term becomes

$$(p-1)p^{e-1}\frac{p^{(p-1)e} - 1}{p^{p-1} - 1}\,\mathrm{tr}(\beta)\alpha^{p+1}.$$

Now both $(p-1)$ and $\frac{p^{(p-1)e}-1}{p^{p-1}-1}$ are invertible mod p^e (the latter by writing it as a geometric sum), implying the claim. ☐

To see how this could be helpful in proving nilpotence, observe that if transfers were sufficiently nontrivial and we could always find β with $\mathrm{tr}(\beta) = \alpha$, the above would show that $p^{e-1}\alpha^{p+2} = 0$, i.e. that some power of α is torsion of strictly smaller order. Iterating this e times, we find a power of α which is 0. The truth is not far from the above.

Theorem 11.8 (Kahn-Priddy)**.** *1. The transfer map $\mathbb{S}_{hC_p} \to \mathbb{S}$ is surjective on the p-local part of π_m for $m > 0$ (C_p acting trivially on $\mathbb{S}$ here).*

2. *More generally, for any $d > 0$ there exists v such that the transfer map*

$$D_p^{C_p}(\mathbb{S}^n) \to \mathbb{S}^{pn}$$

is surjective on the p-local part of π_m for $pn < m < pn + d$ whenever $p^v \mid n$.

The second part is more general than the first since for $n = 0$ the C_p-action on $\mathbb{S}^{\otimes p}$ is the trivial one, so that $D_p(\mathbb{S}) = \mathbb{S}_{hC_p}$. But it turns out the second statement can be deduced from the first by a comparatively easy argument, while the first part seems to be a rather deep fact of stable homotopy theory.

Remark 11.9. Note that the map $\mathbb{S}_{hC_p} \to \mathbb{S}$ (the transfer) here is *not* the one we used above to define $\widetilde{\alpha^p}$ (obtained from triviality of the action). There exists a map $\mathbb{S}_{hG} \to \mathbb{S}$ extending the identity $\mathbb{S} \to \mathbb{S}$ obtained from triviality of the action, and this would be surjective on homotopy groups for trivial reasons (it is a retraction), but this is not the transfer. Instead, for the transfer the composition $\mathbb{S} \to \mathbb{S}_{hG} \to \mathbb{S}$ is the degree $|G|$ map. Also note that the transfer does not coincide with the map $\mathbb{S}_{hG} \to \mathbb{S}$ extended from the degree $|G|$ map using triviality of the action either (that would not be surjective on positive-degree homotopy groups because it is $|G|$ times the retraction above). If this is confusing, remember that we have an equivalence arising from adjunctions

$$\mathrm{Map}_{\mathrm{Sp}}(\mathbb{S}_{hG}, \mathbb{S}) \simeq \mathrm{Map}_{\mathrm{Fun}(BG,\mathrm{Sp})}(\mathbb{S}^{\mathrm{triv}}, \mathbb{S}^{\mathrm{triv}}),$$

but it is *not* true that the latter is just $\mathrm{Map}_{\mathrm{Sp}}(\mathbb{S}, \mathbb{S})$ as it would be the case in ordinary algebra. In a sense, upgrading a map of spectra to a G-equivariant map is not just a property, it is extra structure, and the transfer corresponds to a particular way of viewing a degree $|G|$ map from $\mathbb{S} \to \mathbb{S}$ as equivariant. (Informally, it corresponds more to something like $\sum_{g \in G} \rho_g$ instead of $|G| \cdot \mathrm{id}$.)

For $C_p \subseteq \Sigma_p$, the transfer for Σ_p factors as

$$D_p^{\Sigma_p}(\mathbb{S}^n) \to D_p^{C_p}(\mathbb{S}^n) \to \mathbb{S}^{kn},$$

where the first map is a relative transfer (for any subgroup $H \subseteq G$ and any spectrum X with G-action, there is a transfer map $X_{hG} \to X_{hH}$, constructed similarly to the absolute case $H = e$ we discussed.) So the statement reduces to an analogous statement for $G = \Sigma_p$. In that case, there is an ingenious argument due to Segal, which we express in a slightly more modern way in the next section.

Proof of the second part of the Kahn-Priddy theorem from the first. If the C_p action on $(\mathbb{S}^n)^{\otimes p}$ were equivalent to the trivial action on $\mathbb{S}^{pn}$, this would just follow from a shift of the first statement (by pn). Equivariant maps

$$\mathbb{S}^{pn} \to (\mathbb{S}^n)^{\otimes p}$$

with trivial action on the left, are exactly maps $\mathbb{S}^{pn} \to ((\mathbb{S}^n)^{\otimes p})^{hC_p}$, and so we can try to find an element of $\pi_{pn}(((\mathbb{S}^n)^{\otimes p})^{hC_p}$ refining the obvious non-equivariant equivalence (and see what goes wrong, because there is no such equivalence). The homotopy groups $\pi_*(X^{hG})$ can be calculated by the *homotopy fixed point spectral sequence*

$$H^s(BG; \pi_t(X)) \Rightarrow \pi_{t-s}(X^{hG}).$$

(Adams grading: d_r increases s by r and decreases $t-s$ by 1). This can be constructed by pointwise Whitehead-filtering the diagram $BG \to \mathrm{Sp}$ and passing to the filtration on limits, it is a special case of the filtration on limits of functors $B \to \mathrm{Sp}$ we used to construct the Serre spectral sequence.

The action of C_p on $\pi_{pn}((\mathbb{S}^n)^{\otimes p})$ is given by the sign representation for $p = 2$ and n odd, and the trivial one otherwise. So we have an element $\iota_n \in H^0(BC_p; \pi_{pn}((\mathbb{S}^n)^{\otimes p}))$ coming from the canonical generator of $\pi_{pn}((\mathbb{S}^n)^{\otimes p})$ for n even if $p = 2$ and n arbitrary otherwise. If this element does not lift to homotopy groups, it means that it must support a nontrivial differential d_r. Since the cohomology of BC_p is p-torsion, we must have

$$d_r(\iota_{pn}) = d_r(\iota_n^p) = p d_r(\iota_n)\iota_n^{p-1} = 0,$$

and so $\iota_{pn} \in H^0(BC_p; \pi_{p^2n}((\mathbb{S}^{pn})^{\otimes p}))$ survives even longer in the homotopy fixed point spectral sequence for $((\mathbb{S}^{pn})^{\otimes p})^{hC_p}$. Iterating that, for any r we find v such that all differentials up to d_r vanish on ι_n if $p^v \mid n$. Since $d_r(\iota_n) \in H^r(BC_p; \pi_{pn+r-1}((\mathbb{S}^n)^{\otimes p})$, this means that ι_n is a permanent cycle in the homotopy fixed point spectral sequence for

$$(\tau_{\leq pn+r-2}((\mathbb{S}^n)^{\otimes p}))^{hC_p}.$$

Lifting that class to homotopy, we get an equivariant map

$$\mathbb{S}^{pn} \to \tau_{\leq pn+r-2}((\mathbb{S}^n)^{\otimes p})$$

which is an isomorphism on π_{pn}, hence the underlying nonequivariant map is exactly the truncation map

$$\mathbb{S}^{pn} \to \tau_{\leq pn+r-2}\mathbb{S}^{pn}.$$

Since the target is $(pn+r-2)$-coconnective, it factors furthermore canonically through a C_p-equivariant map

$$\tau_{\leq pn+r-2}\mathbb{S}^{pn} \to \tau_{\leq pn+r-2}((\mathbb{S}^n)^{\otimes p}),$$

which is an equivalence as we can see by again forgetting the C_p action. If we take $r \geq d+1$, then in the diagram of transfers

$$\begin{array}{ccccccc}
\mathbb{S}^{pn}_{hC_p} & \longrightarrow & (\tau_{\leq pn+r-2}\mathbb{S}^{pn})_{hC_p} & \longrightarrow & (\tau_{\leq pn+r-2}((\mathbb{S}^n)^{\otimes p}))_{hC_p} & \longleftarrow & ((\mathbb{S}^n)^{\otimes p})_{hC_p} \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\mathbb{S}^{pn} & \longrightarrow & \tau_{\leq pn+r-2}\mathbb{S}^{pn} & \longrightarrow & \tau_{\leq pn+r-2}((\mathbb{S}^n)^{\otimes p}) & \longleftarrow & (\mathbb{S}^n)^{\otimes p}
\end{array}$$

all horizontal maps are isomorphisms on π_m for $m < pn+d$. Since the first map is just Σ^{pn} of the transfer map in the first part of the Kahn-Priddy theorem, it is surjective on π_m for $pn < m$, showing the second part of the Kahn-Priddy theorem assuming the first. □

Proof of Nishida's nilpotence theorem. Given $\alpha \in \pi_m(\mathbb{S})$ with $n > 0$, we want to prove that some power of α is zero. Since the positive-degree homotopy groups of $\mathbb{S}$ are torsion, α is torsion of some order, and it suffices to show individually that all p-local parts of α are nilpotent, so we may assume $p^e\alpha = 0$ for some e. Starting with $k_0 = 1$, we will inductively prove that for every $i \leq e$, there exists a power α^{k_i} which satisfies $p^{e-i}\alpha^{k_i} = 0$, for $i = e$ we will have found a power which is zero.

By the second statement of the Kahn-Priddy theorem, there exists v such that for every n divisible by p^v, the transfer map

$$\mathrm{tr} : \pi_{pn+m}(D_p^{C_p}(\mathbb{S}^n)) \to \pi_{pn+m}(\mathbb{S}^{pn})$$

has $\alpha \in \pi_{pn+m}(\mathbb{S}^{pn})$ in its image. Choose some $k_i' \geq k_i$ such that $n = mk_i'$, the degree of $\alpha^{k_i'}$, is divisible by p^v, and choose a preimage β under the transfer map. Then by Corollary 11.7 applied to $p^{e-i}\alpha^{k_i'} = 0$, we learn

$$p^{e-i-1}\,\mathrm{tr}_{C_p}(\beta)\alpha^{(p+1)k_i'} = 0,$$

and since $\mathrm{tr}_{C_p}(\beta) = \alpha$ we can take $k_{i+1} = (p+1)k_i' + 1$. □

Remark 11.10. One can extract explicit upper bounds for the resulting exponent depending only on the starting degree m, but they are quite big: The v one gets from the second part of the Kahn-Priddy theorem is of the form $m + O(1)$, so already k_0' will generally have to be to the bigger than p^m, and in each further step we are at least multiplying the exponent by another factor of $(p+1)$. These are very far from optimal as far as we can tell from explicit computations in the range where $\pi_*(\mathbb{S})$ is known.

11.3 Segal's proof of the Kahn-Priddy theorem

In this section, we give a version of Segal's proof of the Kahn-Priddy theorem from [20]. Where Segal uses a version of the Barratt-Priddy-Quillen theorem "in families" which allows him to work in the homotopy category, we instead use the ∞-categorical universal property of $\Omega^\infty\mathbb{S}$ as "free commutative group on one generator". The key idea of the proof remains the same, an ingenious mixture of the multiplicative and additive structures available on $\Omega^\infty R$ for certain commutative ring spectra R.

Here, we will need to black-box some results about commutative algebras which really come from the operad used to represent them. After that operad, they are also called E_∞ algebras in this context. For any spectrum X, there is a free E_∞ ring on X, which has the form

$$\mathrm{Free}^{E_\infty}(X) \simeq \bigoplus_{0\leq k}(X^{\otimes k})_{h\Sigma_k} = \bigoplus_{0\leq k} D_k(X).$$

For formal reasons, one also can form a truncated version

$$\mathrm{Free}^{E_\infty}_{\leq n}(X) \simeq \bigoplus_{0\leq k\leq n} D_k(X)$$

(where "all higher products are taken to be 0"). If R is an E_∞ ring spectrum, $\Omega^\infty R$ admits an E_∞ monoid structure coming from the fact that $\Omega^\infty : \mathrm{Sp} \to \mathcal{S}$ is lax symmetric monoidal, this is different from the usual E_∞ monoid structure one has on $\Omega^\infty X$ of any spectrum. (If we think of that one as "additive", the one for rings can be thought of as "multiplicative", and in fact for R an Eilenberg-MacLane ring, $\Omega^\infty R$ is a discrete ring and these exactly recover the additive and multiplicative monoid structures.) One writes

$$\mathrm{GL}_1(R) \subseteq \Omega^\infty R$$

for the full subspace on all elements of $\pi_0(R)$ which are invertible with respect to that multiplicative monoid structure, this is an E_∞ group. If $R \to \mathbb{S}$ is an augmented E_∞ ring, like the free E_∞ rings above, we will also write $\mathrm{SL}_1(R)$ for the fiber of $\mathrm{GL}_1(R) \to \mathrm{GL}_1(\mathbb{S})$ (this is not standard notation.)

Now let X be connective and $R = \mathrm{Free}^{E_\infty}_{\leq n}(X)$. Since Ω^∞ preserves products, we have that

$$\Omega^\infty R \simeq \prod_{0\leq k\leq n} \Omega^\infty(D_k(X))$$

and since X is connective, $\pi_0(\Omega^\infty R) = \pi_0(R) = \mathrm{Sym}^{\leq n}(\pi_0(X))$, the truncated version of the symmetric algebra on $\pi_0(X)$. Since every element in positive degree is nilpotent, the invertible elements are exactly $\pm 1 + r$ with r of positive degree, and so

$$\mathrm{GL}_1(R) \simeq \mathrm{GL}_1(\mathbb{S}) \times \prod_{1\leq k\leq n} \Omega^\infty(D_k(X)),$$

and

$$\mathrm{SL}_1(R) \simeq \prod_{1\leq k\leq n} \Omega^\infty(D_k(X)),$$

whose π_0 consists of those units in $\pi_0(R)$ of the form $1 + r$ with r of positive degree.

The multiplicative E_∞ structure is generally subtle, but for $n = 1$ it agrees with the additive E_∞ structure on $\Omega^\infty X$.

Finally, we will require that $\Omega^\infty\mathbb{S}$ is the free E_∞ group on one generator. Specifically, this means that there is an equivalence

$$\mathrm{Map}_{\mathrm{Grp}_{E_\infty}(\mathcal{S})}(\Omega^\infty\mathbb{S}, G) \simeq \mathrm{Map}_{\mathcal{S}}(\mathrm{pt}, G) \simeq G,$$

meaning that E_∞ group homomorphisms $\Omega^\infty\mathbb{S} \to G$ are specified up to contractible choice by points in G.

Segal's proof of the Kahn-Priddy theorem now proceeds by proving something slightly stronger, namely that there exists a map

$$\Omega^\infty\mathbb{S} \to \Omega^\infty(\mathbb{S}_{h\Sigma_p})$$

such that the composite

$$\Omega^\infty \mathbb{S} \to \Omega^\infty(\mathbb{S}_{h\Sigma_p}) \xrightarrow{\Omega^\infty(\mathrm{tr}_{\Sigma_p})} \Omega^\infty \mathbb{S}$$

is an isomorphism on the p-local part of π_m for $m > 0$. So the surjectivity in the Kahn-Priddy theorem can actually be upgraded to the existence of a splitting, *but this splitting cannot be chosen as a map of spectra!* Instead it only exists after passing to the underlying spaces.

Writing $R = \mathrm{Free}^{E_\infty}_{\leq p}(\mathbb{S})$, and x for the element in $\pi_0(R)$ coming from the generator $\mathbb{S}$ in the $k = 1$ factor, we have that $1 + x$ is a unit of $\pi_0(R)$ (taken to 1 by the augmentation), and so the universal property of $\Omega^\infty \mathbb{S}$ gives us a map

$$\Omega^\infty \mathbb{S} \xrightarrow{1+x} \mathrm{SL}_1(R) \simeq \prod_{1 \leq k \leq p} \Omega^\infty(\mathbb{S}_{h\Sigma_k}),$$

and Segal proves that the map into the $k = p$ factor here does the job. To see this, we will check that the composite

$$\Omega^\infty \mathbb{S} \xrightarrow{1+x} \prod_{1 \leq k \leq p} \Omega^\infty(\mathbb{S}_{h\Sigma_k}) \to \Omega^\infty(\mathbb{S}_{h\Sigma_p}) \xrightarrow{\Omega^\infty(\mathrm{tr}_{\Sigma_p})} \Omega^\infty \mathbb{S} \qquad (1)$$

is an isomorphism on positive-degree homotopy groups.

Lemma 11.11. *The composition* (1) *is equivalent to the composition*

$$\Omega^\infty \mathbb{S} \xrightarrow{1+x_1+\ldots+x_p} \prod_{1 \leq k \leq p} \Omega^\infty(D_k(\mathbb{S}^{\oplus p})) \to \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) \to \Omega^\infty \mathbb{S},$$

where $x_1 + \ldots + x_p$ is the element of $\pi_0(\mathrm{Free}^{E_\infty}_{\leq p}(\mathbb{S}^{\oplus p}))$ coming from the sum of the generators analogous to x, and where the last map is obtained by applying homotopy orbits to

$$(\mathbb{S}^{\oplus p})^{\otimes p} \simeq \bigoplus_{p^p} \mathbb{S} \to \bigoplus_{p \xrightarrow{\cong} p} \mathbb{S}$$

by projecting onto the summands corresponding to bijections $p \to p$.

Proof. This is the commutative diagram

$$\begin{array}{ccccccc}
\Omega^\infty \mathbb{S} & \xrightarrow{1+x} & \prod_{1 \leq k \leq p}(\Omega^\infty(D_k(\mathbb{S}))) & \longrightarrow & \Omega^\infty(D_p(\mathbb{S})) & & \\
& \searrow^{1+x_1+\ldots+x_p} & \downarrow & & \downarrow & \searrow^{\Omega^\infty \mathrm{tr}_{\Sigma_p}} & \\
& & \prod_{1 \leq k \leq p}(\Omega^\infty(D_k(\mathbb{S}^{\oplus p}))) & \longrightarrow & \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) & \longrightarrow & \Omega^\infty \mathbb{S},
\end{array}$$

where the vertical maps are induced by the diagonal $\mathbb{S} \to \mathbb{S}^{\oplus p}$. Commutativity of the left triangle is clear by the universal property of $\Omega^\infty \mathbb{S}$ and the fact that this diagonal takes $1 + x$ to $1 + x_1 + \ldots + x_p$, whereas commutativity of the right triangle comes from the fact that before homotopy orbits, the composite

$$\mathbb{S} \to \bigoplus_{p^p} \mathbb{S} \to \bigoplus_{p \xrightarrow{\simeq} p} \mathbb{S}$$

of the diagonal followed by the projection is the same map as the one we used to define the transfer. (Since the target has a cofree Σ_p-action, such maps are determined by the map of spectra $\mathbb{S} \to \mathbb{S}$ obtained by composing with any one projection, and these are the identity in both cases.) $\square$

Lemma 11.12. *For $S \subsetneq \{1, \ldots, p\}$ any proper subset, the composite*

$$\Omega^\infty \mathbb{S} \xrightarrow{1+\sum_{s \in S} x_s} \prod_{1 \le k \le p} \Omega^\infty(D_k(\mathbb{S}^{\oplus p})) \to \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) \to \Omega^\infty \mathbb{S},$$

is zero.

Proof. This is the commutative diagram

$$\begin{array}{ccccccc}
\Omega^\infty \mathbb{S} & \xrightarrow{1+\sum x_s} & \prod_{1 \le k \le p}(\Omega^\infty(D_k(\mathbb{S}^{\oplus |S|}))) & \longrightarrow & \Omega^\infty(D_p(\mathbb{S}^{\oplus S})) & & \\
& \searrow^{1+\sum x_s} & \downarrow & & \downarrow & \searrow^{0} & \\
& & \prod_{1 \le k \le p}(\Omega^\infty(D_k(\mathbb{S}^{\oplus p}))) & \longrightarrow & \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) & \longrightarrow & \Omega^\infty \mathbb{S},
\end{array}$$

where the right triangle commutes since the composite

$$(\mathbb{S}^{\oplus |S|})^{\otimes p} \simeq \bigoplus_{|S|^p} \mathbb{S} \to \bigoplus_{p^p} \mathbb{S} \to \bigoplus_{p \xrightarrow{\simeq} p} \mathbb{S}$$

is null: No bijection $p \to p$ factors through a proper subset. $\square$

Lemma 11.13. *Let M, N be E_∞ monoids, and $f, g : M \to N$ E_∞ monoid morphisms. Then $f \cdot g$, defined as composite*

$$M \xrightarrow{(f,g)} N \times N \xrightarrow{\mu} N$$

is also an E_∞ monoid morphism, and the induced maps on $\pi_m(-, 1)$ for $m > 0$ ($1 \in M, N$ denoting the neutral elements) satisfy

$$\pi_m(fg) = \pi_m(f) + \pi_m(g).$$

Proof. As these are E_∞ monoids, $N \times N \to N$ induces a commutative monoid structure on all $\pi_m(-, 1)$, which by Eckmann-Hilton agrees with the existing group structure. Applying π_m to the composite defining $f \cdot g$ now gives

$$\pi_m(M, 1) \xrightarrow{(\pi_m(f), \pi_m(g))} \pi_m(N, 1) \times \pi_m(N, 1) \xrightarrow{+} \pi_m(N, 1),$$

proving the claim. $\square$

(This holds of course more generally for H-spaces instead of E_∞ monoids.)

Remark 11.14. One has to be careful with basepoints here: If $x \in M$ is an arbitrary point, the correct formula expressing $\pi_m(fg, x) : \pi_m(M, x) \to \pi_m(N, f(x)g(x))$ also involves the maps induced by multiplication with $f(x)$ and $g(x)$ on $\pi_m(N)$.

Lemma 11.15. *The map induced on π_m for $m > 0$ by the composition* (1) *is equivalent to the map induced by the composition*

$$\Omega^\infty \mathbb{S} \xrightarrow{\prod_{S\subseteq\{1,\ldots,p\}}(1+\sum_{s\in S} x_s)^{(-1)^{p-|S|}}} \prod_{1\leq k\leq p} \Omega^\infty(D_k(\mathbb{S}^{\oplus p})) \to \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) \to \Omega^\infty \mathbb{S},$$

Proof. As the map on the left is an alternating product, by the previous lemma its effect on π_m is an alternating sum. The second map is just a map of spaces, but of course induces group homomorphisms on π_m for $m > 0$. This means that the effect of the above composite on π_m for $m > 0$ is given by an alternating sum. Here one summand (the one where $S = \{1, \ldots, p\}$) exactly corresponds to the composition (1), whereas the rest acts by 0 by Lemma 11.12. □

Lemma 11.16.

$$\prod_{S\subseteq\{1,\ldots,p\}} (1 + \sum_{s\in S} x_s)^{(-1)^{p-|S|}} = 1 + (-1)^{p-1}(p-1)!x_1 \cdots x_p + \textit{ higher terms}$$

Proof. We may check this in the ring of formal power series $\mathbb{Z}[[x_1, \ldots, x_p]]$, which is a subring of $\mathbb{Q}[[x_1, \ldots, x_p]]$, so we may check it there. Applying logarithms, it amounts to checking that

$$\sum_{S\subseteq\{1,\ldots,p\}} (-1)^{p-|S|} \left(\sum_{s\in S} x_s\right)^k = \begin{cases} 0 \text{ for } k < p \\ p!x_1 \cdots x_p \text{ for } k = p, \end{cases}$$

which is clear, since this alternating sum gives precisely an inclusion-exclusion description isolating those monomials in $(x_1 + \ldots + x_p)^k$ which involve all variables at least once. For $i < k$, there are none such, whereas for $k = p$ we get exactly $x_1 \cdots x_p$ with multiplicity $p!$. So log of the left-hand side evaluates to

$$(-1)^{p-1}\frac{p!x_1 \cdots x_p}{p} + \text{ higher terms.}$$

Applying exp again, we obtain

$$1 + (-1)^{p-1}(p-1)!x_1 \cdots x_p + \text{ higher terms.}$$

□

Proof of the first part of the Kahn-Priddy theorem. From the above discussions, we need to check that the composite

$$\Omega^\infty\mathbb{S} \xrightarrow{1+(-1)^{p-1}(p-1)!x_1\cdots x_p} \prod_{1\le k\le p}\Omega^\infty(D_k(\mathbb{S}^{\oplus p})) \to \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) \to \Omega^\infty\mathbb{S}$$

induces isomorphisms on the p-local part of π_m for $m > 0$. The element $(-1)^{p-1}(p-1)!x_1\cdots x_p \in \pi_0(D_p(\mathbb{S}^{\oplus p}))$ comes from a map $\mathbb{S} \to D_p(\mathbb{S}^{\oplus p})$, which determines a map

$$\mathrm{Free}^{E_\infty}_{\le 1}(\mathbb{S}) \to \mathrm{Free}^{E_\infty}_{\le p}(\mathbb{S}^{\oplus p}).$$

Applying Ω^∞, we obtain a diagram

$$\begin{array}{ccccccc}
\Omega^\infty\mathbb{S} & \xrightarrow{1+t} & \Omega^\infty(\mathbb{S}) & \xrightarrow{=} & \Omega^\infty(\mathbb{S}) & & \\
& \searrow^{1+(-1)^{p-1}(p-1)!x_1\cdots x_p} & \downarrow & & \downarrow & \searrow^{(-1)^{p-1}(p-1)!} & \\
& & \prod_{1\le k\le p}\Omega^\infty(D_k(\mathbb{S}^{\oplus p})) & \longrightarrow & \Omega^\infty(D_p(\mathbb{S}^{\oplus p})) & \longrightarrow & \Omega^\infty\mathbb{S},
\end{array}$$

where the $\Omega^\infty(\mathbb{S})$ at the top middle position arises as $\mathrm{SL}_1(\mathrm{Free}^{E_\infty}_{\le 1}(\mathbb{S}))$, and we write t for the canonical generator of $\pi_0(\mathrm{Free}^{E_\infty}_{\le 1}(\mathbb{S}))$. The left triangle commutes by construction. For the right triangle, observe that the first map can be described before taking homotopy orbits as $(-1)^{p-1}(p-1)!$ times the inclusion

$$\bigoplus_{p\xrightarrow{\simeq}p}\mathbb{S} \to \bigoplus_{p^p}\mathbb{S},$$

and the second as the corresponding projection, so their composite is just $(-1)^{p-1}(p-1)!$ times the identity.

Finally, since the multiplicative E_∞ structure on $\mathrm{SL}_1(\mathrm{Free}^{E_\infty}_{\le 1}(\mathbb{S})) \simeq \Omega^\infty\mathbb{S}$ agrees with the additive one, the top left horizontal map is an equivalence, and since $(-1)^{p-1}(p-1)!$ is coprime to p, the right diagonal map also induces an isomorphism on the p-local part of π_m, finishing the proof. □

12 Synthetic spectra

12.1 Descent and limits of compactly generated categories

We saw that we can get a powerful spectral sequence computing the homotopy groups of $\mathbb{S}$ by exploiting the fact that

$$\mathbb{S}_p^\wedge \simeq \lim_{\Delta}(\ \mathbb{F}_p \rightrightarrows \mathbb{F}_p\otimes_{\mathbb{S}}\mathbb{F}_p \qquad \dots\ )$$

More generally, we can use this to write connective p-complete spectra X as limit over $\mathbb{F}_p\otimes_{\mathbb{S}}X$, $\mathbb{F}_p\otimes_{\mathbb{S}}\mathbb{F}_p\otimes_{\mathbb{S}}X$ and so on. We also saw in Proposition 8.3 that more generally, this works for $S \to R$ any map of connective commutative ring spectra, allowing us to descend bounded below p-complete S-modules if

$\pi_0(S) \to \pi_0(R)$ is surjective and $\pi_0(S)/p \to \pi_0(R)/p$ is an isomorphism, and all bounded below S-modules if $\pi_0(S) \to \pi_0(R)$ is an isomorphism.

We now want to expand on this a bit, and instead of individual S-modules, try to compare the whole category $\mathrm{Mod}_S(\mathrm{Sp})$ with the corresponding limit of categories, i.e. analyze the functor

$$\mathrm{Mod}_S(\mathrm{Sp}) \to \lim_{\Delta}(\ \mathrm{Mod}_R(\mathrm{Sp}) \rightrightarrows \mathrm{Mod}_{R\otimes_S R}(\mathrm{Sp}) \qquad \dots\ ).$$

Informally, we can think of an object in this limit as a tuple of $M_s \in \mathrm{Mod}_{R^{\otimes_S s+1}}(\mathrm{Sp})$ together with coherent identifications of all possible base-changes of the $M_\bullet$ under the ring maps in the cosimplicial diagram. Since all these categories and functors are symmetric-monoidal, the limit inherits a symmetric-monoidal structure as well, its unit $\mathbb{1}$ is the object which informally is given by $R^{\otimes_S s+1}$ in the s-th term.

As in the construction of the Adams spectral sequence, everything will be based on some connectivity argument, which we summarize in the following proposition.

Proposition 12.1. *1. If $R^\bullet$ is a cosimplicial commutative ring spectrum (i.e. a functor $\Delta \to \mathrm{CAlg}(\mathrm{Sp})$), and $M_\bullet$ an object of*

$$\lim_{\Delta} \mathrm{Mod}_{R^\bullet}(\mathrm{Sp}),$$

we have for the underlying cosimplicial diagram in Sp*:*

$$\mathrm{fib}(\lim_{\Delta_{\leq n}} M_\bullet \to \lim_{\Delta_{\leq n-1}} M_\bullet) \simeq (\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)) \otimes_{R^0} M_0.$$

2. *More generally,*

$$\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)$$

admits the structure of an R^0-R^0 bimodule, and for any $N_\bullet, M_\bullet$, we have that

$$\begin{aligned}\mathrm{fib}(&\lim_{\Delta_{\leq n}} \mathrm{map}_{R^\bullet}(N_\bullet, M_\bullet) \to \lim_{\Delta_{\leq n-1}} \mathrm{map}_{R^\bullet}(N_\bullet, M_\bullet))\\ &\simeq \mathrm{map}_{R^0}(N_0, (\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)) \otimes_{R^0} M_0),\end{aligned}$$

where the target on the right inherits its R^0-module structure from the left module structure on the bimodule, and the tensor product uses the right module structure.

3. *If $R^\bullet = R^{\otimes_S \bullet+1}$ is the Cech nerve of a map $S \to R$ of connective commutative ring spectra which is surjective on π_0, then*

$$\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)$$

is connective.

We will not prove this here, as it requires some manipulations of limits using Kan extensions, similar to Proposition 8.1 (of which the last statement is a special case).

Lemma 12.2. *Assume $S \to R$ is a map of connective commutative ring spectra which is surjective on π_0. Then if $M_\bullet \in \lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$ has M_0 k-connective, it admits a "cell structure": A filtration with $\mathrm{gr}_{-n} M_\bullet \simeq \bigoplus_{\Gamma_n} \Sigma^n \mathbb{1}$ for some sets Γ_n, nontrivial only for $n \geq k$.*

Proof. The key observation is that

$$\mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(\mathbb{1}, M_\bullet) \simeq \lim_\Delta M_\bullet$$

using the description of mapping spectra in a limit of stable ∞-categories as limit, and that we are mapping out of the free rank 1 module levelwise. Now the fiber of

$$\lim_{\Delta_{\leq n}} \mathrm{map}_{\mathrm{Mod}_R(\mathrm{Sp})}(R, M_\bullet) \to \lim_{\Delta_{\leq n-1}} \mathrm{map}_{\mathrm{Mod}_R(\mathrm{Sp})}(R, M_\bullet),$$

by Proposition 12.1, is k-connective if M_0 is. Since all of those fibers are k-connective, it follows that

$$\pi_k \, \mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(\mathbb{1}, M_\bullet) \to \pi_k \, \mathrm{Mod}_R(\mathrm{Sp})(R, M_0)$$

is surjective. Picking generators, we find a map $\bigoplus_{\Gamma_k} \Sigma^k \mathbb{1} \to M_\bullet$ whose cofiber is $k+1$-connective in cosimplicial degree 0. Iterating this gives our cell structure. To see that the colimit of the skeleta for $M_\bullet$ constructed this way is indeed $M_\bullet$, we observe that the functors in the diagram of categories giving

$$\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$$

are all colimit-preserving, so colimits are formed "pointwise". Since the cofibers constructed above become increasingly connective in cosimplicial degree 0, their colimit vanishes in cosimplicial degree 0, and since M_i is a base-change of M_0, in all cosimplicial degrees. □

Theorem 12.3. *Let $S \to R$ be a map of connective commutative ring spectra. Then:*

1. *If $\pi_0(S) \to \pi_0(R)$ is an isomorphism, the functor $F : \mathrm{Mod}_S(\mathrm{Sp}) \to \lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}(\mathrm{Sp})}$ induces equivalences*

$$\mathrm{map}_{\mathrm{Mod}_S(\mathrm{Sp})}(X, Y) \to \mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(F(X), F(Y)),$$

whenever Y is bounded below. Furthermore, all objects of $\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$ whose underlying object in $\mathrm{Mod}_R(\mathrm{Sp})$ is bounded below are of the form $F(X)$ for some bounded below $X \in \mathrm{Mod}_S(\mathrm{Sp})$.

2. *If $\pi_0(S) \to \pi_0(R)$ is surjective, $\pi_0(S)/p \to \pi_0(R)/p$ is an isomorphism, and $\pi_0(R)$ is bounded p-torsion, then $F : \mathrm{Mod}_S(\mathrm{Sp}) \to \lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}(\mathrm{Sp})}$ induces equivalences*

$$\mathrm{map}_{\mathrm{Mod}_S(\mathrm{Sp})}(X,Y) \to \mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(F(X),F(Y))$$

whenever Y is bounded below and p-complete. Furthermore, all objects of $\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$ whose underlying object in $\mathrm{Mod}_R(\mathrm{Sp})$ is bounded below are of the form $F(X)$ for some bounded below and p-complete $X \in \mathrm{Mod}_S(\mathrm{Sp})$.

Proof. Since colimits in the limit of categories are formed pointwise, the functor $\mathrm{Mod}_S \to \lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$ preserves colimits. Over the connective ring S, every bounded below module has a cell structure, and in particular we can produce every bounded below module from S using colimits and shifts. So the part of the first statement about mapping spectra reduces to the case $X = S$. Now, Theorem 8.3 tells us that the map

$$\mathrm{map}_{\mathrm{Mod}_S(\mathrm{Sp})}(S,Y) \to \mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(F(S),F(Y))$$

is an equivalence for $X = \mathbb{1} = S$ and Y bounded below. For the part about the essential image, recall from Lemma 12.2 above that any bounded below object in the target has a cell structure, using the fully faithful part we just proved one can inductively lift the attaching maps, the skeleta, and then their entire colimit.

For the second statement, we may again reduce the part about mapping spectra to the case $X = S$, which is covered again by Theorem 8.3. For a p-complete bounded below Y, this implies

$$\begin{aligned}\mathrm{map}_{\mathrm{Mod}_S}(X,Y) &\simeq \lim_n \mathrm{map}_{\mathrm{Mod}_S}(X, Y/p^n)\\ &\simeq \lim_n \mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(F(X), F(Y)/p^n)\\ &\simeq \mathrm{map}_{\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})}(F(X), \lim_n F(Y)/p^n).\end{aligned}$$

Limits in $\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$ are not in general formed pointwise, since it is a limit of categories along functors that don't preserve all limits. However, in this case the natural maps $F(Y) \to F(Y)/p^n$ do give a limit cone in each of the categories $\mathrm{Mod}_{R^{\otimes_S \bullet}}(\mathrm{Sp})$, using that multiplication with p^r for some r is zero on $R \otimes_S Y$ and more generally $R^{\otimes_S s+1} \otimes_S Y$. This is still sufficient to conclude that $F(Y)$ is a limit of the $F(Y)/p^n$'s in $\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}(\mathrm{Sp})}$, without saying anything about existence of more general limits. We also observe from this that $F(-)$ takes an S-module and its p-completion to the same value.

Now for the statement about the essential image, we again use that every object in the target has a cell structure. With the observation just made about F being insensitive to p-completion, we may lift $\bigoplus_{\Gamma_n} \mathbb{1}$ to the p-complete $(\bigoplus_{\Gamma_n} S)^\wedge_p$. The fully faithfullness part now allows us to lift any cell structure to a "p-complete cell structure" where the associated graded parts look like that.

The p-completion of the colimit of these p-complete skeleta now provide the desired lift. ☐

Corollary 12.4. *In the first case of Theorem 12.3, we have an equivalence*

$$\mathrm{Mod}_S(\mathrm{Sp})_+ \to \lim_{\Delta} \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})_+,$$

where + denotes the full subcategories on bounded below spectra. In the second case, we have an equivalence

$$\mathrm{Mod}_S(\mathrm{Sp})^{\wedge}_{p,+} \to \lim_{\Delta} \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})_+,$$

with the left hand side denoting the full subcategory of bounded below p-complete spectra.

Proof. The functors are fully faithful and essentially surjective by Theorem 12.3. ☐

So, more generally than just using the Adams resolution to compute maps between objects, it can even be used to describe objects in terms of "descent data". (Going back to our original motivation of descent, this is also maybe the more interesting half of classical faithfully flat descent.)

We want to turn the above statement into some kind of statement about the whole category $\mathrm{Mod}_S(\mathrm{Sp})$ (dropping the restriction to bounded below objects). A nice systematic way of doing this is using compactly generated categories.

Definition 12.5. Let $\mathcal{C}$ be an ∞-category with all small colimits.

1. An object $X \in \mathcal{C}$ is called *compact* if $\mathrm{Map}_{\mathcal{C}}(X, -)$ preserves *filtered colimits.*

2. $\mathcal{C}$ is called *compactly generated* if there exists a set of compact objects which generate all of $\mathcal{C}$ under small colimits.

Filtered colimits include sequential colimits. For intuition about this definition for compactness, it is instructive to observe that in Set, one has that $\mathrm{Hom}(X, -)$ commutes with sequential colimits if and only if X is finite.

We write $\mathcal{C}^{\omega} \subseteq \mathcal{C}$ for the full subcategory of compact objects. As an important example, one has:

Proposition 12.6. *Let S be a connective ring spectrum. Then for $M \in \mathrm{Mod}_S(\mathrm{Sp})$, the following are equivalent:*

1. *M is compact.*

2. *M is a retract of an M' with a finite filtration with $\mathrm{gr}_{-i} M' \simeq \bigoplus_{\Gamma_i} \Sigma^i S$, a finite direct sum.*

3. *M has a finite filtration where for each i, $\mathrm{gr}_{-i} M$ is a retract of a finite sum of copies of $\Sigma^i S$.*

4. *M is contained in the smallest full subcategory of* $\mathrm{Mod}_S(\mathrm{Sp})$ *which contains S and is closed under cofibers, fibers and retracts. (The thick subcategory generated by S.)*

Proof. For $1 \Rightarrow 2$, one may consider the full subcategory $\mathcal{C} \subseteq \mathrm{Mod}_S(\mathrm{Sp})$ of objects which admit such a filtration ("finite cell complexes") and consider the slice $\mathcal{C}_{/M}$. This is filtered since $\mathcal{C}$ admits pushouts (to see this, one uses a version of "cellular approximation"), and

$$\mathrm{colim}_{K \in \mathcal{C}_{/M}} K \to M$$

is an equivalence, as one checks on homotopy groups: It is surjective on π_n since each $\alpha \in \pi_n(M)$ comes from some $\Sigma^n S \to M$ (representing an object of $\mathcal{C}_{/M}$), and it is injective since if $\alpha \in \pi_n(K)$ is in the kernel of $K \to M$, $K \to M$ factors through $K' = \mathrm{cofib}(\alpha : \Sigma^n S \to K)$ making K' an object of $\mathcal{C}_{/M}$, and $\alpha = 0$ in $\pi_n(K')$. Now the identity $M \to M \simeq \mathrm{colim}_{K \in \mathcal{C}_{/M}} K$ factors through some $K \in \mathcal{C}$ by compactness.

For $2 \Rightarrow 3$, assume M is a retract of M' with such a cellular filtration where $\mathrm{gr}_{-i} M'$ vanishes for i outside of $[a, b]$. The $\pi_0(S)$-module $\pi_a(M')$ is a quotient of $\pi_a(\bigoplus_{\Gamma_a} \Sigma^a S)$, i.e. finitely generated. Also, we have a direct sum decomposition $M' = M \oplus N$, and so $\pi_a(M)$ and $\pi_a(N)$ are also finitely generated. If $a = b$, M is already a retract of a finite sum of copies of $\Sigma^a S$, and has a filtration as claimed with $\mathrm{gr}_{-i} M = M$ for $i = a$ and 0 otherwise. If $a < b$, we argue inductively as follows. Choosing a set of generators A_a for $\pi_a(M)$ and B_a for $\pi_a(N)$, we have:

$$\mathrm{cofib}(\bigoplus_{A_a \coprod B_a} \Sigma^a S \to M') \simeq \mathrm{cofib}(\bigoplus_{A_a} \Sigma^a S \to M) \oplus \mathrm{cofib}(\bigoplus_{B_a} \Sigma^a S \to N)$$

$$\mathrm{cofib}(\bigoplus_{A_a \coprod B_a \coprod \Gamma_a} \Sigma^a S \to M')$$

$$\simeq \mathrm{cofib}(\bigoplus_{A_a} \Sigma^a S \to M) \oplus \mathrm{cofib}(\bigoplus_{B_a} \Sigma^a S \to N) \oplus \bigoplus_{\Gamma_a} \Sigma^{a+1} S$$

$$\mathrm{cofib}(\bigoplus_{A_a \coprod B_a \coprod \Gamma_a} \Sigma^a S \to M')$$

$$\simeq \mathrm{cofib}(\bigoplus_{A_a \coprod B_a} \Sigma^a S \to M'/F_{\leq a} M') \simeq M'/F_{\leq a} M' \oplus \bigoplus_{A_a \coprod B_a} \Sigma^{a+1} S$$

where the second equivalence follows from the first since the left hand term of the second equivalence is obtained from the first one as further cofiber of a map from $\bigoplus_{\Gamma_a} \Sigma^a S$, but its target is $a+1$-connective. A similar argument implies the third equivalence. The third equivalence now shows that

$$\mathrm{cofib}(\bigoplus_{A_a \coprod B_a \coprod \Gamma_a} \Sigma^a S \to M')$$

has a cell structure with cells in the interval $[a+1, b]$, and the second equivalence shows that $\operatorname{cofib}(\bigoplus_{A_a} \Sigma^a S \to M)$ is a retract of it. Inductively, the latter term has a finite filtration with gr_{-i} trivial if $i \notin [a+1, b]$ and given as retracts of finite sums of $\Sigma^i S$ otherwise. This gives a filtration on M with $\mathrm{gr}_{-i} M \simeq \bigoplus_{A_a} \Sigma^a S$ as desired.

$3 \Rightarrow 4$ is clear since this thick subcategory contains retracts of finite sums of $\Sigma^a S$, and then inductively every object which admits a finite filtration with such associated graded terms.

Finally, $4 \Rightarrow 1$ follows from the fact that S itself is a compact objects and compact objects are closed under cofibers, shifts and retracts. □

One also denotes $\mathrm{Mod}_S(\mathrm{Sp})^\omega$ by $\mathrm{Perf}(S)$.

As homotopy classes of S-module maps $S^m \to S^n$ are in bijection to $n \times m$-matrices with entries in $\pi_0(S)$, it is not hard to see that homotopy classes of S-module idempotents on S^n correspond to $\pi_0(S)$-module idempotents on $\pi_0(S)^n$. In particular, retracts of finite sums of S correspond up to homotopy equivalence to isomorphism classes of finitely generated projective $\pi_0(S)$-modules. We will therefore also call an S-module "finitely generated projective" if it is a retract of S^n for some n.

Example 12.7. As $\pi_0(\mathbb{S}) = \mathbb{Z}$, and every finitely generated projective $\mathbb{Z}$-module is free, compact objects in $\mathrm{Sp} = \mathrm{Mod}_{\mathbb{S}}(\mathrm{Sp})$ are exactly the finite cell complexes.

A surprising feature of this story is that a compactly generated category $\mathcal{C}$ is fully determined by $\mathcal{C}^\omega$. For an essentially small ∞-category $\mathcal{C}_0$, one can form a category $\mathrm{Ind}(\mathcal{C}_0)$ by "freely adding filtered colimits". Formally, this is described as a full subcategory of $\mathrm{Fun}(\mathcal{C}_0^{\mathrm{op}}, \mathcal{S})$ obtained by closing the image of the Yoneda embedding $\mathcal{C}_0 \to \mathrm{Fun}(\mathcal{C}_0^{\mathrm{op}}, \mathcal{S})$ under filtered colimits. Another (more informal, but useful) perspective is that objects in $\mathrm{Ind}(\mathcal{C}_0)$ are represented by "formal filtered colimits" $\text{"}\mathrm{colim}\text{"}_I X_i$, with

$$\mathrm{Map}_{\mathrm{Ind}(\mathcal{C}_0)}(\text{"}\mathrm{colim}\text{"}_I X_i, \text{"}\mathrm{colim}\text{"}_J Y_j) \simeq \lim_I \mathrm{colim}_J \mathrm{Map}_{\mathcal{C}_0}(X_i, Y_j).$$

This construction $\mathrm{Ind}(\mathcal{C}_0)$ can also be characterized by a universal property: If $\mathcal{D}$ is an ∞-category with all filtered colimits, then arbitrary functors $\mathcal{C}_0 \to \mathcal{D}$ extend uniquely to filtered colimit preserving functors $\mathrm{Ind}(\mathcal{C}_0) \to \mathcal{D}$, which informally are obtained by taking formal colimits to actual colimits. In particular, for any $\mathcal{C}$ with all small colimits, one has a canonical functor

$$\mathrm{Ind}(\mathcal{C}^\omega) \to \mathcal{C},$$

and this turns out to be an equivalence if and only if $\mathcal{C}$ is compactly generated. If $\mathcal{C}$ is stable, then $\mathcal{C}^\omega$ is stable, and if $\mathcal{C}_0$ is stable, then $\mathrm{Ind}(\mathcal{C}_0)$ is stable (and actually turns out to have all small colimits, since these can be decomposed into filtered colimits and pushouts). One can organize these observations into the following statement:

Theorem 12.8. *The constructions* $(-)^\omega$ *and* Ind *provide inverse equivalences between the following* ∞*-categories:*

1. $\mathrm{Pr}^L_{\omega,\mathrm{st}}$, *whose objects are compactly generated stable ∞-categories with all small colimits, and where the morphisms are functors which preserve colimits and compact objects.*

2. $\mathrm{Cat}^{\mathrm{perf}}_\infty$, *whose objects are small stable ∞-categories which are idempotent complete (all idempotents split as projections to direct summands), and where the morphisms are exact functors.*

For example, for a map of ring spectra $S \to R$, the base-change functor $\mathrm{Mod}_S(\mathrm{Sp}) \to \mathrm{Mod}_R(\mathrm{Sp})$ preserves colimits and compact objects. It can therefore be recovered from its restriction $\mathrm{Perf}(S) \to \mathrm{Perf}(R)$ by applying Ind. Relatedly, base-change functors $\mathcal{D}(R) \to \mathcal{D}(S)$ for ordinary rings, or $\mathcal{D}^{\mathrm{gr}}(R) \to \mathcal{D}^{\mathrm{gr}}(S)$ for ordinary graded rings, are morphisms in $\mathrm{Pr}^L_{\omega,\mathrm{st}}$, and hence can be recovered from their restrictions to compact objects by applying Ind.

Now, when we talk about some limit like

$$\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*}) := \lim_{\Delta} \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*^{\otimes \bullet}),$$

we can interpret the diagram in different "categories of categories":

1. $\widehat{\mathrm{Cat}}_\infty$, the category of (possibly large) ∞-categories. Here, limits have the familiar description where objects are compatible families of objects in the diagram, and mapping spaces can be computed as limit.

2. $\mathrm{Pr}^L_{\mathrm{st}}$, the category of *presentable* stable ∞-categories. A presentable ∞-category is one which has all small colimits and is κ-compactly generated for some cardinal κ, and morphisms in $\mathrm{Pr}^L_{\mathrm{st}}$ are just colimit-preserving functors. Note that $\mathrm{Pr}^L_{\mathrm{st},\omega}$ is not a full subcategory of $\mathrm{Pr}^L_{\mathrm{st}}$, since in the former morphisms also have to preserve compact objects. The forgetful functor $\mathrm{Pr}^L_{\mathrm{st}} \to \widehat{\mathrm{Cat}}^{\infty}$ preserves limits, so interpreting the limit in $\mathrm{Pr}^L_{\mathrm{st}}$ doesn't change anything, and we still get the same category.

3. $\mathrm{Pr}^L_{\mathrm{st},\omega}$. Here, the forgetful functor $\mathrm{Pr}^L_{\mathrm{st},\omega} \to \mathrm{Pr}^L_{\mathrm{st}}$ does not preserve limits, so we get a different result. The limit can be computed using the fact that we have an equivalence $\mathrm{Pr}^L_{\omega,\mathrm{st}} \simeq \mathrm{Cat}^{\mathrm{perf}}_\infty$, and the forgetful functor $\mathrm{Cat}^{\mathrm{perf}}_\infty \to \mathrm{Cat}_\infty$ does actually preserve limits. So

$$\overset{\mathrm{Pr}^L_{\omega,\mathrm{st}}}{\lim_{\Delta}} \mathcal{D}^{\mathrm{gr}}(\mathcal{A}_*^{\otimes \bullet}) \simeq \mathrm{Ind}\left(\lim_{\Delta} \mathrm{Perf}(\mathcal{A}_*^{\otimes \bullet})\right) =: \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*}).$$

The limit on the inside is a full subcategory of $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$, but it is not the full subcategory of compact objects in $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$. Instead, it is the full subcategory consisting of all objects whose underlying derived graded $\mathbb{F}_p$-module in $\mathcal{D}^{\mathrm{gr}}(\mathbb{F}_p)$ is compact. In more elementary terms, if we think of objects in $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ as represented by complexes of comodules, these are the objects which can be represented by a bounded complex of finite-dimensional comodules. The resulting big category (after applying Ind) is generally better behaved than the naive $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$. For example, in $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$, the monoidal unit $\mathbb{1}$ is not compact, so if we want to think of the Ext-groups $[\Sigma^{-s}\mathbb{1}(t), -]$ as a version of "homotopy groups of derived comodules", we get unintuitive behavior with respect to filtered colimits as long as we work in $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$

We now want to see that this gives a nice reinterpretation of the limits in Corollary 12.4.

Proposition 12.9. *If $S \to R$ is a map of connective commutative ring spectra which is surjective on π_0, then*

$$\lim_{\Delta} \mathrm{Perf}(R^{\otimes_S \bullet+1}),$$

the full subcategory of $\lim_\Delta \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$ on all objects $M_\bullet$ where each M_i is compact, is generated as a thick subcategory by $\mathbb{1}$.

Proof. Observe that since the M_i are base-changes of M_0, all M_i are compact if and only if M_0 is compact. Now if M_0 is compact, by the characterisation of compact R-modules for connective R, M_0 is k-connective for some k and $\pi_k(M_0)$ is finitely generated as $\pi_0(R)$-module. Using Proposition 12.1, we again get that

$$\pi_k \operatorname{map}_{\lim_\Delta \mathrm{Perf}(R^{\otimes_S \bullet})}(\mathbb{1}, M_\bullet) \to \pi_k \operatorname{map}_{\mathrm{Perf}(R)}(R, M_0)$$

is surjective, and hence that we can find a map $\bigoplus \Sigma^k \mathbb{1} \to M_\bullet$ from a finite sum whose component $\bigoplus \Sigma^k R \to M_0$ is surjective on π_k. The cofiber is still compact but has degree 0 part $(k+1)$-connective. Now assume the compact R-module M_0 is a retract of a finite cell complex with cells in the range $[a, b]$ in the sense of Proposition 12.6, and that by iterating the above procedure we have built a cofiber sequence

$$K_\bullet \to M_\bullet \to M'_\bullet$$

where $K_\bullet$ is in the thick subcategory generated by $\mathbb{1}$, and $M'_\bullet$ is $b+2$-connective. The second part of Proposition 12.1 shows that each

$$\lim_{\Delta_{\leq n}} \operatorname{map}_{R^\bullet}(M_\bullet, M'_\bullet)$$

is at least 2-connective, since $\operatorname{map}_R(M_0, -)$ takes $b+2$-connective objects to 2-connective objects, as M_0 is a retract of an object with a finite filtration whose associated graded terms are sums of $\Sigma^i R$ with $i \leq b$. So their limit is at least 1-connective (here, a more precise $\lim^1$ argument would also work with 1 instead of 2), and so π_0 of the mapping spectrum from $M_\bullet$ to $M'_\bullet$ vanishes. It follows that the map $M_\bullet \to M'_\bullet$ above is nullhomotopic, and hence that $M_\bullet$ is a retract of $K_\bullet$, i.e. also contained in the thick subcategory generated by $\mathbb{1}$. □

Theorem 12.10. *Let $S \to R$ be a map of connective commutative ring spectra. Then:*

1. *If $\pi_0(S) \to \pi_0(R)$ is an isomorphism, the canonical functor*

$$\mathrm{Mod}_S(\mathrm{Sp}) \to \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$$

   is an equivalence.

2. *If $\pi_0(S) \to \pi_0(R)$ is surjective, $\pi_0(S)/p \to \pi_0(R)/p$ an isomorphism, and $\pi_0(R)$ is bounded p-torsion, then there is an equivalence*

$$\mathrm{Mod}_{S^\wedge_p}(\mathrm{Sp}) \to \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$$

Proof. We first prove the first statement. Under the equivalence between $\mathrm{Cat}^{\mathrm{perf}}_\infty$ and Pr^L_ω, this corresponds to the claim that

$$\mathrm{Perf}(S) \simeq \lim_{\Delta} \mathrm{Perf}(R^{\otimes_S \bullet+1})$$

is an equivalence. The functor is "fully faithful on $\mathbb{1}$", i.e. that

$$\mathrm{map}_S(S, S) \to \lim_{\Delta} \mathrm{map}_R(R, R)$$

is an equivalence, which follows from Theorem 8.3 applied to $M = S$. Since $\mathrm{Perf}(S)$ is generated by S as thick subcategory, this immediately implies that the functor is fully faithful in general. The essential image of a fully faithful exact functor is then closed under cofibers and fibers, and in this case also under retracts since $\mathrm{Perf}(S)$ is idempotent-complete. So the essential image is a thick subcategory, and since the target category is generated by $\mathbb{1}$ as thick subcategory, the functor is also essentially surjective.

For the second statement, the same argument (using Theorem 8.10 for the fully faithfulness) provides instead an equivalence between the limit and $\mathrm{Thick}(S^\wedge_p) \subseteq \mathrm{Mod}_S(\mathrm{Sp})$, the thick subcategory generated by the p-completion of $S^\wedge_p$. We now observe that the forgetful functor

$$\mathrm{Perf}(S^\wedge_p) \to \mathrm{Mod}_S(\mathrm{Sp})$$

is fully faithful with essential image $\mathrm{Thick}(S^\wedge_p)$, again by a similar argument. This finishes the proof. □

12.2 Filtered and graded spectra

We have seen filtered spectra before: A filtered spectrum is just a diagram

$$\ldots \to X_{n+1} \to X_n \to \ldots$$

of spectra. Formally, the category FilSp of filtered spectra is just given by

$$\mathrm{FilSp} = \mathrm{Fun}(\mathbb{Z}^{\mathrm{op}}, \mathrm{Sp}),$$

where we view $\mathbb{Z}$ as category with the usual partial ordering. The *underlying spectrum* is just $\operatorname{colim}_n X_n$, this gives a functor

$$\operatorname{Fun}(\mathbb{Z}^{\mathrm{op}}, \mathrm{Sp}) \to \mathrm{Sp}\,.$$

We also saw that a spectrum has an associated graded $\operatorname{gr}_* X$, with $\operatorname{gr}_n X = \operatorname{cofib}(X_{n+1} \to X_n)$. This can be interpreted as functor $\mathrm{FilSp} \to \operatorname{gr}\mathrm{Sp}$ with $\operatorname{gr}\mathrm{Sp}$ the category of graded spectra, since a graded spectrum is just a $\mathbb{Z}$-indexed family of spectra, this can be defined as

$$\operatorname{gr}\mathrm{Sp} = \operatorname{Fun}(\mathbb{Z}^{\delta}, \mathrm{Sp}),$$

where $\mathbb{Z}^{\delta}$ denotes the "discrete" category whose objects are integers, with no nonidentity morphisms between them.

FilSp and $\operatorname{gr}\mathrm{Sp}$ have symmetric monoidal structures (called Day convolution), in $\operatorname{gr}\mathrm{Sp}$ this takes the familiar form $(X \otimes Y)_n \simeq \oplus_{i+j=n} X_i \otimes Y_j$, while in FilSp the description is more complicated: $(X \otimes Y)_n$ admits a description as colimit of a diagram consisting of all $X_i \otimes Y_j$ with $i + j \geq n$. By formal properties of Day convolution, both the colimit functor $\operatorname{Fun}(\mathbb{Z}^{\mathrm{op}}, \mathrm{Sp}) \to \mathrm{Sp}$ and the associated graded functor $\operatorname{Fun}(\mathbb{Z}^{\mathrm{op}}, \mathrm{Sp}) \to \operatorname{Fun}(\mathbb{Z}^{\delta}, \mathrm{Sp})$ turn out to be symmetric-monoidal.

In FilSp, we write $X(i)$ for the degree shift of an object, $X(i)_n = X_{n-i}$. If X is an ordinary spectrum, we will also use the notation $X(i)$ to denote the filtered spectrum with

$$X(i)_n = \begin{cases} X \text{ if } n \leq i, \\ 0 \text{ otherwise} \end{cases}$$

and identity structure maps. With this notation, the unit of the monoidal structure is $\mathbb{S}(0)$, so $\mathbb{S}(0) \otimes X \simeq X$, and more generally we have $\mathbb{S}(i) \otimes X \simeq X(i)$, so the degree shift can also be described as tensoring with the objects $\mathbb{S}(i)$. A useful property of the $\mathbb{S}(i)$ is that

$$\operatorname{map}_{\mathrm{FilSp}}(\mathbb{S}(i), X) \simeq X_i.$$

We also write $\mathbb{S}^{n,i} := \Sigma^n \mathbb{S}(i)$, and define

$$\pi_{n,i}(X) \cong [\mathbb{S}^{n,i}, X] \cong [\mathbb{S}^n, X_i] \cong \pi_n(X_i).$$

We have a canonical map $\tau : \mathbb{S}(-1) \to \mathbb{S}(0)$. Tensoring an arbitrary filtered spectrum X with this, we get the natural map

$$X(-1) \to X$$

which is just given by the structure maps $X_{n+1} \to X_n$ levelwise. Writing $\mathbb{S}/\tau$ for the cofiber of τ, we therefore see that $\mathbb{S}/\tau \otimes X$ is a filtered spectrum which in degree n is just given by $\operatorname{gr}_n X$, so we can think of the associated graded construction as a combination of tensoring with $\mathbb{S}/\tau$, and then restricting along $\mathbb{Z}^{\delta} \to \mathbb{Z}^{\mathrm{op}}$ (i.e. forgetting the structure maps).

For example for obstruction-theoretic reasons (using that $(\mathbb{S}/\tau)_n$ is zero for $n \neq 0$), one actually has that $\mathbb{S}/\tau$ has a commutative algebra structure in FilSp. The associated graded functor therefore factors as

$$\mathrm{FilSp} \xrightarrow{\mathbb{S}/\tau \otimes -} \mathrm{Mod}_{\mathbb{S}/\tau}(\mathrm{FilSp}) \to \mathrm{gr}\,\mathrm{Sp},$$

where the second functor just passes to the underlying filtered spectrum and forgets the structure maps. In fact, this second functor is an equivalence: One checks that $\mathrm{Mod}_{\mathbb{S}/\tau}(\mathrm{FilSp})$ is generated by the objects $\mathbb{S}/\tau(i)$ under desuspensions and colimits, and $\mathrm{gr}\,\mathrm{Sp}$ analogously by their images (which are the graded spectra which have $\mathbb{S}$ in degree i and 0 in all other degrees). So to see that the functor is an equivalence, it suffices to check that it is "fully faithful on those objects", which amounts to

$$\mathrm{map}_{\mathrm{Mod}_{\mathbb{S}/\tau}}(\mathbb{S}/\tau(i), \mathbb{S}/\tau(j)) \simeq \begin{cases} \mathbb{S} \text{ if } i = j \\ 0 \text{ otherwise} \end{cases}$$

which one sees again from the fact that $\mathbb{S}/\tau$ is $\mathbb{S}$ in degree 0 and 0 in all other degrees.

So we can identify $\mathrm{gr}\,\mathrm{Sp}$ with $\mathrm{Mod}_{\mathbb{S}/\tau}(\mathrm{FilSp})$, and the associated graded functor as the "induction" functor $\mathrm{FilSp} \to \mathrm{Mod}_{\mathbb{S}/\tau}(\mathrm{FilSp})$.

Similarly, we can also form a filtered spectrum

$$\mathbb{S}[\tau^{-1}] := \mathrm{colim}(\mathbb{S}(0) \xrightarrow{\tau} \mathbb{S}(1) \xrightarrow{\tau} \mathbb{S}(2) \xrightarrow{\tau} \ldots),$$

which is of course just the constant diagram with all entries $\mathbb{S}$. Tensoring an arbitrary filtered spectrum with $\mathbb{S}[\tau^{-1}]$, we see (by commuting out the colimit) that $\mathbb{S}[\tau^{-1}] \otimes X$ is the constant diagram with all entries $\mathrm{colim}_n X_n$. We see from this that $\mathbb{S}[\tau^{-1}] \otimes \mathbb{S}[\tau^{-1}] \simeq \mathbb{S}[\tau^{-1}]$, i.e. $\mathbb{S}[\tau^{-1}]$ is an idempotent algebra in FilSp. This means that a filtered spectrum X admits a (unique) $\mathbb{S}[\tau^{-1}]$ module structure if and only if $X \to \mathbb{S}[\tau^{-1}] \otimes X$ is an equivalence, i.e. if X is already a constant diagram, and hence that

$$\mathrm{Mod}_{\mathbb{S}[\tau^{-1}]}(\mathrm{FilSp}) \simeq \mathrm{Sp}\,.$$

The functor $\mathrm{FilSp} \to \mathrm{Mod}_{\mathbb{S}[\tau^{-1}]}(\mathrm{FilSp})$ given by tensoring with $\mathbb{S}[\tau^{-1}]$ corresponds under this equivalence to the functor which takes a filtered spectrum to its colimit.

In total, we observe that we can identify the diagram of symmetric-monoidal categories

$$\mathrm{Sp} \leftarrow \mathrm{FilSp} \to \mathrm{gr}\,\mathrm{Sp}$$

with the diagram

$$\mathrm{Mod}_{\mathbb{S}[\tau^{-1}]}(\mathrm{FilSp}) \leftarrow \mathrm{FilSp} \to \mathrm{Mod}_{\mathbb{S}/\tau}(\mathrm{FilSp}),$$

where both functors are just induction.

Example 12.11. For an ordinary spectrum X, we have the filtered spectrum $\tau_{\geq *}X$ (the Whitehead tower). Here,

$$\mathbb{S}[\tau^{-1}] \otimes (\tau_{\geq *}X)$$

is the constant filtered spectrum X, and hence corresponds just to $X \in \mathrm{Sp}$, whereas

$$\mathbb{S}/\tau \otimes (\tau_{\geq *}X)$$

corresponds to the graded spectrum which is $\Sigma^n \pi_n(X)$ in degree n.

Example 12.12. For a an algebra object R in filtered spectra, we can form $R[\tau^{-1}] = \mathbb{S}[\tau^{-1}] \otimes R$ and $R/\tau = \mathbb{S}/\tau \otimes R$, which inherit algebra structures in the respective categories by symmetric monoidality. We get an R-linear version of the above diagram of categories as

$$\mathrm{Mod}_{R[\tau^{-1}]}(\mathrm{FilSp}) \leftarrow \mathrm{Mod}_R(\mathrm{FilSp}) \rightarrow \mathrm{Mod}_{R/\tau}(\mathrm{FilSp}).$$

We can write the left category also as $R[\tau^{-1}]$-modules in $\mathbb{S}[\tau^{-1}]$-modules, and so this identifies with modules over $\mathrm{colim}_n R_n$ in Sp. Similarly, the right category can be identified with $\mathrm{gr}_* R$-modules in $\mathrm{gr}\,\mathrm{Sp}$.

Example 12.13. Specialising the previous example to the Whitehead tower $\tau_{\geq *}R$ of a ring spectrum R, $\mathrm{Mod}_{\tau_{\geq *}R[\tau^{-1}]}(\mathrm{FilSp})$ is equivalent to with $\mathrm{Mod}_R(\mathrm{Sp})$, while $\mathrm{Mod}_{R/\tau}(\mathrm{FilSp})$ is equivalent to $\mathrm{Mod}_{\Sigma^*\pi_*(R)}(\mathrm{gr}\,\mathrm{Sp})$. This latter category is equivalent to $\mathcal{D}^{\mathrm{gr}}(\pi_*(R))$. In fact, here we get a better explanation for some of the choices we made when we first discussed $\mathcal{D}^{\mathrm{gr}}(\pi_*(R))$. Indeed, $\mathrm{gr}\,\mathrm{Sp}$ admits a t-structure where X is (co)connective if for each n, $\Sigma^{-n}X_n$ is (co)connective. This is designed in such a way that $\Sigma^*\pi_*(R)$ is in the heart, and can be used to identify $\mathrm{Mod}_{\Sigma^*\pi_*(R)}$ with $\mathcal{D}^{\mathrm{gr}}(\pi_*(R))$. The heart is equivalent to the category $\mathrm{gr}\,\mathrm{Ab}$ of graded abelian groups, and interestingly the symmetric-monoidal structure induced on the heart is automatically the one with the Koszul sign. There are other t-structures which differ from this one by some sort of shearing equivalence, for example one has the pointwise t-structure in which X is (co)connective if and only if for all n, X_n is (co)connective. In that t-structure, the heart is also $\mathrm{gr}\,\mathrm{Ab}$, but here the induced symmetric-monoidal structure is the one with the trivial sign. (A related observation is that the shearing equivalence $\mathrm{gr}\,\mathrm{Sp} \to \mathrm{gr}\,\mathrm{Sp}$ which takes $X_* \mapsto \Sigma^{-*}X_*$, is monoidal, but not symmetric monoidal.)

In summary, the most natural explanation for the Koszul sign appearing when defining the graded rings $\pi_*(R)$ and their module categories is that they appear canonically as the objects $\Sigma^*\pi_*(R)$ in $\mathrm{gr}\,\mathrm{Sp}$, and the corresponding embedding of $\mathrm{gr}\,\mathrm{Ab}$ into $\mathrm{gr}\,\mathrm{Sp}$ leads to the Koszul sign symmetric monoidal structure on $\mathrm{gr}\,\mathrm{Ab}$.

12.3 Synthetic spectra

We have seen that for the individual filtered ring spectrum $\tau_{\geq *}R$, $\mathrm{Mod}_{\tau_{\geq *}R}(\mathrm{FilSp})$ is a category which somehow interpolates between $\mathrm{Mod}_R(\mathrm{Sp})$ and $\mathrm{Mod}_{\Sigma^*\pi_*(R)}(\mathrm{gr}\,\mathrm{Sp}) \simeq$

$\mathcal{D}^{\mathrm{gr}}(\pi_*(R))$. The underlying object of such a $\tau_{\geq *}R$ module is a filtered spectrum, but its colimit comes with an R-module structure, and its associated graded can be interpreted as an object of $\mathcal{D}^{\mathrm{gr}_*}(\pi_*(R))$. One way to construct such an object is by starting with an R-module spectrum M and passing to its Whitehead tower $\tau_{\geq *}M$, but there are more interesting objects as well. For example, for commutative R, $\mathrm{Mod}_{\tau_{\geq *}R}(\mathrm{FilSp})$ is a closed symmetric-monoidal category, i.e. it has a symmetric-monoidal structure and an internal Hom characterized as right adjoint to tensoring. For R-module spectra M and N, we can then construct an object of the form

$$(\tau_{\geq *}M) \otimes_{\tau_{\geq *}R} (\tau_{\geq *}N)$$

which is an interesting filtered spectrum with underlying object $M \otimes_R N$, and associated graded $\pi_*(M) \otimes^L_{\pi_*(R)} \pi_*(N)$. This object leads to another construction of the "Tor spectral sequence" mentioned in Theorem 7.8. Analogously, one can use the internal Hom to construct the "Ext spectral sequence" from Theorem 7.7.

The idea behind synthetic spectra is to provide a similar categorification of Adams spectral sequences. This will be a category which interpolates between Sp (or p-complete spectra) and some derived category of comodules. We will develop this in the same generality as we have discussed Adams spectral sequences, for general maps $S \to R$ of connective commutative ring spectra (plus additional π_0 assumptions), before specializing to the main examples $\mathbb{S} \to \mathbb{F}_p$ and $\mathbb{S} \to MU$.

Definition 12.14. Let $R^\bullet$ be a cosimplicial filtered commutative ring spectrum, that is a functor $\Delta \to \mathrm{CAlg}(\mathrm{FilSp})$. We say $R^\bullet$ is *vertically connective* if, for each n, the filtered spectrum

$$\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)$$

is pointwise connective.

The name is justified by the fact that in the "Adams grading" way of displaying $\pi_{*,*}(\lim_\Delta R^\bullet)$, where one puts $\pi_{n,w}$ in bidegree (n,s) with $s = w - n$, this condition implies that $\pi_{*,*}(\lim_\Delta R^\bullet)$ is concentrated to the right of the vertical line $n = 0$.

Lemma 12.15. *If $R^\bullet$ is vertically connective, then in particular each R^n is (pointwise) connective.*

Proof. The fiber

$$\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)$$

has an interpretation as Ω^n of the total fiber of an n-dimensional cube-shaped diagram involving the degeneracy maps of the cosimplicial diagram (this is the same ingredient as what goes into Propositions 8.1 and 12.1, which we don't make explicit here). A consequence is that

$$\pi_{k,w}(\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet))$$

can be identified with the joint kernel of the degeneracy maps $\pi_{k+n,w}(R^n) \to \pi_{k+n,w}(R^{n-1})$. This is the degree n part of the cochain complex corresponding to $\pi_{k+n,w}(R^\bullet)$ under Dold-Kan. The vertical connectivity condition implies that this vanishes for $k < 0$. In particular, the cochain complex of graded abelian groups corresponding to $\pi_{*,w}(R^\bullet)$ under Dold-Kan vanishes in negative degrees. Since the inverse of the Dold-Kan correspondence takes a cochain complex C^* to a cosimplicial object which is levelwise just a direct sum of terms of the cochain complex, this shows that also $\pi_{k,w}(R^n)$ vanishes whenever $k < 0$. □

Lemma 12.16. *Let $S \to R$ be a map of connective commutative ring spectra which is surjective on π_0. Then the cosimplicial filtered spectrum obtained from the Cech nerve by applying $\tau_{\geq *} : \mathrm{Sp} \to \mathrm{FilSp}$,*

$$R^\bullet := \tau_{\geq *}(R^{\otimes_S \bullet+1}),$$

is vertically connective.

Proof. Again, we can interpret

$$\pi_{k,w}(\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet))$$

as the degree n part of the cochain complex corresponding to $\pi_{k+n,w}(R^\bullet)$ under Dold-Kan. If $R^s = \tau_{\geq *}(R^{\otimes_S s+1})$, then of course

$$\pi_{k+n,w}(R^s) \cong \pi_{k+n}\tau_{\geq w}(R^s) \cong \begin{cases} \pi_{k+n}(R^s) \text{ if } k+n \geq w \\ 0 \text{ otherwise} \end{cases}$$

and so we get

$$\begin{aligned} &\pi_{k,w}(\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet)) \\ \cong &\begin{cases} \pi_k(\mathrm{fib}(\lim_{\Delta_{\leq n}} R^{\otimes_S \bullet+1} \to \lim_{\Delta_{\leq n-1}} R^{\otimes_S \bullet+1})) \text{ if } k+n \geq w \\ 0 \text{ otherwise} \end{cases} \end{aligned}$$

since the latter thing can also be identified with the degree n-part of the cochain complex corresponding to $\pi_{k+n}(R^{\otimes_S \bullet})$ under Dold-Kan. Now we know this to be connective by Proposition 8.1, and so the right hand side is only nonzero if $k \geq 0$, finishing the claim. □

Remark 12.17. In fact, there is an ∞-categorical Dold-Kan equivalence (discussed for example in [21]), which takes a cosimplicial object $X_\bullet$ in an additive idempotent-complete ∞-category to a coherent cochain complex $\mathrm{DK}(X_\bullet)$. Just as in the usual Dold-Kan, the inverse takes a coherent cochain complex C^* to a cosimplicial object which in degree n is given as $\bigoplus_{[n]\twoheadrightarrow[m]} C^m$, in particular $\mathrm{DK}(X_\bullet)_n$ is a retract of X_n. For cosimplicial diagrams in a stable ∞-category, one has

$$\mathrm{fib}(\lim_{\Delta_{\le n}} X_\bullet \to \lim_{\Delta_{\le n-1}} X_\bullet) \simeq \Omega^n \mathrm{DK}(X_\bullet)_n,$$

and the above discussion can be explained more homotopy-coherently by observing that $\tau_{\ge *} : \mathrm{Sp} \to \mathrm{FilSp}$ is additive, so it commutes with the Dold-Kan equivalence, so

$$\mathrm{DK}(R^\bullet)_n \simeq \tau_{\ge *} \mathrm{DK}(R^{\otimes_S \bullet+1})_n.$$

Then of course knowing that the right-hand side is n-connective before applying $\tau_{\ge *}$ shows that it gives a pointwise n-connective filtered spectrum.

Proposition 12.18. *If $R^\bullet$ is vertically connective, then*

$$\lim_{\Delta} \mathrm{Perf}(R^\bullet)$$

is generated as thick subcategory by $\mathbb{1}(w)$ for all w, where $\mathrm{Perf}(R^n) \subseteq \mathrm{Mod}_{R^n}(\mathrm{FilSp})$ denotes the full subcategory on compact objects.

Proof. Since R^n is pointwise connective, a similar argument as for Proposition 12.6 can be used to check that every compact R^n-module M admits a finite filtration where $\mathrm{gr}_{-k} M$ is a finite sum of objects of the form $\Sigma^k R^n(w)$, for varying w. We may think of this as some kind of cell structure where cells have, in addition to dimension k, also weight w. If M is pointwise k-connective, one can see from this that there exists a map

$$\bigoplus \Sigma^k R^n(w_i) \to M$$

surjective on $\pi_{k,w}$ for all w.

If $M_\bullet \in \lim_\Delta \mathrm{Perf}(R^\bullet)$, with M_0 pointwise k-connective, then by a filtered spectrum version of Proposition 12.1, the map

$$[\Sigma^k \mathbb{1}(w), M_\bullet] \to [\Sigma^k R^0(w), M_0]$$

is surjective. In particular we find a map $\bigoplus \Sigma^k \mathbb{1}(w_i) \to M_\bullet$ such that the underlying map in cosimplicial degree 0 is surjective on $\pi_{k,w}$ for all w. Its cofiber is then in cosimplicial degree 0 pointwise $k+1$-connective. As in the proof of Proposition 12.9, we can use this to produce a cofiber sequence

$$K_\bullet \to M_\bullet \to M'_\bullet$$

with $K_\bullet$ in the thick subcategory generated by all $\mathbb{1}(w)$, and $M'_\bullet$ is pointwise $b+2$-connective, with b such that M_0 is a retract of an object with a finite filtration by $\Sigma^i R^0(w)$ with $i \le b$. Then the filtered analogue of Proposition 12.1 shows that the mapping spectrum for $M_\bullet \to M'_\bullet$ is 1-connective, so that map is nullhomotopic and $M_\bullet$ a retract of $K_\bullet$. $\square$

Theorem 12.19. *If $R^\bullet$ is a cosimplicial filtered commutative ring spectrum which is vertically connective, then*

$$\mathrm{Mod}_{\lim_\Delta R^\bullet}(\mathrm{FilSp}) \simeq \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^\bullet}(\mathrm{FilSp}).$$

Proof. This amounts again to checking that

$$\mathrm{Perf}(\lim_\Delta R^\bullet) \simeq \lim_\Delta \mathrm{Perf}(R^\bullet).$$

By construction, the functor gives an equivalence on $\mathrm{map}(\mathbb{1}(i), \mathbb{1}(j))$: On both sides, this just evaluates to the degree $(i-j)$-part of $\lim_\Delta R^\bullet$. Now, the source is generated by the $\mathbb{1}(i)$ as thick subcategory, so the functor in general is fully faithful. Since the target is generated by $\mathbb{1}(i)$ as a thick subcategory and the source is idempotent complete, it then follows that the functor is also essentially surjective, hence an equivalence. $\square$

We now specialize to the case where

$$R^\bullet := \tau_{\geq *} R^{\otimes_S \bullet+1}$$

is the Cech nerve of a map of connective commutative ring spectra $S \to R$ which is surjective on π_0. Then $R^\bullet/\tau = R^\bullet \otimes \mathbb{S}/\tau$ is an object of $\mathrm{Fun}(\Delta, \mathrm{CAlg}_{\mathbb{S}/\tau}(\mathrm{FilSp})) \simeq \mathrm{Fun}(\Delta, \mathrm{CAlg}(\mathrm{gr\,Sp}))$, and corresponds to the cosimplicial graded ring spectrum $\bigoplus \Sigma^* \pi_*(R^{\otimes_S \bullet+1})$. As

$$\mathrm{Mod}_{\bigoplus \Sigma^* \pi_*(R^{\otimes_S n+1})}(\mathrm{gr\,Sp}) \simeq \mathcal{D}^{\mathrm{gr}}(\pi_*(R^{\otimes_S n+1})),$$

we have

$$\lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^\bullet/\tau}(\mathrm{FilSp}) \simeq \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathcal{D}^{\mathrm{gr}}(\pi_*(R^{\otimes_S n+1})).$$

If $\pi_*(R \otimes_S R)$ is a flat $\pi_*(R)$-module, one has a similar Künneth result as we already saw in the case $S = \mathbb{S}$, $R = \mathbb{F}_p$, where

$$\pi_*(R^{\otimes_S n+1}) \cong \pi_*(R \otimes_S R) \otimes_{\pi_*(R)} \cdots \otimes_{\pi_*(R)} \pi_*(R \otimes_S R).$$

Here, one has to be careful, as the two maps $\pi_*(R) \to \pi_*(R \otimes_S R)$ are not necessarily the same anymore.

Under the flatness assumption, we also have that

$$\lim_\Delta \mathcal{D}^{\mathrm{gr}}(\pi_*(R^{\otimes_S n+1})).$$

carries a "pointwise" t-structure, with heart

$$\lim_\Delta \mathrm{Mod}_{\pi_*(R^{\otimes_S n+1})}.$$

Under Künneth, we may interpret objects in there as a type of comodules, where the structure map

$$M \to \pi_*(R \otimes_S R) \otimes_{\pi_*(R)} M$$

is $\pi_*(R)$-linear with respect to the left module structure, and the tensor product is formed using the right module structure. This can be made precise in the language of *Hopf algebroids* and their comodules. In analogy with $\mathcal{D}(\mathrm{Comod}_{\mathcal{A}_*})$ and $\mathrm{Comod}_{\mathcal{A}_*}$, we can write

$$\mathcal{D}(\mathrm{Comod}_{\pi_*(R\otimes_S R)}) := \lim_{\Delta} \mathcal{D}^{\mathrm{gr}}(\pi_*(R^{\otimes_S n+1})),$$

and $\mathrm{Comod}_{\pi_*(R\otimes_S R)}$ for its heart.

Definition 12.20.

$$\mathcal{D}_\omega(\mathrm{Comod}_{\pi_*(R\otimes_S R)}) := \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathcal{D}^{\mathrm{gr}}(\pi_*(R^{\otimes_S n+1})).$$

Explicitly, this is a compactly generated stable ∞-category, whose compact objects are given by

$$\lim_{\Delta} \mathrm{Perf}^{\mathrm{gr}}(\pi_*(R^{\otimes_S n+1})),$$

the full subcategory of $\mathcal{D}(\mathrm{Comod}_{\pi_*(R\otimes_S R)})$ on pointwise compact objects.

Remark 12.21. For $S = \mathbb{S}$, this is what's called $\mathrm{Stable}(R_*R)$ in parts of the literature. It is the Ind-category of the full subcategory of $\mathcal{D}(\mathrm{Comod}_{\pi_*(R\otimes_S R)})$ on those objects whose underlying object in $\mathcal{D}^{\mathrm{gr}}(\pi_*(R))$ is compact (i.e. in $\mathrm{Perf}(\pi_*(R))$).

Definition 12.22. For a map $S \to R$ of connective commutative ring spectra which is surjective on π_0 and satisfies that $\pi_*(R \otimes_S R)$ is flat as $\pi_*(R)$-module, we define $R^\bullet := \tau_{\geq *} R^{\otimes_S \bullet+1}$ as the Whitehead tower of the Cech nerve. With that, we set

$$\mathrm{Syn}_{R/S} := \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^\bullet}(\mathrm{FilSp}) \simeq \mathrm{Mod}_{\lim_\Delta R^\bullet}(\mathrm{FilSp}).$$

This $\mathrm{Syn}_{R/S}$ has a symmetric-monoidal structure, and a corresponding unit $\mathbb{1}$. More generally, there are filtration shifts $\mathbb{1}(i)$ (base-changed from $\mathbb{S}(i)$, and given pointwise by $R^n(i)$), and we can form bigraded "homotopy groups"

$$\pi_{n,w}(X) := [\Sigma^n \mathbb{1}(w), X]$$

for any $X \in \mathrm{Syn}_{R/S}$.

We have a map $\tau : \mathbb{1}(-1) \to \mathbb{1}$, base-changed from the corresponding map $\mathbb{S}(-1) \to \mathbb{S}$. There is a commutative algebra structure on $\mathbb{1}/\tau$ (for example because it can be written as $\mathbb{S}/\tau \otimes \lim_\Delta R^\bullet$ in $\mathrm{Mod}_{\lim_\Delta R^\bullet}(\mathrm{FilSp})$.

Theorem 12.23. *For a map $S \to R$ of connective commutative ring spectra which is surjective on π_0 and satisfies that $\pi_*(R \otimes_S R)$ is flat as $\pi_*(R)$-module, we have*

$$\mathrm{Mod}_{\mathbb{1}/\tau}(\mathrm{Syn}_{R/S}) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{\pi_*(R\otimes_S R)})$$

and

1. *If $\pi_0(S) \to \pi_0(R)$ is an isomorphism,* $\mathrm{Mod}_{\mathbb{1}[\tau^{-1}]}(\mathrm{Syn}_{R/S}) \simeq \mathrm{Mod}_S(\mathrm{Sp})$.
2. *If $\pi_0(S)/p \to \pi_0(R)/p$ is an isomorphism and p is nilpotent in $\pi_0(R)$,* $\mathrm{Mod}_{\mathbb{1}[\tau^{-1}]}(\mathrm{Syn}_{R/S}) \simeq \mathrm{Mod}_{S^\wedge_p}(\mathrm{Sp})$.

Proof. $R^\bullet = \tau_{\geq *} R^{\otimes_S \bullet+1}$ is vertically connective, and so is $R^\bullet[\tau^{-1}]$ and $R^\bullet/\tau$: Both the colimit involved in inverting τ and the cofiber of τ commute with the finite limits involved in forming

$$\mathrm{fib}(\lim_{\Delta_{\leq n}} R^\bullet \to \lim_{\Delta_{\leq n-1}} R^\bullet).$$

This means that $\mathrm{Syn}_{R/S} \simeq \mathrm{Mod}_{\lim_\Delta R^\bullet}(\mathrm{FilSp})$, and

$$\mathrm{Mod}_{\mathbb{1}/\tau}(\mathrm{Syn}_{R/S}) \simeq \mathrm{Mod}_{\lim_\Delta R^\bullet/\tau} \simeq \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^\bullet/\tau}(\mathrm{FilSp})$$

and

$$\mathrm{Mod}_{\mathbb{1}[\tau^{-1}]}(\mathrm{Syn}_{R/S}) \simeq \mathrm{Mod}_{\lim_\Delta R^\bullet[\tau^{-1}]} \simeq \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^\bullet[\tau^{-1}]}(\mathrm{FilSp})$$

The first limit identifies with

$$\lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathcal{D}^{\mathrm{gr}}(\pi_* R^\bullet) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{\pi_*(R \otimes_S R)})$$

and the second with

$$\lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{R^{\otimes_S \bullet+1}}(\mathrm{Sp})$$

which by Theorem 12.10 has the claimed form. □

For any S-module M, we can produce an object in $\mathrm{Syn}_{R/S}$ by forming the object

$$(\tau_{\geq *} R^{\otimes_S \bullet} \otimes_S M).$$

This defines a functor

$$\nu : \mathrm{Mod}_S \to \mathrm{Syn}_{R/S},$$

the *synthetic analogue* functor. It takes $\mathbb{S} \mapsto \mathbb{1}$, and $\mathbb{S}^n \mapsto \Sigma^n \mathbb{1}(n)$. It is not exact (since $\tau_{\geq *}$ isn't), but it does preserve those cofiber sequences

$$M_0 \to M_1 \to M_2$$

which induce short exact sequences on $\pi_*(R \otimes_S -)$, by flatness of $\pi_*(R \otimes_S R)$ and the fact that $\tau_{\geq *}$ preserves cofiber sequences which are short exact sequences on homotopy.

12.4 $\mathbb{F}_p$-synthetic spectra

We now specialize to the case $S = \mathbb{S}$. In that case, we write Syn_R for $\mathrm{Syn}_{R/\mathbb{S}}$. For $R = \mathbb{F}_p$, we get a category $\mathrm{Syn}_{\mathbb{F}_p}$ with:

$$\mathrm{Mod}_{\mathbb{1}[\tau^{-1}]}(\mathrm{Syn}_{\mathbb{F}_p}) \simeq \mathrm{Mod}_{\mathbb{S}_p^\wedge}(\mathrm{Sp})$$
$$\mathrm{Mod}_{\mathbb{1}/\tau}(\mathrm{Syn}_{\mathbb{F}_p}) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$$

We will now also write $\mathbb{S}_p^{0,0}$ for $\mathbb{1}$ and $\mathbb{S}_p^{n,w}$ for $\Sigma^n \mathbb{1}(w)$. Similarly, we will write $\Sigma^{n,w} X$ for $\Sigma^n X(w)$.

The equivalence between $\mathrm{Mod}_{\mathbb{1}/\tau}(\mathrm{Syn}_{\mathbb{F}_p}) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ involves a shift, arising from the translation between modules over the graded ring $\Sigma^* \pi_*(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet+1})$ and $\pi_*(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet+1})$. As a reality check, observe that

$$\mathrm{map}_{\mathrm{Syn}_{\mathbb{F}_p}}(\mathbb{S}_p^{0,0}(w), \mathbb{S}_p^{0,0}/\tau) \simeq \lim_{\Delta} \Sigma^w \pi_w(\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet+1}),$$

a shift of the limit which we previously identified with the cobar complex of $\mathcal{A}_*$. So its π_n can be identified with $H^{-(n-w)}$ of the weight w part of the cobar complex, thus

$$\pi_{n,w}(\mathbb{S}_p^{0,0}/\tau) \cong \mathrm{Ext}_{\mathcal{A}_*}^{w-n,w}(\mathbb{F}_p, \mathbb{F}_p).$$

If we want to relate $\mathrm{Mod}_{\mathbb{S}_p^{0,0}/\tau}(\mathrm{Syn}_{\mathbb{F}_p})$ to the Adams spectral sequence, $\pi_{n,w}$ should therefore be placed at coordinates $(t-s, s)$ with $s = w-n$ and $t = w$, so at coordinates $(n, w-n)$. Said more succinctly, n is the horizontal coordinate, but w is constant along codiagonals. For example, the element h_1 detecting η in the $\mathbb{F}_2$-based Adams spectral sequence corresponds to an element of $\pi_{1,2}(\mathbb{S}_p^{0,0}/\tau)$.

For any object $X \in \mathrm{Syn}_{\mathbb{F}_p}$, $X[\tau^{-1}]$ corresponds to an object of $\mathrm{Mod}_{\mathbb{S}_p^\wedge}(\mathrm{Sp})$, which we will also write by $X[\tau^{-1}]$, and X/τ to an object of $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$. Forgetting structure on X along $\mathrm{Syn}_{\mathbb{F}_p} \to \mathrm{FilSp}$, the resulting filtered spectrum has underlying spectrum given precisely by $X[\tau^{-1}]$, and associated graded given by X/τ. Concretely, this means that

$$\pi_{n,w}(X/\tau) = [\mathbb{S}_p^{n,w}, X/\tau]_{\mathrm{Syn}_{\mathbb{F}_p}} \cong [\Sigma^{n-w}\mathbb{F}_p(w), X/\tau]_{\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})}$$

has some interpretation as $\mathrm{Ext}_{\mathcal{A}_*}$. So X is a filtration on $X[\tau^{-1}]$ which represents a variant of the Adams spectral sequence.

For example, the unit $\mathbb{S}_p^{0,0} \in \mathrm{Syn}_{\mathbb{F}_p}$ itself is the filtered spectrum which gives rise to the Adams spectral sequence, with $\mathbb{S}_p^{0,0}[\tau^{-1}] \simeq \mathbb{S}_p^\wedge$, and

$$\pi_{n,w}(\mathbb{S}^{0,0}/\tau) \cong [\Sigma^{n-w}\mathbb{F}_p(w), \mathbb{F}_p]_{\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})} \cong \mathrm{Ext}_{\mathcal{A}_*}^{w-n,w}(\mathbb{F}_p, \mathbb{F}_p).$$

For $X \in \mathrm{Syn}_{\mathbb{F}_p}$, we should think of $\pi_{*,*}(X/\tau)$ as a generalized version of $\mathrm{Ext}_{\mathcal{A}_*}^{*,*}(\mathbb{F}_2, \mathbb{F}_2)$. How should we think about $\pi_{*,*}(X)$ itself?

We may view

$$\ldots \xrightarrow{\tau} X(-2) \xrightarrow{\tau} X(-1) \xrightarrow{\tau} X \xrightarrow{\mathrm{id}} X \xrightarrow{\mathrm{id}} \ldots$$

as filtration

$$F^{\geq i}X = \begin{cases} X(-i) \text{ if } i \geq 0 \\ X \text{ otherwise.} \end{cases}$$

on X in $\mathrm{Syn}_{\mathbb{F}_p}$ (introducing a new filtration direction). In analogy with the usual Bockstein spectral sequence (which is obtained from a similar filtration using the endomorphism p instead of τ on a spectrum), we call this the "τ-Bockstein filtration".

Even though this is just a formal trick, as the maps τ essentially reuse the existing structure maps of the underlying filtered spectrum of X (under $\mathrm{Syn}_{\mathbb{F}_p} \simeq \mathrm{Mod}_{\lim \tau_{\geq *}\mathbb{F}_p^{\otimes_{\mathbb{S}} \bullet+1}}(\mathrm{FilSp})$), it gives us a useful new perspective on the spectral sequence corresponding to X.

The τ-Bockstein filtration satisfies $\mathrm{gr}_i(F^{\geq *}X) \simeq (X/\tau)(-i)$ for $i \geq 0$. We get a spectral sequence (now with a third grading)

$$\pi_{n,w}(\mathrm{gr}_u F^{\geq *}X) \Rightarrow \pi_{n,w}(X).$$

For this, we don't even have to generalize the construction of spectral sequences from filtered spectra, it suffices to observe that for each w,

$$\mathrm{map}(\mathbb{1}(w), F^{\geq *}X)$$

is a filtered spectrum whose underlying spectrum is $\mathrm{map}(\mathbb{1}(w), X) = X_w$. In fact, it just is the filtered spectrum

$$\ldots \to X_{w+2} \to X_{w+1} \to X_w \to X_w \to \ldots$$

which is X_{w+i} in degree $i \geq 0$ and constant X_w in negative degrees. Each of these gives the weight w part of the above spectral sequence. We also see that the weight w part of the above spectral sequence is just a truncated and shifted version of the spectral sequence of X itself: It is the truncation where the filtration is constant below X_w, so where gr_{w-1} and below are made 0. The map

$$\mathrm{map}(\mathbb{1}(w), F^{\geq *}X) \to \mathrm{map}(\mathbb{1}(w-1), F^{\geq *}X)$$

of filtered spectra obtained by precomposing with τ is just the canonical map between these different truncations of X.

In fact, none of the above requires that we are in $\mathrm{Syn}_{\mathbb{F}_p}$, and works for an arbitrary filtered spectrum. If $X \in \mathrm{Syn}_{\mathbb{F}_p}$, the resulting spectral sequence will however be a module over the spectral sequence for $\mathbb{S}_p^{0,0}$.

Theorem 12.24. *For a filtered spectrum X, the τ-Bockstein spectral sequence*

$$\pi_{n,w}(\mathrm{gr}_u F^{\geq *}X) \Rightarrow \pi_{n,w}(X)$$

has the following description:

1. *The E^1 page is isomorphic to*

$$\pi_{*,*}(X/\tau)[\overline{\tau}]$$

with $\pi_{n,w+u}(X/\tau) \cdot \overline{\tau}^u$ corresponding to $\pi_{n,w}(\mathrm{gr}_u F^{\geq *} X)$.

2. *The d_r differential preserves w, decreases n by 1, and increases u by r.*

3. *Inverting $\overline{\tau}$, the resulting spectral sequence with E^1 page*

$$\pi_{*,*}(X/\tau)[\overline{\tau}^{\pm 1}]$$

is periodic in the w direction, and each constant w slice is isomorphic to the spectral sequence associated to the filtered spectrum X itself,

$$\pi_n(\mathrm{gr}_u X) \Rightarrow \pi_n(X).$$

Here $\overline{\tau}$ can be thought of as a formal element which detects the action of the actual $\tau \in \pi_{0,-1}(\mathbb{S}_p^{0,0})$ on $\pi_{*,*}(X)$ in the spectral sequence.

Say $\alpha \in \pi_{n,w}(X/\tau)$ and $\beta \in \pi_{n-1,w+r}(X/\tau)$ are related by a differential $d_r(\alpha) = \beta$ in the spectral sequence associated to X. Then α corresponds to an element in the E^1 page $\pi_{*,*}(X/\tau)[\overline{\tau}]$ of the τ-Bockstein spectral sequence of X, with u-degree 0. The element β gives rise to an element $\overline{\tau}^r\beta$ in the same weight w, and in the τ-Bockstein spectral sequence we must have $d_r(\alpha) = \overline{\tau}^r\beta$ (since inverting $\overline{\tau}$ must recover the spectral sequence of X).

On the E_∞ page, the elements $\beta, \ldots, \overline{\tau}^{r-1}\beta$ are still nonzero, at least if β was not a boundary on an earlier page of the spectral sequence associated to X. This means that in the τ-Bockstein spectral sequence, boundaries never become fully zero, just $\overline{\tau}$-torsion, with the order relating to the length of the responsible differentials. Thus, in the E_∞ page of the τ-Bockstein spectral sequence of X, the τ-torsion remembers something about the history of the spectral sequence.

In homotopy $\pi_{*,*}(X)$, we may still think of τ-torsion as related to the history of differentials in the spectral sequence. Say $\widetilde{\beta} \in \pi_{n-1,w+r}(X)$ is τ^r-torsion, so that it is in the kernel of

$$\pi_{n-1}(X_{w+r}) \to \pi_{n-1}(X_w).$$

This means that $\tau^{r-1}\widetilde{\beta} \in \pi_{n-1}(X_{w+1})$ is in the kernel of τ, so in the image of the connecting homomorphism $\pi_n(\mathrm{gr}_w X) \to \pi_{n-1}(X_{w+1})$. From the description of differentials in terms of the unrolled exact couple, this means that a preimage $\alpha \in \pi_n(\mathrm{gr}_w X)$ is an r-cycle with $d_r(\alpha) = \beta$, where β is the image of $\widetilde{\beta}$ in $\pi_{n-1}(\mathrm{gr}_{w+r} X)$.

The advantage of "reifying" spectral sequence differentials into information in homotopy groups is that the latter can be studied in different ways. For example, the latter behaves well under cofiber sequences of filtered spectra, whereas there is no corresponding notion of "long exact sequence of spectral sequences".

Another advantage is that we can craft objects in $\mathrm{Syn}_{\mathbb{F}_p}$ whose associated spectral sequence behaves better than the Adams spectral sequence of an object. We illustrate this by revisiting the discussion from Section 11.1.

To start, observe that the element $2 \in \pi_{0,0}(\mathbb{S}_2^{0,0})$ maps to 0 in $\pi_{0,0}(\mathbb{S}_2^{0,0}/\tau)$, as the latter is $\mathrm{Ext}_{\mathcal{A}_*}^{0,0}(\mathbb{F}_2, \mathbb{F}_2) \cong \mathbb{F}_2$. This means that there exists a preimage $\widetilde{2} \in \pi_{0,0}(\mathbb{S}_2^{0,-1})$ with $\tau \cdot \widetilde{2} = 2$. In fact, this preimage is unique, as the leftmost term in the long exact sequence

$$\pi_{1,0}(\mathbb{S}_2^{0,0}/\tau) \to \pi_{0,0}(\mathbb{S}_2^{0,-1}) \to \pi_{0,0}(\mathbb{S}_2^{0,0}) \to \pi_{0,0}(\mathbb{S}_2^{0,0}/\tau)$$

is $\mathrm{Ext}_{\mathcal{A}_*}^{-1,0}(\mathbb{F}_2, \mathbb{F}_2) = 0$.

The cofiber of $\widetilde{2} : \mathbb{S}_2^{0,1} \to \mathbb{S}_2^{0,0}$ does in fact have a more direct description: It is the synthetic analogue of $\mathbb{S}/2$. To see this, observe that

$$\mathbb{S} \to \mathbb{S}/2 \to \mathbb{S}^1$$

is a cofiber sequence which gives a short exact sequence on $\pi_*(\mathbb{F}_2 \otimes_{\mathbb{S}} -)$, so

$$\nu(\mathbb{S}) \to \nu(\mathbb{S}/2) \to \nu(\mathbb{S}^1)$$

is a cofiber sequence in $\mathrm{Syn}_{\mathbb{F}_2}$. Since $\nu(\mathbb{S}) = \mathbb{S}^{0,0}$ and $\nu(\mathbb{S}^1) = \mathbb{S}^{1,1}$, this rotates to a cofiber sequence

$$\mathbb{S}^{0,1} \to \mathbb{S}^{0,0} \to \nu(\mathbb{S}/2).$$

The left hand map lifts $2 : \mathbb{S} \to \mathbb{S}$, so it is $\widetilde{2}$ by the uniqueness discussed above. So the spectral sequence of $\mathbb{S}_2^{0,0}/\widetilde{2} \simeq \nu(\mathbb{S}/2)$ is just the usual Adams spectral sequence of $\mathbb{S}/2$.

If we had instead considered the cofiber $\mathbb{S}_2^{0,0}/2$ of $2 = \tau \cdot \widetilde{2}$, its mod τ reduction would split as $\mathbb{S}_2^{0,0}/\tau \oplus \mathbb{S}_2^{1,0}/\tau$. So the first page of the spectral sequence associated to $\mathbb{S}_2^{0,0}/2$ actually looks like the sum of two copies of $\mathrm{Ext}_{\mathcal{A}_*}(\mathbb{F}_2, \mathbb{F}_2)$.

In usual Adams grading, one of them starts in $(t-s, s)$-degree $(0, 0)$, the other in $(1, -1)$. The identity $2 = 0$ in $\pi_*(\mathbb{S}/2)$ here is implemented by a differential, from $(t-s, s)$-degree $(1, -1)$ to $(0, 1)$. In terms of the filtered spectrum, this is a d_1 differential, but in usual Adams grading, the pages are reindexed so that it is a d_2 differential. This means that instead of the element detecting 2 in the Adams spectral sequence for $\mathbb{S}/2$ being zero algebraically on the first page already, the effect is "delayed" to a differential. If we took the cofiber of $\tau^{r-1} \cdot 2 : \mathbb{S}^{0,-r+1} \to \mathbb{S}^{0,0}$ instead, the same effect would be delayed to a later differential (d_{r+1} in Adams grading), and the shifted copy of $\mathrm{Ext}_{\mathcal{A}_*}(\mathbb{F}_2, \mathbb{F}_2)$ would be moved even lower.

What's more useful is that sometimes we can also sometimes shorten differentials. Recall that there is an element in degree $(8, 4)$ in the Adams spectral sequence represented by a map $x : \mathbb{S}^8 \to \mathbb{S}/2$ which extends over $\mathbb{S}^8/2 \to \mathbb{S}/2$. Writing X for its cofiber, $\nu(X)$ sits in a cofiber sequence

$$\nu(\mathbb{S}/2) \to \nu(X) \to \Sigma^{9,9}\nu(\mathbb{S}/2).$$

The map $\Sigma^{8,9}\nu(\mathbb{S}/2) \to \nu(\mathbb{S}/2)$ is in fact zero mod τ, so the first page of the spectral sequence for $\nu(X)$ splits into two copies of the one for $\nu(\mathbb{S}/2)$. The second one is shifted so that in Adams grading it begins in degree $(9,0)$. The element in Adams degree $(8,4)$ that we used to form X is only killed by a differential which is a d_4 in Adams grading.

The element in degree $(8,4)$ in the Adams E_2 page is represented by a map

$$\mathbb{S}^{8,12} \to \nu(\mathbb{S}/2)/\tau.$$

Using that it is at the top of the Adams spectral sequence, we can see from the τ-Bockstein spectral sequence that this lifts to a map

$$\widetilde{x} : \mathbb{S}^{8,12} \to \nu(\mathbb{S}/2),$$

and that $\widetilde{2} \cdot \widetilde{x} : \mathbb{S}^{8,13} \to \nu(\mathbb{S}/2)$ is 0. So it extends over the cofiber $\mathbb{S}^{8,12}/\widetilde{2} \simeq \Sigma^{8,12}\nu(\mathbb{S}/2)$, to a self-map

$$\widetilde{\theta} : \Sigma^{8,12}\nu(\mathbb{S}/2) \to \nu(\mathbb{S}/2).$$

This lifts the algebraic self-map of $\mathbb{F}_2/h_0$ from Section 11.1 (the differences in grading are related to the equivalence $\mathrm{Mod}_{\Sigma^*\pi_*(R)}(\mathrm{gr}\,\mathrm{Sp}) \simeq \mathrm{Mod}_{\pi_*(R)}(\mathrm{gr}\,\mathrm{Sp}) \simeq \mathcal{D}^{\mathrm{gr}}(R_*)$ involved in the identification of $\mathbb{S}^{0,0}/\tau$-modules.)

Write $\widetilde{X}$ for the cofiber of $\widetilde{\theta}$. Its underlying spectrum $\tau^{-1}\widetilde{X}$ is X. But its associated spectral sequence has first page given by the cofiber of the algebraic self-map of Adams periodicity, whose cofiber is concentrated below a line of slope $\frac{1}{5}$. Thus, a huge family of differentials in the spectral sequence of $\nu(X)$ have been "accelerated" in such a way that they already happen algebraically in the first page.

12.5 Vanishing lines and self-maps

We just saw that we can build a synthetic spectrum of the form $\mathbb{S}_2^{0,0}/(\widetilde{2},\widetilde{\theta})$ whose bigraded homotopy groups $\pi_{n,w}$ are concentrated below a line of slope $\frac{1}{5}$ in the $(n, w-n)$ Adams grading. The long exact sequence for the cofiber sequence

$$\Sigma^{8,12}(\mathbb{S}_2^{0,0}/\widetilde{2}) \to \mathbb{S}_2^{0,0}/\widetilde{2} \to \mathbb{S}_2^{0,0}/(\widetilde{2},\widetilde{\theta})$$

shows that this is equivalent to saying that above a line of slope $\frac{1}{5}$, θ acts through isomorphisms. Note that θ is "parallel" to the vanishing line of slope $\frac{1}{2}$ that $\mathbb{S}_2^{0,0}/\widetilde{2}$ has.

A natural question is if this process continues: Can we find some kind of self-map on $\mathbb{S}_2^{0,0}/(\widetilde{2}/\widetilde{\theta})$, of slope $\frac{1}{5}$ (which means a suspension of $\Sigma^{5k,6k}$, since this is a shift of $(5k,k)$ in Adams grading), which induces an isomorphism above a line of even lower slope? Does this continue on?

This and related questions are settled in the author's thesis [11], building on work by Palmieri [19]. The approach is roughly to reduce these questions in $\mathrm{Syn}_{\mathbb{F}_p}$ to $\mathrm{Mod}_{\mathbb{S}/\tau}(\mathrm{Syn}_{\mathbb{F}_p}) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$. In this algebraic world, one additional degree of freedom is that one can compare with quotient Hopf algebras $\mathcal{A}_*/I$ for suitable ideals: For each of those, one has a symmetric-monoidal category $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*/I})$, and various adjoint functors (corestriction/coinduction) between them. These typically do not arise from spectra in the same way that $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ arises from $\mathbb{S} \to \mathbb{F}_p$, so this is really an algebraic phenomenon.

We sketch the relevant ingredients at $p = 2$. The quotient algebra

$$\mathcal{A}_*/(\zeta_i^{2^{e_i}})$$

for a sequence e_i, $i \geq 1$, inherits a Hopf algebra structure if and only if for all $i, j \geq 1$ either $e_i \leq j + e_{i+j}$ or $e_j \leq e_{i+j}$. In fact, by a theorem of Adams-Margolis [1], this classifies all Hopf algebra quotients of $\mathcal{A}_*$.

Now consider two quotients $B = \mathcal{A}_*/(\zeta_i^{2^{e_i}})$ and $C = \mathcal{A}_*/(\zeta_i^{2^{e'_i}})$ where if e_i and e'_i differ for exactly one i, and there $e'_i = e_i - 1$. In that case, we have a quotient map $B \to C$ whose kernel is generated by the single element $b = \zeta_i^{2^{e_i}}$ which squares to zero, and this is necessarily primitive in B, i.e. $\Delta(b) = b \otimes 1 + 1 \otimes b$. We may think of this as exhibiting B as some kind of Hopf algebra extension of C by an exterior algebra on $\Lambda_{\mathbb{F}_2}(b)$.

For $f : B \to C$ such a quotient map, there is a corestriction functor $f^* : \mathcal{D}_\omega(\mathrm{Comod}_B) \to \mathcal{D}_\omega(\mathrm{Comod}_C)$, and a right adjoint coinduction functor $f_* : \mathcal{D}_\omega(\mathrm{Comod}_C) \to \mathcal{D}_\omega(\mathrm{Comod}_B)$. These satisfy various identities, for example the *projection formula* $f_*(X) \otimes Y \simeq f_*(X \otimes f^*(Y))$, and f^* is symmetric-monoidal. Using these, one can write

$$\begin{aligned}
\mathrm{map}_{\mathcal{D}_\omega(\mathrm{Comod}_C)}(\mathbb{1}(w), f^*X) &\simeq \mathrm{map}_{\mathcal{D}_\omega(\mathrm{Comod}_C)}(f^*\mathbb{1}(w), f^*X) \\
&\simeq \mathrm{map}_{\mathcal{D}_\omega(\mathrm{Comod}_B)}(\mathbb{1}(w), f_*f^*X) \\
&\simeq \mathrm{map}_{\mathcal{D}_\omega(\mathrm{Comod}_B)}(\mathbb{1}(w), f_*(\mathbb{1}) \otimes X).
\end{aligned}$$

This means that $\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_C)}(f^*X)$, the homotopy of X over C, can be written as $\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_B)}(f_*(\mathbb{1}) \otimes X)$, the "$f_*(\mathbb{1})$-homology" of X in B.

In the case where $B \to C$ are as described above, $f_*(\mathbb{1})$ does in fact have a description as an extension of $\mathbb{1}$ and a degree shift of $\mathbb{1}$ (in fact, it is represented by an ordinary B-comodule whose underlying $\mathbb{F}_2$-module is given by $\Lambda_{\mathbb{F}_2}(b)$. So $f_*(\mathbb{1}) \otimes X$ is an extension of two copies of X (one degree-shifted). In particular, if $\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_B)}(X)$ has a vanishing line, $\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_C)}(f^*X)$ has the same vanishing line.

We can also try to go back:

Proposition 12.25. *For $f : B \to C$ as above and $X \in \mathcal{D}_\omega(\mathrm{Comod}_B)$ a compact object, there is a spectral sequence*

$$\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_C)}(f^*X)[\beta] \to \pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_B)}(X)$$

where $f^* : \mathcal{D}_\omega(\mathrm{Comod}_B) \to \mathcal{D}_\omega(\mathrm{Comod}_C)$ is the corestriction functor.

This spectral sequence can be constructed as an Adams spectral sequence in $\mathcal{D}_\omega(\mathrm{Comod}_B)$ with respect to the algebra object $f_*(\mathbb{1})$. The polynomial generator β is a permanent cycle which arises from $\mathrm{Ext}_{\Lambda(b)}^{*,*}(\mathbb{F}_2, \mathbb{F}_2) \cong \mathbb{F}_2[\beta]$.

If X is an algebra object, this spectral sequence is multiplicative. Now assume X has an algebra structure and f^*X has a vanishing line of some slope over C. Then there are the following possibilities:

1. The slope of β is lower (or equal) than the slope of the vanishing line of $f^*(X)$ over C. Then the entire first page of the spectral sequence lies under that line, and so X has the same vanishing line over B.

2. The slope of β is higher, but some power β^k becomes zero on some page. Then all multiples of β^k are zero from that page on, and so that entire page lies below a shift of the vanishing line of $f^*(X)$ over C. This then implies the same result for the E^∞ page and so X has a vanishing line of the same slope (but possibly shifted offset) over B.

3. The slope of β is higher and no power of β becomes zero, so all of them are nonzero on the E^∞ page. Then the E_∞ page still has a vanishing line of slope equal to the slope of β (and no smaller slope).

So a vanishing line over C implies one over B, but the slope or offset can increase. Applying this inductively to a chain of finite quotients of $\mathcal{A}_*$

$$\mathbb{F}_2 \leftarrow \mathbb{F}_2[\zeta_1]/\zeta_1^2 \leftarrow \dots$$

we get that any compact algebra object $X \in \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ has a vanishing line over each finite quotient algebra of $\mathcal{A}_*$. Since the degrees of the elements b in such a chain of quotients goes to ∞, the slope and offset of the vanishing line does in fact only increase finitely many times, and there is a limiting argument to conclude that one also has a vanishing line over $\mathcal{A}_*$.

As written, this works only for algebra objects, but there is a powerful trick to extend to all compact objects. Since a object $X \in \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ can be written as a retract of a finite extension of shifts of $\mathbb{1}$, $X \otimes Y$ for any Y can be written as a retract of a finite extension of shifts of Y. If Y has a vanishing line of some slope, $X \otimes Y$ has a vanishing line of the same slope, so "all Y which have a vanishing line of a given slope" behaves like an ideal (in compact objects).

Compact objects in $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ are also dualizable: We call an object X of a symmetric-monoidal ∞-category dualizable if there is an object $X^\vee$ (the dual) and maps

$$\mathbb{1} \to X \otimes X^\vee$$
$$X^\vee \otimes X \to \mathbb{1}$$

such that the two composites

$$X \to\to X \otimes X^\vee \otimes X \to X$$
$$X^\vee \to X^\vee \otimes X \otimes X^\vee \to X^\vee$$

are homotopic to the identities. The motivating example is the category of vector spaces with symmetric-monoidal structure given by tensor products, where the dualizable objects are exactly the finite-dimensional vector spaces.

In our categories, $\mathbb{1}(w)$ is dualizable with dual $\mathbb{1}(-w)$, and this implies that all compact objects are dualizable. One also has that $X^\vee \otimes X$ has an algebra structure (in the vector space example, this is the endomorphism algebra of a finite-dimensional vector space). Now, if X has a vanishing line, $X^\vee \otimes X$ has a vanishing line of the same slope, by the "ideal" property above. But the converse is also true, since the composite

$$X \to\to X \otimes X^\vee \otimes X \to X$$

expresses X as retract of $X \otimes X^\vee \otimes X$, which inherits a vanishing line from $X^\vee \otimes X$ again by the ideal property. This allows us to extend the existence of vanishing lines from algebra objects to all X. Along a chain of finite Hopf algebra quotients of $\mathcal{A}_*$, these slopes only increase, and only change finitely often, so every X has a minimal vanishing line whose slope agrees with the slope of one of the β encountered. Let us write slope(X) for that slope, more generally $\mathrm{slope}_B(X)$ for the slope over some quotient Hopf algebra. Let us also write slope(x) for the slope of an element (with respect to Adams grading, so $\frac{w-n}{n}$ if $x \in \pi_{n,w}$).

Because τ has degree $(0,-1)$ in Adams grading, the existence of vanishing lines in $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ for $\pi_{*,*}(X)$ for any compact object implies the same in $\mathrm{Syn}_{\mathbb{F}_p}$, using the τ-Bockstein spectral sequence.

Next, we deal with the question of self-maps parallel to the minimal vanishing line. Let $X \in \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{A}_*})$ be compact, and again first assume that X has an algebra structure. By the above, $\pi_{*,*}(X)$ has a vanishing line of some slope, arising as slope(β) in one of the spectral sequences from 12.25 in a chain of finite quotients of $\mathcal{A}_*$. Fixing such a chain, let us look at the step $B \to C$ where the slope of the vanishing line increases for the last time, so $\mathrm{slope}(X) = \mathrm{slope}_B(X) > \mathrm{slope}_C(X)$. Then, the element β in the first page

$$\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_C)}(f^*X)[\beta]$$

is a permanent cycle (that is the only case in which the slope increases), and any element $v \in \pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_B)}(X)$ detected by β acts on X as an endomorphism (using the algebra structure). The cofiber can be identified with $f_*(\mathbb{1}) \otimes X$, in particular it has a vanishing line of slope $\mathrm{slope}_C(X) < \mathrm{slope}(X)$, and so multiplication by v acts as an isomorphism on $\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_B)}(X)$ above a line of slope $\mathrm{slope}_C(X)$. In the E_∞ page, β is central (commutes with everything), one can see that this allows us to pass to a power of v which is actually central in $\pi_{*,*}(X)$.

Going further up along the sequence of Hopf algebra quotients of $\mathcal{A}_*$, since the slope doesn't increase again, we know that in the spectral sequence for a later extension $B \to C$, we are always in one of the following two situations:

1. We have $\mathrm{slope}(\beta) \leq \mathrm{slope}(X)$. In fact, it cannot be equal, since the elements $\zeta_i^{2^j}$ appearing as b in the extensions of Hopf algebras in the chain are all of different degree, so every value of slope(β) is encountered

at most once along our chain, and slope(X) agrees with the slope of a previous β. So slope(β) < slope(X). In the region where the powers v^n live, this means that only multiples of finitely many powers of β appear. The spectral sequence is trigraded, with third degree coming from the powers of β, and d_r increases this third grading by r. This means that in the region where the powers v^n live, this third grading is bounded and so the spectral sequence degenerates in that region after a finite page. Say v is $r-1$-cycle, but $d_r(v) \neq 0$. Then, because we are over $\mathbb{F}_p$, $d_r(v^p) = pv^{p-1}d_r(v^p)$ (using that v is central), so v^p is an r-cycle. Together with boundedness of the third grading near the region where the v^n live, it follows that some power v^{p^i} is a permanent cycle, represented in homotopy by some v'. Over C, $\mathrm{slope}_C(X/v') = \mathrm{slope}_C(X/v^{p^i}) = \mathrm{slope}_C(X/v)$, so $\mathrm{slope}_B(X/v') < \mathrm{slope}(X)$ and v' is still an isomorphism above a line of slope smaller than slope(X). Again one can see that some power of v' is also central in $\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_B)}(X)$.

2. We have slope(β) > slope(X), but some power β^k becomes zero on some page. Then we have the same boundedness in the third grading of the spectral sequence, even globally, and the same argument shows that some power v^{p^i} is a permanent cycle. We again have $\mathrm{slope}_C(X/v') = \mathrm{slope}_C(X/v)$, and by the general discussion of how the slopes develop along extensions $B \to C$ we must have $\mathrm{slope}_B(X/v') = \mathrm{slope}_C(X/v')$ or $\mathrm{slope}_B(X/v') = \mathrm{slope}(\beta)$. But the latter is absurd, as $\mathrm{slope}_B(X/v') \leq \mathrm{slope}_B(X)$.

This means that in the chain of extensions, after the extension where the slope of the vanishing line of X increased for the last time, we may always lift a power of the periodicity class v up the extension. Once the (horizontal) degrees of the elements β become large enough, the offset of the vanishing line doesn't increase anymore, and eventually they become high enough (the precise degree depending on the offset of the vanishing line) that in the first case above, there are no nontrivial powers of β along the region where the elements v^n live, so v itself lifts. Then, a limiting argument allows us to lift v all the way.

Using multiplicativity of the τ-Bockstein spectral sequence, we may perform the same argument to lift a periodicity class $v \in \pi_{*,*}(X/\tau)$ to one in $\pi_{*,*}(X)$: Again in any region parallel to the vanishing line only finitely many powers of τ appear, and so the spectral sequence filtration is bounded in that region. Some power v^{p^i} is then a permanent cycle.

For arbitrary X, we may use $\pi_{n,w}(X^\vee \otimes X) \simeq [\Sigma^{n,w}X, X]$ (a general consequence of duality) to translate such a $v \in \pi_{*,*}(X^\vee \otimes X)$ to a self-map $\theta : \Sigma^{n,w}X \to X$ whose cofiber X/θ satisfies slope(X/θ) < slope(X).

For odd p, there are similar arguments, but the exterior algebras $\Lambda_{\mathbb{F}_2}(b)$ for even $|b|$ get replaced by truncated polynomial algebras $\mathbb{F}_p[b]/b^p$, and in the spectral sequence comparing derived comodules over two adjacent Hopf algebra quotients $B \to C$ the polynomial algebra generated by β (which came from $\mathrm{Ext}_{\Lambda_{\mathbb{F}_2}(b)}(\mathbb{F}_2, \mathbb{F}_2)$) gets replaced by $\mathrm{Ext}_{\mathbb{F}_p[b]/b^p}(\mathbb{F}_p, \mathbb{F}_p) \cong \Lambda(\alpha) \otimes \mathbb{F}_p[\beta]$, with β now coming from Ext^2.

We summarize the results as follows:

Theorem 12.26 ([11]). *Let $X \in (\mathrm{Syn}_{\mathbb{F}_p})^\omega$. Then $\pi_{*,*}(X)$ admits a minimal vanishing line, whose slope (in Adams grading) is of the form*

$$\frac{1}{(2^i-1)2^j-1}$$

with $i \geq 1$ and $j \geq 0$ for $p = 2$, and of the form

$$\frac{1}{2(p^i-1)} \ \mathit{or}\ \frac{1}{(p^i-1)p^{j+1}-1}$$

with $i \geq 0$ in the first case and $i \geq 1$ and $j \geq 0$ in the second for odd p.

Furthermore, X admits a self-map parallel to that vanishing line, whose cofiber has a vanishing line of strictly lower slope.

In fact, not all of the above slopes do appear (by a theorem of Miller-Wilkerson, some of the associated self-maps become nilpotent no matter what), and the self-map is "asymptotically unique" in the sense that two different such self-maps agree after taking suitable powers of them. The slopes here arise from the slopes of nonnilpotent elements in Ext of the Hopf algebras $\Lambda_{\mathbb{F}_2}(b)$ for $b = \zeta_i^{2^j}$ at $p = 2$ and $\Lambda_{\mathbb{F}_p}(b)$ for $b = \tau_i$ resp. $\mathbb{F}_p[b]/b^p$ for $b = \xi_i^{p^j}$ at odd p.

12.6 MU-synthetic spectra

We now consider Syn_{MU}, i.e. the version of synthetic spectra related to the MU-based Adams spectral sequence. (This is also called the Adams-Novikov spectral sequence.) This is a category where

$$\mathrm{Mod}_{\mathbb{1}/\tau}(\mathrm{Syn}_{MU}) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU})$$

and

$$\mathrm{Mod}_{\mathbb{1}[\tau^{-1}]}(\mathrm{Syn}_{MU}) \simeq \mathrm{Sp}\,.$$

Again, we have a "synthetic analogue" functor ν with $\nu(\Sigma X) = \Sigma^{1,1}(X)$.

Remark 12.27. In some parts of the literature, MU-synthetic spectra are instead constructed in terms of the "double speed Postnikov filtration" $\tau_{\geq *}MU^{\otimes_{\mathbb{S}}\bullet+1}$, which yields a slightly different category and a synthetic analogue functor with $\nu(\Sigma^{2n}X) \simeq \Sigma^{2n,n}\nu(X)$ and $\nu(\Sigma^{2n+1}X) \simeq \Sigma^{2n+1,n}\nu(X)$. This is closer to phenomena observed in motivic homotopy theory over $\mathbb{C}$, and one of the original application of synthetic spectra was to model phenomena observed there. The filtrations obtained from the "double speed Postnikov filtration" also have a natural interpretation in terms of Hahn-Raksit-Wilson's *even filtration.* However, for our applications here, it is sufficient to stick with the single-speed filtration which works for $\mathrm{Syn}_{R/S}$ in general.

In the following, we will make nontrivial use of some deep results about MU_*MU. These rely crucially on Quillen's interpretation of MU_*MU in terms of formal group laws, and are best understood in that language. This is discussed in depth in most introductions to chromatic homotopy theory, but we will take some results for granted and just give an overview over the key computational ingredients that we are using.

$MU_*MU = \pi_*(MU \otimes MU)$ is truly only a Hopf algebroid, with the two ring maps $\eta_L, \eta_R : \pi_*MU \to MU_*MU$ ("left unit" and "right unit") quite different. In fact,

$$\mathrm{Ext}^{0,*}_{MU_*MU}(\pi_*MU, \pi_*MU) \cong \ker(\eta_L - \eta_R)$$

is just $\mathbb{Z}$, generated by 1, so much smaller than the "ground ring" π_*MU. This means that the unit $\mathbb{1}$ of $\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU})$ has very few endomorphisms. Of course, there is a very nontrivial higher Ext.

The cofiber $\mathrm{cofib}(\mathbb{1} \xrightarrow{p} \mathbb{1}) \simeq \mathbb{1}/p$ is represented by the comodule π_*MU/p, and it turns out in

$$\mathrm{Ext}^{0,*}_{MU_*MU}(\pi_*MU, \pi_*MU/p) \cong \ker(\eta_L - \eta_R : \pi_*MU \to MU_*MU/p) \cong \mathbb{F}_p[v_1],$$

one suddenly finds a family of additional elements, powers of a v_1 of degree $2(p-1)$ (of course, this depends on p, and there is a separate such element for each p. Usually, the p is fixed implicitly and not part of the notation, like for the Steenrod algebra $\mathcal{A}$). This corresponds to an element $\widetilde{v}_1 \in \pi_*(MU)$ which is in the kernel of $\eta_L - \eta_R$ only mod p. If we view v_1 as endomorphism of $\mathbb{1}/p$ (with a shift: In our (n, w)-grading, it would correspond to $\Sigma^{2(p-1),2(p-1)}\mathbb{1}/p \to \mathbb{1}/p$), we can form its cofiber $\mathbb{1}/(p, v_1)$, which is represented by $\pi_*MU/(p, v_1)$. The fact that $\eta_L(v_1) = \eta_R(v_1)$ mod p is directly related to the fact that this still has a well-defined comodule structure.

It turns out that this pattern continues, and

$$\begin{aligned} &\mathrm{Ext}^{0,*}_{MU_*MU}(\pi_*MU, \pi_*MU/(p, \dots, v_{n-1})) \\ &\cong \ker(\eta_L - \eta_R : \pi_*MU \to MU_*MU/(p, \dots, v_{n-1})) \cong \mathbb{F}_p[v_n], \end{aligned}$$

for some element v_n of degree $2(p^n - 1)$. Writing $I_n = (p, \dots, v_{n-1})$ and $I_\infty = \bigcup I_n$ for the ideals appearing here, we get an object in the colimit which is represented by the comodule π_*MU/I_∞. We have Hopf algebroid structures on MU_*MU/I_n (where because of the fact that $\eta_L(v_n) = \eta_R(v_n)$ mod I_n it doesn't matter which module structure we use to form the quotient), and the π_*MU/I_n comodules are the image of the unit under the right adjoints to

$$\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU}) \to \mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU/I_n}).$$

There is now a quite nontrivial equivalence

$$\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU/I_\infty}) \simeq \mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{P}_*}),$$

where $\mathcal{P}_* \subseteq \mathcal{A}_*$ is the Hopf subalgebra generated by the ζ_i^2 resp. ξ_i. This equivalence does not hold at the level of Hopf algebroids themselves, as $\pi_*(MU)/I_\infty$

is still quite a lot bigger than $\mathbb{F}_p$ (it is a polynomial ring on one generator of degree $2k$ for each k where $k+1$ is not a power of p), and similarly MU_*MU/I_∞ is bigger than $\mathcal{P}_*$. The cleanest explanation of this equivalence is that the cosimplicial graded rings

$$\pi_*MU/I_\infty \rightrightarrows MU_*MU/I_\infty \Rrightarrow \dots.$$

and

$$\mathbb{F}_p \rightrightarrows \mathcal{P}_* \Rrightarrow \dots.$$

present the same *stacks*: Applying the Yoneda embedding $\mathrm{gr\,CRing}^{\mathrm{op}} \to \mathrm{Fun}(\mathrm{gr\,CRing}, \mathcal{S})$ and taking colimits in this presheaf category, the above diagrams yield equivalent "functors of points". This hinges on Quillen's interpretation of MU_*MU in terms of formal group laws: The functor of points obtained from the first cosimplicial ring takes a ring R to the groupoid consisting of formal group laws isomorphic to the additive formal group law, whereas the second one yields the full subgroupoid consisting just of the additive formal group law itself and isomorphisms. These are of course equivalent, and one can show that $\mathcal{D}_\omega(-)$ depends only on these functors of points, because it can be interpreted as Kan extension.

Taking these results for granted, we get that we can go from $\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU})$ to $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{P}_*})$ by a (infinite) sequence of quotients by v_n.

In $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{P}_*})$, the same arguments as for $\mathcal{A}_*$ can be used to prove existence of minimal vanishing lines and periodicity self-maps parallel to those. In fact, the description becomes a bit more uniform at odd p, since $\mathcal{P}_*$ does not have the exterior generators from $\mathcal{A}_*$.

Theorem 12.28 ([11])**.** *Let $X \in (\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{P}_*}))^\omega$. Then $\pi_{*,*}(X)$ admits a minimal vanishing line, whose slope (in Adams grading) is of the form*

$$\frac{1}{(p^i-1)p^{j+1}-1}$$

for some p, $i \geq 1$ and $j \geq 0$.

Furthermore, X admits a self-map parallel to that vanishing line, whose cofiber has a vanishing line of strictly lower slope.

By a limiting argument, for any given $X \in (\mathrm{Syn}_{MU})^\omega$, these results applied to the image of X in $\mathcal{D}_\omega(\mathrm{Comod}_{\mathcal{P}_*})$ imply the corresponding result for the image of X in $\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU/I_n})$ for some n (depending on the slope and offset of the vanishing line). Then, there is a spectral sequence

$$\pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU/I_{n+1}})}(X/I_{n+1})[v_n] \to \pi_{*,*}^{\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU/I_n})}(X/I_n)$$

which can be constructed analogously to the spectral sequence of Proposition 12.25 from an adjunction (but also can be identified with a "v_n-Bockstein spectral sequence"). Since the slope of v_n is 0, we inductively get the same minimal vanishing line as we had for X/I_∞ for X/p, and a periodicity map parallel to that vanishing line, by similar but easier arguments as for the Hopf algebra extensions in the previous section.

From the existence of minimal vanishing lines and self-maps over MU_*MU/p, we can also infer an "integral" statement over MU_*MU itself. For this, observe that for $X \in (\mathcal{D}_\omega(\mathrm{Comod}_{MU_*MU}))^\omega$, each X/p has a minimal vanishing line, and the highest possible slope for it is $\frac{1}{(p-1)p-1}$, which goes to 0 as $p \to \infty$. In particular, one of these lines attains the maximal slope, say the one for the prime p_0. One can also see that all potential slopes are different. Now $\pi_{*,*}(X/p)$ for $p \neq p_0$ is not necessarily concentrated below that line, but one can see that there is some minimal shift of this line which serves as a vanishing line for all $\pi_{*,*}(X/p)$ at the same time. One also sees that

$$\pi_{*,*}(X \otimes \mathbb{Q})$$

for compact X is concentrated in finitely many degrees (this is closely related to the fact that $\mathrm{Ext}_{MU_*MU}(MU_*, MU_*)$ is torsion in all degrees except $(0,0)$). We may shift the vanishing line in a way to also include these rationalized homotopy groups. Now the cofiber of $X \to X \otimes \mathbb{Q}$ may be written as filtered colimit of objects of the form X/n for varying n, which may in turn each be written as extension of objects of the form X/p, from which it follows that this vanishing line also serves as vanishing line for $\pi_{*,*}(X)$.

By a more subtle connectivity argument, one can in fact see that in any region parallel to the minimal vanishing line for $\pi_{*,*}(X)$, the groups are torsion of uniformly bounded exponent. Using this, one can lift a periodicity self-map of X/p_0 to one of X, by lifting it in the p-Bockstein spectral sequence.

Finally, from $\mathcal{D}_\omega(MU_*MU)$ to Syn_{MU}, one can lift both the vanishing line and the self-map along the τ-Bockstein spectral sequence.

Theorem 12.29 ([11])**.** *Let $X \in (\mathrm{Syn}_{MU})^\omega$. Then $\pi_{*,*}(X)$ admits a minimal vanishing line, whose slope (in Adams grading) is of the form*

$$\frac{1}{(p^i-1)p^{j+1}-1}$$

for some p, $i \geq 1$ and $j \geq 0$.

Furthermore, X admits a self-map parallel to that vanishing line, whose cofiber has a vanishing line of strictly lower slope.

Finally, what can we say about non-compact objects of Syn_{MU}? Clearly, there doesn't have to be a vanishing line in full generality, as we have objects like $\bigoplus_{n,w} \mathbb{S}^{n,w}$. In order to get controllable vanishing lines for a synthetic spectrum X, we want some control over what cells $\mathbb{S}^{n,w}$ can be built from.

For an ordinary connective spectrum X with a chosen cell structure, cellular homology tells us that every n-cell roughly contributes either a generator to $H_n(X)$ or a relation to $H_{n-1}(X)$. It turns out that by a combination of Hurewicz and the construction from the CW approximation theorem, this admits a kind of converse: A connective spectrum X admits a cell structure with one n-cell for each generator of $H_n(X)$ and relation of $H_{n-1}(X)$ in chosen presentations of those groups. In particular, every connective X admits a cell structure with cells only in those degrees n for which $H_n(X)$ or $H_{n-1}(X)$ are nonzero.

It turns out we can play a similar trick in Syn_{MU}. The Adams-Novikov spectral sequence for $H\mathbb{Z}$ itself degenerates, with all of $\mathrm{Ext}^{*,*}_{MU_*MU}(MU_*, MU_*H\mathbb{Z})$ given by $\mathbb{Z}$ concentrated in degree $(0,0)$. This means that $\nu(H\mathbb{Z})/\tau$ has bigraded homotopy groups given by $\mathbb{Z}$ concentrated in degree $(0,0)$. We also have that

$$\nu(H\mathbb{Z})/\tau \otimes \mathbb{S}^{n,w}$$

has bigraded homotopy groups given by $\mathbb{Z}$ in degree (n,w). So if X admits a cell structure, every $\mathbb{S}^{n,w}$ contributes a generator to $\pi_{n,w}(\nu(H\mathbb{Z})/\tau \otimes X)$ or a relation to $\pi_{n-1,w}(\nu(H\mathbb{Z})/\tau \otimes X)$.

Conversely, using if X satisfies some connectivity condition (vertically bounded below and τ-complete, to be precise), it admits a cell structure with one (n,w)-cell for each generator of $\pi_{n,w}(\nu(H\mathbb{Z})/\tau \otimes X)$ and relation of $\pi_{n-1,w}(\nu(H\mathbb{Z})/\tau \otimes X)$ in chosen presentations of those abelian groups.

Proposition 12.30. *For X a connective spectrum, $\nu(X)$ admits a cell structure with cells $\mathbb{S}^{n,w}$ with $n \geq w \geq 0$. (In particular, in the lower right quadrant $n \geq 0$, $s \leq 0$ in Adams grading)*

Proof. We have $\mathrm{Mod}_{\nu(H\mathbb{Z})/\tau}(\mathrm{Syn}_{MU}) \simeq \lim_{\Delta}^{\mathrm{Pr}^L_\omega} \mathrm{Mod}_{\Sigma^*\pi_*(MU^{\otimes_{\mathbb{S}}\bullet+1}\otimes_{\mathbb{S}} H\mathbb{Z})}(\mathrm{gr}\,\mathrm{Sp})$, and so we have the following commutative diagram

$$\begin{array}{ccc}
\mathrm{Syn}_{MU} & \longrightarrow & \mathrm{Mod}_{\tau_{\geq *}MU}(\mathrm{FilSp}) \\
\downarrow & & \downarrow \\
\mathrm{Mod}_{\mathbb{S}^{0,0}/\tau}(\mathrm{Syn}_{MU}) & \longrightarrow & \mathrm{Mod}_{\Sigma^*\pi_*MU}(\mathrm{gr}\,\mathrm{Sp}) \\
\downarrow & & \downarrow \\
\mathrm{Mod}_{\nu(H\mathbb{Z})/\tau}(\mathrm{Syn}_{MU}) & \longrightarrow & \mathrm{Mod}_{\Sigma^*\pi_*(MU\otimes_{\mathbb{S}} H\mathbb{Z})}(\mathrm{gr}\,\mathrm{Sp}).
\end{array}$$

As $\nu(H\mathbb{Z})/\tau$ as a graded spectrum is simply $\mathbb{Z}$ concentrated in degree 0, the bottom horizontal functor is split by the base-change $\mathrm{Mod}_{\Sigma^*\pi_*(MU\otimes_{\mathbb{S}} H\mathbb{Z})}(\mathrm{gr}\,\mathrm{Sp}) \to \mathrm{Mod}_{\mathbb{Z}(0)}(\mathrm{gr}\,\mathrm{Sp})$. This lets us identify the "homology" functor

$$\mathrm{Syn}_{MU} \to \mathrm{Mod}_{\nu(H\mathbb{Z})/\tau}(\mathrm{Syn}_{MU})$$

with the composite

$$\mathrm{Syn}_{MU} \to \mathcal{D}^{\mathrm{gr}}(\pi_*MU) \to \mathcal{D}^{\mathrm{gr}}(\mathbb{Z})$$

which first passes to the associated graded of the underlying $\tau_{\geq *}MU$-module, and then performs the (derived) base-change $\mathbb{Z} \otimes^L_{\pi_*MU} (-)$.

The first of these functors takes $\nu(X)$ to the ordinary π_*MU-module MU_*X, and the second to the derived tensor product $\mathbb{Z} \otimes^L_{\pi_*MU} MU_*X$. Under the regrading arising from identifying $\mathrm{Mod}_{\Sigma^*\pi_*MU}(\mathrm{gr}\,\mathrm{Sp}) \simeq \mathcal{D}^{\mathrm{gr}}(\pi_*MU)$, it follows that

$$\pi_{n,w}(\nu(H\mathbb{Z})/\tau \otimes \nu(X)) \cong (\mathrm{Tor}^{\pi_*MU}_{n-w}(\mathbb{Z}, MU_*X))_w,$$

which is in particular zero if $n - w < 0$ or $w < 0$ (since by connectivity of X, MU_*X is concentrated in degrees ≥ 0). ☐

Corollary 12.31. *If $X \in \mathrm{Syn}_{MU}$ has a vanishing line of positive slope, then $X \otimes \nu(Y)$ has the same vanishing line.*

Proof. For $n \geq w \geq 0$, $X \otimes \mathbb{S}^{n,w}$ is a shift of X to the right and down (in Adams grading), and in particular concentrated under the same vanishing line as X. Since $\nu(Y)$ is built from such cells under extensions and colimits, this follows more generally for $X \otimes \nu(Y)$. ☐

13 Burklund's proof of the nilpotence theorem

In this section, we will finally prove the nilpotence theorem, following ideas of Robert Burklund. Where the original proof of [7] uses a sequence of very specific spectra $X(n)$ and a subtle analysis of their Adams-Novikov spectral sequences, Burklund's idea is to replace these by a much more arbitrary sequence of *synthetic* spectra which have the desired properties by design. Their existence and properties are an easy consequence of Theorem 12.29 and Nishida nilpotence. Another interesting feature of this proof is that it uses Nishida nilpotence as the key homotopy-theoretic ingredient, in a way that could just be cited as a black-box result. This appears more indirectly in Devinatz-Hopkins-Smith proof: For the inductive step relating nilpotence over $X(n+1)$ to nilpotence over $X(n)$, they use ideas from Nishida's original paper [18], but with nontrivial modifications.

The nilpotence theorem can be stated in several equivalent forms:

Theorem 13.1 (Nilpotence theorem, smash form)**.** *Assume X and Y are spectra with X finite, and $f : X \to Y$ a map such that $X \to MU \otimes Y$ is nullhomotopic. Then f is nilpotent under $\otimes$, i.e. there exists k such that*

$$f^{\otimes k} : X^{\otimes k} \to Y^{\otimes k}$$

is nullhomotopic.

Theorem 13.2 (Nilpotence theorem, connective ring form)**.** *Assume R is a connective associative ring spectrum, and $f \in \pi_n(R)$ an element of homotopy such that the image of f in $MU_n(R)$ is 0. Then f is nilpotent with respect to the ring structure on $\pi_*(R)$.*

Theorem 13.3 (Nilpotence theorem, ring form)**.** *Assume R is an associative ring spectrum, and $f \in \pi_n(R)$ an element of homotopy such that the image of f in $MU_n(R)$ is 0. Then f is nilpotent with respect to the ring structure on $\pi_*(R)$.*

The version that we will prove is the connective ring form. It is obviously a special case of the ring form, but in fact all three versions are equivalent.

Proposition 13.4. *The ring form, connective ring form and smash forms of the nilpotence theorem are all equivalent.*

Proof. Assuming the connective ring form, we prove the smash form. If $X \to Y$ is a map from a finite spectrum to an arbitrary spectrum, it factors through some $\tau_{\geq n} Y$ since X is bounded below. By compactness of X and $MU \otimes Y \simeq \mathrm{colim}_n\, MU \otimes \tau_{\geq n} Y$, some map $X \to MU \otimes \tau_{\geq n} Y$ is already nullhomotopic, and so we can choose a factorisation $X \to \tau_{\geq n} Y$ which still satisfies the assumptions of the smash nilpotence theorem. This means we may assume Y bounded below without loss of generality. Using dualizability of X, the map $X \to Y$ corresponds to a map $\mathbb{S} \to X^\vee \otimes Y$, where the target is still bounded below, or a map $\alpha : \mathbb{S}^n \to \Sigma^n X^\vee \otimes Y$ for suitable n, where the target is now connective. By spelling out this map in terms of the original $X \to Y$ and the duality coevaluation $\mathbb{S} \to X \otimes X^\vee$, we also see that the composite $\mathbb{S}^n \to MU \otimes \Sigma^n X^\vee \otimes Y$ is nullhomotopic. Now letting

$$R = \mathrm{Free}^{E_1}(\Sigma^n X^\vee \otimes Y) \simeq \bigoplus_{k \geq 0} (\Sigma^n X^\vee \otimes Y)^{\otimes k},$$

we may view this map as an element of $\pi_n(R)$ for a connective associative ring spectrum R, and so it is nilpotent. As the k-th power α^k is exactly the map $\mathbb{S}^{kn} \to (\Sigma^n X^\vee \otimes Y)^{\otimes k}$, the connective ring form of the nilpotence theorem implies the smash form.

Now assuming the smash form, we prove the general ring form. If R is some associative ring spectrum, and $\alpha \in \pi_n(R)$ some element which becomes zero in $MU_n(R)$, this means exactly that $\mathbb{S}^n \to R$ is a map which becomes nullhomotopic in $MU \otimes R$. So the smash form of the nilpotence theorem applies to show that $\alpha^{\otimes k} : \mathbb{S}^{nk} \to R^{\otimes k}$ is nullhomotopic for some k. As α^k is the composite $\mathbb{S}^{nk} \to R^{\otimes k} \to R$ with the multiplication map, the ring form of the nilpotence theorem follows.

Finally, the ring form obviously implies the connective ring form. □

Remark 13.5. Note that in the implication from the smash form to the general ring form of the nilpotence theorem, we do not use an E_1 structure, and not even associativity or unitality up to homotopy. All we use is that α^k can be described as composite of $\alpha^{\otimes k}$ with a map $R^{\otimes k} \to R$. This stays true if we just have a spectrum R with any "multiplication map" $R \otimes R \to R$, provided we define $\alpha^k \in \pi_*(R)$ by say $\alpha \cdot (\alpha^{k-1})$ inductively, and the map $R^{\otimes k} \to R$ correspondingly.

The nilpotence theorem was proved by Devinatz-Hopkins-Smith to answer in the affirmative an earlier conjecture by Ravenel, which in turn was strongly motivated by Adams periodicity (which we discussed earlier). It deals with self-maps of finite spectra, and as such is a bit less general than the smash form and the ring form.

Theorem 13.6 (Nilpotence theorem, self-map form)**.** *Assume X is a finite spectrum, and $f : \Sigma^n X \to X$ a map which is zero on MU_*-homology. Then f is* *nilpotent under composition, i.e. there exists k such that $f^{\circ k} : \Sigma^{kn} X \xrightarrow{f} \dots \xrightarrow{f} X$ is nullhomotopic.*

Proof assuming the ring form of the nilpotence theorem. First, we observe that if each of the maps

$$X \xrightarrow{f} \Sigma^{-n} X \xrightarrow{f} \Sigma^{-2n} X \xrightarrow{f} \dots$$

is zero on MU-homology, the MU-homology of the colimit $X[f^{-1}]$ vanishes, and hence $MU \otimes X[f^{-1}] \simeq 0$. By compactness of X, this means that already one of the maps $X \to MU \otimes \Sigma^{-kn} X$ is nullhomotopic. Replacing f by some finite composition power, we may therefore replace the condition that $MU_*(f) = 0$ with the a priori stronger condition that $\Sigma^n X \to MU \otimes X$ is nullhomotopic, without loss of generality. Using dualizability of X, the map $\Sigma^n X \to X$ corresponds to a map $S^n \to X^\vee \otimes X$, which still becomes null when tensoring the target with MU. Now, $X^\vee \otimes X \simeq \mathrm{map}(X, X)$ has a ring structure coming from composition, and so the ring form of the nilpotence theorem proves that the map is nilpotent. As the ring structure comes from composition, this implies that the original $\Sigma^n X \to X$ is nilpotent under composition. $\square$

The proof of the connective ring form of the nilpotence theorem starts by reinterpreting it as a statement about the Adams-Novikov spectral sequence computing $\pi_*(R)$ from $MU_*(R)$. The Adams-Novikov 0-line $\mathrm{Ext}^0_{MU_*MU}(MU_*, MU_*R)$ is a subgroup of $MU_*(R)$ (those elements m sent to $1 \otimes m$ under the coaction), and the kernel of $\pi_*(R) \to MU_*(R)$ can be interpreted as those elements which go to 0 in the top associated graded term of the abutment filtration, i.e. which are detected above the 0-line. If $\alpha \in \pi_n(R)$ is a nonzero element in the kernel, it is therefore detected in bidegree (n, s) for some $s > 0$. The powers α^k are detected in bidegree (nk, sk) or even higher filtration, in any case above this line of slope $\frac{s}{n}$. If we had a vanishing line for the Adams-Novikov spectral sequence of R whose slope is smaller than $\frac{s}{n}$, that would suffice: The first power α^k for which (nk, sk) lies above the vanishing line would necessarily be zero.

For the sphere spectrum itself, the best vanishing line for the Adams-Novikov spectral sequence is of slope 1. However, this improves on later pages, and in fact the methods we will use below to prove the nilpotence theorem using synthetic spectra can be strengthened to show that the E_∞ page has a nonlinear vanishing curve whose slope tends to 0. This is clearly enough to prove nilpotence of any class detected in positive filtration, as any line of positive slope will eventually exceed this vanishing line. Little is known about the actual behavior of this vanishing curve: One can get an effective upper bound from the synthetic methods, but since Nishida's nilpotence theorem is an ingredient there and the exponents obtained there are very big, that bound is probably far from optimal. Hopkins has conjectured that the vanishing curve should behave somewhat like $\sqrt{n}$, but no upper or lower bounds close to this are known.

Instead of directly establishing a sublinear vanishing curve for the E^∞ page of the Adams-Novikov spectral sequence for any connective ring R, Devinatz-Hopkins-Smith produce instead a sequence of spectra $X(i)$ which somewhat miraculous properties. On one hand, the Adams-Novikov spectral sequences of

$X(i)$ have vanishing lines whose slopes tend to 0 as $i \to \infty$. On the other hand, they are able to prove a relative version of the nilpotence theorem, which says that if the image of $\alpha \in \pi_n(R)$ is nilpotent in $X(i+1)_*R$, so is the image of α in $X(i)_*R$. For large i, the nilpotence follows from the arbitrarily low slope of the vanishing line, whereas $X(0) = \mathbb{S}$ leads to nilpotence of α itself. The proof of the inductive step, i.e. that $X(i) \to X(i+1)$ detects nilpotence, is however an extremly intricate computational argument.

Burklund's proof of the nilpotence theorem. We begin by applying $\nu : \mathrm{Sp} \to \mathrm{Syn}_{MU}$ to $\alpha : S^n \to R$ to obtain a map $\nu(\alpha) : \mathbb{S}^{n,n} \to \nu(R)$. In terms of the underlying filtered spectrum, this corresponds to a map

$$\Sigma^n \mathbb{1}(n) \to \lim_{\Delta} \tau_{\geq *}(MU^{\otimes_{\mathbb{S}} \bullet+1} R),$$

or a map

$$\mathbb{S}^n \to \lim_{\Delta} \tau_{\geq n}(MU^{\otimes_{\mathbb{S}} \bullet+1} R)$$

into the n-th stage of the filtered spectrum underlying $\nu(R)$. The n-th graded is computed by

$$\Sigma^n \lim_{\Delta} \pi_n(MU^{\otimes_{\mathbb{S}} \bullet+1} R),$$

whose homotopy groups are computed by the cochain complex corresponding to this cosimplicial abelian group, its π_n is therefore a subgroup of $MU_n R$. The original assumption was that the image of α in $MU_n R$ is zero, and tracing back through the definitions we see that the above map factors through the $(n+1)$st stage of the filtration $\nu(R)$, or equivalently that $\nu(\alpha) : \mathbb{S}^{n,n} \to \nu(R)$ can be divided by τ. Write $\widetilde{\alpha} : \mathbb{S}^{n,n+1} \to \nu(R)$ for a choice of map with $\tau \cdot \widetilde{\alpha} \simeq \nu(\alpha)$.

In Adams grading, we have lifted α to an element in bidegree $(n, 1)$. This map $\widetilde{\alpha} : \mathbb{S}^{n,n+1} \to \nu(R)$ still has the property that under the "underlying spectrum" functor τ^{-1} it turns back into α, which we are going to use to see that α is nilpotent.

As our next ingredient, we will use Theorem 12.29 to contruct $C_i \in \mathrm{Syn}_{MU}$ with a map $\mathbb{S}^{0,0} \to C_i$, with the following properties:

1. All C_i are compact.
2. The slope of the minimal vanishing line of C_i goes to 0 as $i \to \infty$.
3. The image $\tau^{-1} C_i \in \mathrm{Sp}$ of C_i splits as a nonzero sum of finitely many spheres, with $\mathbb{S}^{0,0} \to C_i$ becoming the inclusion of one of the summands.

We construct these inductively, starting with $C_0 = \mathbb{S}^{0,0}$ itself. Suppose we have found C_i. Then, by Theorem 12.29, C_i has a minimal vanishing line, and a self-map $v : \Sigma^{|v|} C_i \to C_i$ whose cofiber has a vanishing line of strictly smaller slope. If we didn't require the third condition, $C_{i+1} = C_i/v$ would work, since the set of possible vanishing line slopes does not have other accumulation points than 0. To understand the third condition, we analyze $\tau^{-1} v : \Sigma^{|\tau^{-1} v|} \tau^{-1} C_i \to \tau^{-1} C_i$. This is a map between finite sums of spheres, so we can write it as matrix with entries in $\pi_*(\mathbb{S})$. Replacing v by a power if necessary, we may assume that $|\tau^{-1}v|$ is large enough that all these entries have positive degrees. By Nishida, these entries and hence the entire matrix are nilpotent. So we can replace v by a power for which $\tau^{-1}v$ is nullhomotopic. Then if we let $C_{i+1} = C_i/v$, we have $\tau^{-1}C_{i+1} \simeq \tau^{-1}C_i \oplus \Sigma^{|\tau^{-1}v|+1}\tau^{-1}C_i$, a sum of spheres.

Now everything is in place: Choose i large enough that the slope of the minimal vanishing line of C_i is lower than the slope $\frac{1}{n}$ of $\widetilde{\alpha}$. By Corollary 12.31, $C_i \otimes \nu(R)$ has the same vanishing line. So there exists k such that $\widetilde{\alpha}^k$ has trivial image in $\pi_{*,*}(C_i \otimes \nu(R))$. Inverting τ, we learn that α^k has trivial image in $\pi_*(\tau^{-1}C_i \otimes R)$. But since the map $\mathbb{S} \to \tau^{-1}C_i$ splits, this means that α^k has trivial image in $\pi_*(R)$, finishing the proof. □

Essentially, all that happens in the proof is that we use Theorem 12.29 to produce quotients of $\mathbb{S}^{0,0}$ with flatter and flatter vanishing lines, but then use Nishida to recognize these quotients as really just being interesting modified Adams-Novikov filtrations on sums of spheres.

It is surprising that we only really used two properties of MU: Some kind of regularity of Ext_{MU_*MU} in the form of Theorem 12.29 (we may think of the successive quotienting of self-maps which leads to smaller and smaller vanishing lines as saying that "up to nilpotents" we are passing from Ext_{MU_*MU} closer to $\mathbb{Z}$ by something like a quotient by a regular sequence), and $\pi_0(MU) = \mathbb{Z}$. The latter part is crucial: If we replaced MU by $\mathbb{F}_p$, the first part would still be satisfied, but if we tried to construct a sequence C_i as in the above proof, to pass from C_0 to C_1 we have to quotient by a power of $\widetilde{p}$. This maps to $\pi_0(\mathbb{S})$, and as such Nishida doesn't apply to it. No matter which power we take, the cofiber C_1 will have underlying spectrum of the form $\mathbb{S}/p^n$, which has its own non-nilpotent self-map (from Adams periodicity), etc. The fact that $\pi_0(MU) \cong \pi_0(\mathbb{S})$ ensures that the first self-map we have to quotient by over MU already maps to $\pi_{>0}(\mathbb{S})$, and hence is nilpotent.

14 Applications of the nilpotence theorem

In this section, we want to discuss two important applications of the nilpotence theorem, the thick subcategory and periodicity theorems. The structure follows roughly [10], but we remove some of the external dependencies. Most notably, the existence of a type exactly n complex with a v_n self-map, classically based on an explicit construction of such an object due to Mitchell, is done here through a discussion of vanishing lines and self-maps in $\mathrm{Syn}_{\mathbb{F}_p}$.

We also give a version of obstruction theory for E_1 structures on MU algebras directly based on the E_1 cotangent complex (by working in filtered spectra over $\tau_{\geq *}MU$).

14.1 Morava K-theory and the $K(n)$-form of nilpotence

The nilpotence theorem tells us that any non-nilpotent $f : X \to Y$ (for the smash version) or $f : \Sigma^n X \to X$ (for the self-map version) has to be detected in MU-homology. But $\pi_*(MU)$ is a big ring, so it would be nice to have a clearer picture of the different ways these can be detected. It turns out there is a powerful alternative version of the nilpotence theorem where we get to replace MU by certain "quotients" of MU which have much smaller homotopy groups.

Of course, "quotients" of ring spectra are more subtle than of ordinary rings, for example we mentioned previously that $\mathbb{S}/p$, the cofiber of $p : \mathbb{S} \to \mathbb{S}$, has no ring structure. Over MU, the situation is a bit better, and we have the following:

Proposition 14.1. *Let $x \in \pi_{2k}(MU)$ be a nonzerodivisor in $\pi_*(MU)$. Then the MU-module $MU/x = \mathrm{cofib}(\Sigma^{2k}MU \xrightarrow{x} MU)$ admits the structure of an associative MU-algebra.*

We will sketch a proof of this below, as an application of *deformation theory.* Deformation theory allows us to describe, for certain maps of commutative ring spectra $\widetilde{R} \to R$, and an R-algebra A, the space of all *deformations* of A: $\widetilde{R}$-algebras together with an equivalence $R \otimes_{\widetilde{R}} \widetilde{A}$. A great source for this is [13, Section 7.4], but unfortunately the version developed there is not quite sufficient for our needs here, since it only covers the situation where $\widetilde{R}$, R, $\widetilde{A}$ and A are all E_∞ algebras. We will instead need a version for A and $\widetilde{A}$ E_1 algebras over E_∞ algebras R and $\widetilde{R}$. Instead of studying a space of diagrams

$$\begin{array}{ccc} \widetilde{R} & \longrightarrow & \widetilde{A} \\ \downarrow & & \downarrow \\ R & \longrightarrow & A, \end{array}$$

(which is not the right thing anymore, since there is a difference between "associative algebras over R" and "associative rings with a homomorphism from R"), we will have to study the fiber of the functor $\mathrm{Alg}(\mathrm{Mod}_{\widetilde{R}}(\mathrm{Sp})) \to \mathrm{Alg}(\mathrm{Mod}_R(\mathrm{Sp}))$. We will try to sketch the relevant parts of this version of deformation theory below, without going into all the technical details.

Lemma 14.2. *Let $\widetilde{R} \to R$ be a square-zero extension of connective E_∞ ring spectra, with fiber given by the connective R-module I. Then we have limit diagrams*

$$\begin{array}{ccc} \widetilde{R} & \longrightarrow & R \\ \downarrow & & \downarrow{\scriptstyle s} \\ R & \xrightarrow{\partial} & R \oplus \Sigma I \end{array}$$

and

$$\begin{array}{ccccccc}
\widetilde{R} & & \longrightarrow & & R & & \\
\downarrow & \searrow & & & \downarrow & \searrow & \\
& & R & \longrightarrow & & & R \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
R & & \longrightarrow & & R & & \\
& \searrow & \downarrow & & & \searrow & \downarrow \\
& & R & \longrightarrow & & & R \oplus \Sigma^2 I
\end{array}$$

which give limit diagrams of symmetric-monoidal categories when applying $\mathrm{Mod}_{(-)}(\mathrm{Sp})_+$ *(i.e. forming bounded below modules.)*

Proof sketch. Square zero extensions (c.f. [13, Section 7.4]) are by definition given as a pullback of the first form, with $R \oplus \Sigma I$ endowed with the "split square-zero extension" E_∞ structure (informally, the multiplication is determined by the R-module structure on I and that ΣI squares to 0), and $s, \partial : R \to R \oplus \Sigma I$ two E_∞ maps that project to the identity in the first factor, with s the trivial such map and ∂ an interesting one.

Writing $R \oplus \Sigma I$ itself as pullback $R \times_{R \oplus \Sigma^2 I} R$ analogously, here using the trivial map two times, we may express the map $\partial : R \to R \oplus \Sigma I$ through a square

$$\begin{array}{ccc}
R & \longrightarrow & R \\
\downarrow & & \downarrow{\scriptstyle s} \\
R & \xrightarrow{s} & R \oplus \Sigma^2 I,
\end{array}$$

with the map ∂ now hidden in the choice of homotopy making the square commute. The first pullback may then also be expressed as the cube-shaped limit diagram, by writing it as $R \times_{R \times_{R \oplus \Sigma^2 I} R} (R \times_R R)$.

The fact that these pullback diagrams give pullbacks of bounded below module categories follows from the following general observation: With $R_{00} = R_{01} \times_{R_{11}} R_{10}$ a pullback of ring spectra, $\mathrm{Mod}_{R_{00},+} \to \mathrm{Mod}_{R_{01},+} \times_{\mathrm{Mod}_{R_{11},+}} \mathrm{Mod}_{R_{10},+}$ is always fully faithful and admits a right adjoint. It is an equivalence if and only if its right adjoint, which takes compatible modules M_{01}, M_{10}, M_{11} to the pullback $M_{01} \times_{M_{11}} M_{10}$, is conservative. In our case, M_{01} and M_{10} are the same R-module M, and $M_{11} \simeq M \oplus \Sigma I \otimes_R M$. The underlying spectrum of the pullback P fits then into a fiber sequence

$$I \otimes_R M \to P \to M.$$

If M is k-connective and $\pi_k(M) \neq 0$, then $\pi_k(P) \neq 0$, in particular the functor is conservative (on bounded below modules). For the cube diagram, one gets the statement by iterating pullbacks again. □

Proposition 14.3. *Let $\widetilde{R} \to R$ be a square-zero extension of connective commutative ring spectra, with fiber given by the connective R-module I. Let $A \in \mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}_R)$, and write $\mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}_R)_{/A}$ for the slice category of $\mathcal{O}$-R-algebras with a map to A. Then the fiber $\mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}_{\tilde{R}}) \times_{\mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}_R)} \{A\}$ is an ∞-groupoid, which is equivalent to the space of homotopies between a certain map $A \to (R \oplus \Sigma^2 I) \otimes_R A$ and the canonical such map (in $\mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}_R)_{/A}$). In particular, there exists a lift $\widetilde{A}$ if and only if these maps are homotopic.*

Proof sketch. From the fact that we can write $\mathrm{Mod}(\widetilde{R})$ as limit of symmetric-monoidal categories in a cube-shaped diagram as in Lemma 14.2, we see that we can analogously write $\mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}(\widetilde{R}))$ as such a limit. This means that in order to lift A to a $\widetilde{R}$-algebra, we need to construct a diagram

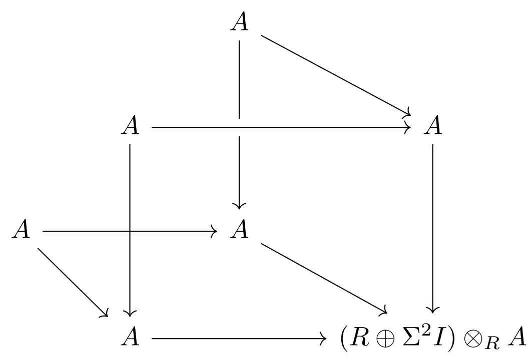

where the maps are linear over R or semilinear over $R \to R \oplus \Sigma^2 I$ respectively. (Formally, this means we are constructing a cocartesian section to the cocartesian fibration $\mathrm{Alg}_{\mathcal{O}}(\mathrm{Mod}) \to \mathrm{CAlg}(\mathrm{Sp})$ over the diagram of commutative rings.)

We work out the data required to fill out this diagram. Working in the slice over A identifies the outer hexagon with the identity on A. Filling the front right vertical arrow (with the unique cocartesian lift) identifies the bottom front right vertex with the basechange $(R \oplus \Sigma^2 I) \otimes_R A$. Now, the back right and front squares can be filled uniquely (since the outer morphisms are equivalences). Finally, filling the remaining square amounts exactly to a homotopy as stated. □

Proof of Lemma 14.1. We will use the formalism of filtered spectra to turn the problem into an inductive one to which we can apply deformation theory. Namely, consider $R := \tau_{\geq *} MU \in \mathrm{FilSp}$. This is a commutative algebra object by lax symmetric-monoidality of the Whitehead filtration. We can then refine x to an element $\widetilde{x} \in \pi_{2k,2k}(\tau_{\geq *} MU) \cong \pi_{2k}(\tau_{\geq 2k} MU)$, or a module map

$$\widetilde{x} : \Sigma^{2k,2k} \tau_{\geq *} MU \to MU$$

Let us write $M := (\tau_{\geq *} MU)/\widetilde{x}$ for the cofiber. Since x was a nonzerodivisor in $\pi_* MU$, $\widetilde{x} : \Sigma^{2k,2k} R/\tau \to R/\tau$ is injective on homotopy groups, and so M/τ is really just given by the graded $\Sigma^*\pi_*MU$-module $\Sigma^*\pi_*MU/x$. Clearly, this has an algebra structure over R/τ, as all of this is just taking place in the heart of some t-structure in graded spectra, and hence in a category equivalent to the usual 1-category of graded MU_*-modules.

We will inductively give M/τ^n the structure of an R/τ^n-algebra. In the limit, this will give M the structure of an R-algebra, and inverting τ, we will have found on MU/x the structure of an MU-algebra. As we saw in Proposition 14.3, the obstruction to lifting the R/τ^n-algebra M/τ^n to an R/τ^{n+1}-algebra is given by an E_1-R/τ^n-algebra map $M/\tau^n \to M/\tau^n \oplus \Sigma^2 M/\tau$.

For an E_1 algebra A over some base ring k, and a bimodule N (i.e. $A \otimes_k A^{\mathrm{op}}$ module), there is a split square-zero E_1 algebra $A \oplus N$, and a bimodule $L^{E_1}_{A/k}$ with the following universal property: E_1-algebra maps $A \to A \oplus N$ lifting the identity $A \to A$ are the same as $A \otimes_k A^{\mathrm{op}}$-module maps $L^{E_1}_{A/k} \to N$. This is called the E_1 cotangent complex. It is determined by a fiber sequence

$$L^{E_1}_{A/k} \to A \otimes_k A \to A$$

of bimodules ([13, Theorem 7.3.5.1] for $k = 1$). As $A \otimes_k A$ is the free bimodule, we also get that algebra maps $A \to A \oplus N$ lifting the identity are the cofiber of $\mathrm{map}_{A \otimes_k A^{\mathrm{op}}}(A, N) \to N$. The first term here is also known as Hochschild cohomology with coefficients: $HH^*(A/k; N)$ (of course, we mean the complex or spectrum here rather than just the cohomology groups, but since a bare HH could also refer to Hochschild homology, we will keep the $(-)^*$ in the notation).

In our case, this means that the obstruction to lifting the E_1-algebra M/τ^n over the square-zero extension $R/\tau^{n+1} \to R/\tau^n$ with fiber $\Sigma^{0,-n}R/\tau$ is an E_1-algebra map

$$M/\tau^n \to M/\tau^n \oplus \Sigma^2(\Sigma^{0,-n}R/\tau) \otimes_{R/\tau^n} M/\tau^n \simeq M/\tau^n \oplus \Sigma^{2,-n}M/\tau,$$

and hence lives in π_0 of the cofiber of a map

$$HH^*((M/\tau^n)/(R/\tau^n); \Sigma^{2,-n}M/\tau) \to \Sigma^{2,-n}M/\tau,$$

where the left term through some base-change formulas for Hochschild cohomology identifies with

$$\Sigma^{2,-n}HH^*((M/\tau)/(R/\tau); M/\tau) \simeq \Sigma^{2,-n}HH^*((M/\tau)/(R/\tau)).$$

Under the identification of $\mathrm{Mod}_{\mathbb{1}/\tau}(\mathrm{FilSp}) \simeq \mathrm{gr}\,\mathrm{Sp}$, both R/τ and M/τ correspond to Σ^*MU_* and Σ^*MU_*/x. A standard Hochschild cohomology computation shows that

$$HH^*((MU_*/x)/MU_*).$$

is concentrated in even degrees, the above cofiber is along a map which is surjective on homotopy groups since MU_*/x is a commutative MU_*/x-algebra, and so the cofiber is concentrated purely in odd homotopical degrees (the first coordinate of the grading). So the obstruction for lifting M/τ^n to a R/τ^{n+1} algebra vanishes.

Since any such lift $\widetilde{M}$ sits in a cofiber sequence

$$\Sigma^{0,-n} M/\tau \to \widetilde{M} \to M/\tau^n,$$

which shows that $\pi_{2k,2k}(\widetilde{M}) \to \pi_{2k,2k}(M/\tau^n)$ is injective, we also see that $\widetilde{x} \in \pi_{2k,2k}(R)$ still maps to 0 in $\pi_{2k,2k}(\widetilde{M})$. This gives a map $(R/\tau^{n+1})/\widetilde{x} \to \widetilde{M}$ base-changing to the canonical equivalence $(R/\tau^n)/\widetilde{x} \to M/\tau^n$, and so the underlying R/τ^{n+1}-module of $\widetilde{M}$ can still be identified with M/τ^{n+1}. □

Now recall $\pi_* MU \cong \mathbb{Z}[x_1, x_2, \ldots]$ with $|x_i| = 2i$. Recall that

$$\mathrm{Ext}_{MU_*MU}(MU_*, MU_*/I_n) \cong \mathbb{F}_p[v_n^{\pm 1}],$$

so MU_*/I_n contains an element v_n in the kernel of $\eta_L - \eta_R$. It turns out that modulo I_n and x_i for $i < p^n - 1$, this v_n agrees with a unit multiple of x_{p^n-1}. Conversely, there exists a unit multiple of $v_n \in MU_*/I_n$ whose image in $MU_*/(I_n, x_1, \ldots, x_{p^n-2})$ is exactly x_{p^n-1}, and hence it admits a preimage in MU_* which is of the form $x_{p^n-1} +$ decomposables. We will abusively call this element of MU_* also v_n (justified by the fact that its image in MU_*/I_n is at least a unit multiple of v_n). Note that we had never fixed v_n more precisely than up to multiplication with a unit mod p (and adding elements from I_n) anyways.

Next we observe that if we replace x_{p^n-1} by the just defined element $v_n \in \pi_* MU$, the elements x_i still form a set of polynomial generators (in general, we may change the generators of a graded polynomial ring on positive-degree generators by adding decomposables). So we may change the generators of $\pi_*(MU) = \mathbb{Z}[x_1, \ldots]$ in such a way that x_{p^n-1} agrees with v_n mod I_n, and we will do this from now on, for all p at the same time. In order to lessen confusion, we will write $v_{p,n}$ for the v_n at the prime p.

It is worth pointing out that there are other conventions for fixing generators x_i for $\pi_*(MU)$ and elements v_n (in terms of the interpretations through formal group laws), and we make no claim that these conventions are compatible with the choices made above. (For example, the Hazewinkel or Araki generators v_n obtained from the theory of p-typical formal group laws probably don't lift to $\pi_*(MU)$, only to $\pi_*(MU_{(p)})$.)

Definition 14.4. Fix p. For $n > 0$, we define an associative MU-algebra

$$K(n) = MU[v_{p,n}^{-1}] \otimes_{MU} \bigotimes_{i \neq p^n-1} MU/x_i,$$

where $MU[v_{p,n}^{-1}]$ inherits a unique commutative MU-algebra structure as localisation, and MU/x_i is endowed with a (non-unique) associative MU-algebra structure through Lemma 14.1.

For $n = 0$, we set $K(0) = H\mathbb{Q}$ (independent of the prime p).

It is customary to just write $K(n)$, making the dependence on p implicit. Similar to the elements v_n, we will sometimes encounter statements where that would lead to confusion, in which case we will write $K(p, n)$ instead.

Remark 14.5. A priori, $K(n)$ depends not only on p and n but also on the choice of generators x_i. The MU-algebra structure depends furthermore on choosing one of the uncountably many algebra structures on each MU/x_i provided by the obstruction theory. It is known that the underlying spectrum of $K(n)$ does not depend on any of these choices, and that the E_1-algebra structure definitely does, see [3]. (It is not clear to us how unique the MU-module structure is.)

The $K(n)$ give (co)homology theories which are much better behaved than MU-homology, since $\pi_* K(n)$ is very simple: It is just $\mathbb{F}_p[v_n^{\pm 1}]$. A peculiar property of these graded rings is that any graded module over them is free, so they are called *graded fields*. A consequence is the following:

Proposition 14.6.

1. *Every $K(n)$-module is of the form $\bigoplus \Sigma^{n_i} K(n)$.*
2. *We have a Künneth isomorphism $K(n)_*(X \otimes Y) \cong K(n)_* X \otimes_{\pi_* K(n)} K(n)_* Y$.*

Proof. The first statement follows immediately from Lemma 7.3, the second from Corollary 7.4. □

This makes the homology theories $K(n)_*$ much easier to compute with than MU_*. Somewhat surprisingly, they still detect nilpotence:

Theorem 14.7 (Nilpotence theorem, $K(n)$ version, smash form)**.**

1. *If $f : X \to Y$ is a map of spectra with X compact, and $K(p,n)_*(f)$ and $(H\mathbb{F}_p)_*(f)$ are zero for all p and n, then there exists k with $f^{\otimes k} : X^{\otimes k} \to Y^{\otimes k}$ nullhomotopic.*
2. *If $f : X \to Y$ is a map of compact p-local spectra (compact objects in $\mathrm{Mod}_{\mathbb{S}_{(p)}}(\mathrm{Sp})$), $K(p,n)_*(f)$ is zero for all n, and $(H\mathbb{F}_p)_*(f)$ is zero, then there exists k with $f^{\otimes k} : X^{\otimes k} \to Y^{\otimes k}$ nullhomotopic.*

As in the MU case, there is also a ring form:

Theorem 14.8 (Nilpotence theorem, $K(n)$ version, ring form)**.**

1. *If R is a ring spectrum and $\alpha \in \pi_*(R)$ an element which has zero image in $K(p,n)_*(R)$ for all p, n and $(H\mathbb{F}_p)_*(R)$ for all p, then α is nilpotent.*
2. *If R is a p-local ring spectrum and $\alpha \in \pi_*(R)$ an element which has zero image in $K(n)_*(R)$ and $(H\mathbb{F}_p)_*(R)$, then α is nilpotent.*

Analogously to Proposition 13.4, one shows that these are equivalent to each other. We also observe the following useful sharpening: In the smash form, if X and Y are both compact spectra, then the condition on $H\mathbb{F}_p$ is automatic, since the Atiyah-Hirzebruch spectral sequence gives a natural iso

$$K(p,n)_* X \cong (H\mathbb{F}_p)_* X[v_n^{\pm 1}]$$

for n large enough (specifically, so that $|v_n| = 2(p^n - 1)$ is bigger than the difference between the largest and smallest dimension of cells of X, plus some small constant).

We will deduce the smash form from the MU-version of nilpotence after some helpful intermediate results. The most surprising part here is probably that the $K(n)$ only explicitly see a very small part of the generators of $\pi_* MU$. This is related to the fact that most of these also cannot contribute to $\mathrm{Ext}^{*,*}_{MU_* MU}$, as we saw previously. The key algebraic fact here is that most MU_*-modules just don't appear as $MU_* X$, i.e. that the MU_*-modules which come from $MU_* MU$-comodules are very constrained.

Theorem 14.9 (Landweber filtration theorem)**.** *Assume M is a graded $MU_* MU$ comodule which is finitely generated as MU_*-module. Then M admits a finite filtration all of whose associated graded terms have the form $MU_*/I_{p,n}$ for some n, where $I_{p,n} = (p, v_{p,1}, \ldots, v_{p,n-1})$ denotes the ideal I_n at p.*

We will not prove this, but note that it is based on (and can be deduced from) the fact that $\mathrm{Ext}_{MU_* MU}(MU_*, MU_*/I_{p,n}) \cong \mathbb{F}_p[v_n]$, for example one can prove from this inductively that all radical finitely generated ideals of $\pi_* MU$ are $I_{p,n}$ for some p and n.

This has as immediate consequence a kind of universal coefficient theorem:

Proposition 14.10. *If R is an MU-algebra such that the sequences $(p, v_{p,2}, v_{p,3}, \ldots)$ are regular in $\pi_*(R)$, then $\pi_*(R) \otimes_{MU_*} -$ preserves all exact sequences of $MU_* MU$-comodules, and for any spectrum X,*

$$R_*(X) \cong \pi_*(R) \otimes_{MU_*} MU_*(X),$$

Proof. For the first part, if

$$0 \to M_0 \to M_1 \to M_2 \to 0$$

is an exact sequence of $MU_* MU$-comodules, the failure of exactness after basechange is measured by $\mathrm{Tor}^1_{MU_*}(\pi_*(R), M_2)$. Every $MU_* MU$-comodule is a filtered colimit of subcomodules which are finitely generated as MU_*-modules (this is a general fact about comodules over Hopf algebras/algebroids), and so by Landweber's filtration theorem it suffices to know that $\mathrm{Tor}^1_{MU_*}(\pi_*(R), MU_*/I_{p,n}) = 0$ for all p and n. But this is exactly the flatness condition.

For the second statement, there is a natural transformation (from the right to the left) which we want to show is an isomorphism, since both sides commute with filtered colimits we may assume X compact. This also implies that $MU_*(X)$ is a finitely generated MU_*-module (using that MU_* satisfies a sort of "graded noetherian" condition).

By Landweber filtration, $MU_*(X)$ admits a finite filtration whose associated graded terms are all of the form $MU_*/I_{p,n}$. The regularity assumption implies that

$$\mathrm{Tor}^i_{MU_*}(R_*, MU_*/I_{p,n}) = 0 \text{ for } i > 0,$$

and so the same is true for $MU_*(X)$ replacing $MU_*/I_{p,n}$. So the Tor spectral sequence for $R \otimes_{MU} (MU \otimes X)$ degenerates and implies the desired basechange formula. □

Remark 14.11. The first part remains true if we replace $\pi_*(R)$ by any MU_*-algebra R_* in which the sequences $(p, v_{p,1}, \ldots)$ are regular, and then the right hand side in the above isomorphism still defines a homology theory, a statement known as *Landweber exactness theorem*. Since (co)homology theories can be represented by spectra, this makes it possible to lift various MU_*-modules to spectra. This is not very homotopy-coherent though, and for example it does not produce MU-modules. We will not need the Landweber exactness theorem here, and will make no further mention of it.

Definition 14.12. We say an MU-algebra R is *Landweber flat* if all the sequences $(p, v_{p,1}, \ldots)$ are regular in $\pi_*(R)$.

For a nonzero Landweber flat MU-algebra R, its *height at p* $\mathrm{ht}_p(R)$ denotes the largest n for which $\pi_*(R)/I_{p,n}$ is nonzero, or ∞ if that is the case for all n.

For example, MU itself has $\mathrm{ht}_p(MU) = \infty$ for all p, while the p-localisation $MU_{(p)}$ (obtained by inverting all primes except for p) has height ∞ at p and 0 at all other primes, and $MU_{(p)}[v_n^{-1}]$ has height n at p and 0 at all other primes.

Remark 14.13. There is a more general way to assign heights to MU-algebras (through formal group laws), and in that language what we define as "height n at p" would more accurately be expressed as "height $\leq n$ at p", and something like $K(n)$ will actually have "height exactly n". But for Landweber-flat R, there is only one sensible height condition (equivalent to the above), and we won't consider any other notion of height, so the terminology is unproblematic.

The heights of R characterize exactly for which spectra X we have $R \otimes X = 0$ (the R-acyclic spectra).

Proposition 14.14. *Assume R, R' are Landweber flat MU-algebras with $\mathrm{ht}_p(R) \geq \mathrm{ht}_p(R')$ for each p. Then:*

1. *If $X \to Y$ is a map of compact spectra such that $X \to R \otimes Y$ is nullhomotopic, then $X \to R' \otimes Y$ is nullhomotopic.*
2. *If X is any spectrum with $R \otimes X = 0$, then $R' \otimes X = 0$.*

Proof. For the first statement, observe that $X \to R \otimes Y$ is nullhomotopic if and only if $\mathbb{S} \to R \otimes X^\vee \otimes Y$ is, so we may assume $X = \mathbb{S}$ without loss of generality. Now assume $\mathbb{S} \to Y$ becomes nullhomotopic when tensored with R. Consider the image N of the MU-homology map

$$MU_*(\mathbb{S}) \to MU_*(Y).$$

N is an MU_*MU-comodule which is finitely generated as MU_*-module, and so N admits a finite filtration whose associated graded consists of $MU_*/I_{p,m}$ for various p and m. By the base-change formula from Proposition 14.10, $R_* \otimes_{MU_*} N = 0$, and since the basechange also takes the filtration through a filtration, this must mean that $R_* \otimes_{MU_*} MU_*/I_{p,m} = 0$ for all the $I_{p,m}$ appearing in the filtration on N. By the definition of height, this means $m > \mathrm{ht}_p(R)$ for each $I_{p,m}$ appearing in the filtration on N. But then also $R'_* \otimes_{MU_*} N = 0$, and so $R'_*(\mathbb{S}) \to R'_*(Y)$ is zero, so $\mathbb{S} \to Y$ becomes nullhomotopic when tensored with R'.

For the second statement, write X as filtered colimit of compact spectra X_i, indexed over some diagram I. For each i, since X_i is compact and $\operatorname{colim} R \otimes X_j \simeq R \otimes X = 0$, we find $i \to j$ in I such that $X_i \to R \otimes X_j$ is nullhomotopic. But then also $X_i \to R' \otimes X_j$ is nullhomotopic. This means that $R \otimes X \simeq \operatorname{colim} R' \otimes X_i = 0$. □

Example 14.15. $MU_{(p)}$ and

$$BP := MU_{(p)} \otimes_{MU} \bigotimes_{i \neq p^n - 1} MU/x_i$$

annihilate the same spectra, so $MU_{(p)} \otimes X = 0$ if and only if $BP \otimes X = 0$.

Here, BP can be thought of as the spectrum obtained by removing from $MU_{(p)}$ all the "inessential" generators x_i which do not contribute to the sequence v_n. For some parts of the theory, BP is convenient, since it is smaller, but one advantage of working with MU and its localisations instead is that MU is actually a commutative ring spectrum, whereas BP is known to not admit an E_∞ structure.

Example 14.16. $MU_{(p)}[v_n^{-1}]$ and $E(n) := MU_{(p)}[v_n^{-1}] \otimes_{MU} \bigotimes_i MU/x_i$, where i ranges over all positive integers except those of the form $p^m - 1$ for $m \leq n$, annihilate the same spectra.

Example 14.17. If $E(n) \otimes X = 0$, then also $E(m) \otimes X = 0$ for $m < n$.

We now want to relate $E(n) \otimes X = 0$ with the vanishing of $K(m)_*X$ for $m \leq n$. We will use the following lemma:

Lemma 14.18. *Let M be an MU-module and $\alpha \in \pi_*(MU)$. We write $M[\alpha^{-1}] := M \otimes_{MU} MU[\alpha^{-1}]$, and $M/\alpha := M \otimes_{MU} MU/\alpha$. Then for any spectrum X, the following are equivalent:*

1. *$M \otimes X = 0$.*
2. *$(M/\alpha) \otimes X = 0$ and $M[\alpha^{-1}] \otimes X = 0$.*

Proof. The first statement trivially implies the second, since both M/α and $M[\alpha^{-1}]$ are obtained by tensoring M with some MU-module. In the other direction, observe that the MU-module maps $MU \otimes X \to MU \otimes X$ given by multiplication with α have cofiber $(MU/\alpha) \otimes X$, and we can write $MU[\alpha^{-1}] \otimes X$ as colimit of a sequence of such maps. Applying $M \otimes_{MU} -$, we may interpret

$M[\alpha^{-1}]\otimes X$ as colimit of a sequence maps $M\otimes X \to M\otimes X$ which have cofibers $(M/\alpha)\otimes X$. If the latter vanishes, all these maps are equivalences, identifying

$$M \otimes X \simeq M[\alpha^{-1}] \otimes X \simeq 0.$$

$\square$

Proposition 14.19. *Fix p and n. For a spectrum X, the following are equivalent:*

1. *$E(n) \otimes X = 0$.*
2. *$K(n) \otimes X = 0$ and $E(n-1) \otimes X = 0$.*
3. *$K(m) \otimes X = 0$ for all $m \leq n$.*

Proof. We show that the first two statements are equivalent. The third statement is then obtained inductively.

To see that the first statement implies the second, observe first that $E(n)\otimes X = 0$ implies $E(n-1)\otimes X = 0$ because both are Landweber flat, the other with smaller height. $E(n)\otimes X = 0$ also implies $MU_{(p)}[v_n^{-1}]\otimes X = 0$, and so we have

$$K(n) \otimes X = K(n) \otimes_{MU_{(p)}[v_n^{-1}]} (MU_{(p)}[v_n^{-1}] \otimes X) = 0.$$

For the other direction, assume $E(n-1) \otimes X = 0$ and $K(n) \otimes X = 0$. Iterating Lemma 14.18, the vanishing of $E(n)\otimes X$ is equivalent to the vanishing of $R_\sigma \otimes X$ where σ ranges over all 2^n functions $\{0,\dots,n-1\} \to \{0,1\}$, and R_σ denotes the tensor product

$$R_\sigma = E(n) \otimes_{MU} \bigotimes_{i=0}^{n-1} \left\{ \begin{array}{l} MU/v_i \text{ if } \sigma(i) = 0 \\ MU[v_i^{-1}] \text{ if } \sigma(i) = 1 \end{array} \right\}.$$

Those R_σ where $\sigma(m) = 1$ modules over $MU_{(p)}[v_m^{-1}]$, but because this has smaller or equal height than $E(n-1)$, for those R_σ we automatically have $R_\sigma \otimes X = 0$. The remaining σ is the one with $\sigma(m) = 0$ for all m, but then $R_\sigma = K(n)$. $\square$

We can also give a similar criterion for the vanishing of $MU_{(p)} \otimes X$ itself:

Proposition 14.20. *Let $P(m) = MU_{(p)} \otimes_{MU} \bigotimes_i MU/x_i$ where i ranges over all numbers not of the form p^n-1 for $n \geq m$. So $\pi_*(P(m)) = \mathbb{Z}_{(p)}[v_m, v_{m+1}, \dots]$. Then for any spectrum X and arbitrary m, the following are equivalent:*

1. *$MU_{(p)} \otimes X = 0$.*
2. *$K(n) \otimes X = 0$ for all $n \leq m-1$ and $P(m) \otimes X = 0$.*

Proof. The first statement implies the second because all $K(n)$ and $P(m)$ are modules over $MU_{(p)}$. In the other direction, let us write (similarly to the previous proof)

$$R_\sigma = BP \otimes_{MU} \bigotimes_{i=0}^{m-1} \left\{ \begin{array}{l} MU/v_i \text{ if } \sigma(i) = 0 \\ MU[v_i^{-1}] \text{ if } \sigma(i) = 1 \end{array} \right\}.$$

for a function $\sigma : \{0, \ldots, m-1\} \to \{0, 1\}$. Here, iterated application of Lemma 14.18 shows that $BP \otimes X$ (and hence $MU_{(p)} \otimes X$) vanishes if and only if $R_\sigma \otimes X = 0$ for all σ. Now assume $K(n) \otimes X = 0$ for all $n \leq m-1$ and $P(m) \otimes X = 0$. Applying the previous proposition shows that $E(m-1) \otimes X = 0$, hence $MU_{(p)}[v_n^{-1}] \otimes X = 0$ for all $n \leq m-1$. This means that $R_\sigma \otimes X = 0$ for all σ with $\sigma(n) = 1$ for some $n \leq m-1$. The only remaining R_σ is then the one with $\sigma(n) = 0$ for all $0 \leq n \leq m-1$, but this is precisely $P(m)$. ☐

Proof of the $K(n)$-version of the smash nilpotence theorem. We do the "global" version, the p-local version follows analogously, working in $\mathrm{Sp}_{(p)} = \mathrm{Mod}_{\mathbb{S}_{(p)}}(\mathrm{Sp})$ instead.

Let $f : X \to Y$ be a map of spectra with X finite, which induces 0 on $K(p,n)$-homology for all p, n and $H\mathbb{F}_p$-homology for all p. Using dualizability of X (and Künneth), we may replace f by $\mathbb{S} \to X^\vee \otimes Y$, and hence assume that $X = \mathbb{S}$ without loss of generality.

So let $f : \mathbb{S} \to Y$ be a map which is null on $K(p,n)$-homology and $H\mathbb{F}_p$-homology. To show that it is nilpotent, we form

$$T := \mathrm{colim}(\mathbb{S} \to Y \to Y^{\otimes 2} \to \ldots)$$

where the map $Y^{\otimes k} \to Y^{\otimes k+1}$ is $f \otimes \mathrm{id}_{Y^{\otimes k}}$.

We first fix p. By assumption on f, we have that $K(n) \otimes T \simeq 0$ and $H\mathbb{F}_p \otimes T \simeq 0$. Observe that $\mathrm{colim}_m P(m) \simeq H\mathbb{F}_p$, and so

$$\mathrm{colim}_m P(m) \otimes T \simeq 0.$$

We may also write this as

$$\mathrm{colim}_{m,k} P(m) \otimes Y^{\otimes k} \simeq 0,$$

and since $\mathbb{S}$ is compact, this means that there exists m and k such that the map $\mathbb{S} \to P(m) \otimes Y^{\otimes k}$ arising from $f^{\otimes k}$ is zero. But since the k-fold composites in the colimit defining T can be expressed in terms of $f^{\otimes k}$, this then implies $P(m) \otimes T \simeq 0$. Since we also know $K(n) \otimes T \simeq 0$ for all n, it follows that $MU_{(p)} \otimes T \simeq 0$. This means that $MU \otimes T$ becomes zero when localized at any prime, but that just means it is zero (for example since localisations just localize the homotopy groups, and the analogous statement holds in abelian groups).

Finally, if $MU \otimes T = 0$, there exists k such that already $\mathbb{S} \to MU \otimes Y^{\otimes k}$ is nullhomotopic, by compactness of $\mathbb{S}$. But then $\mathbb{S} \xrightarrow{f^{\otimes k}} Y^{\otimes k}$ is nilpotent, and we are done. ☐

14.2 The thick subcategory theorem

Recall that for a stable ∞-category $\mathcal{C}$, a *thick subcategory* is a full subcategory closed under cofibers, fibers and retracts. For example, we saw that Sp^ω is the thick subcategory of Sp generated by $\mathbb{S}$.

Not every nonzero compact spectrum generates Sp^ω as thick subcategory. For example, the thick subcategory generated by $\mathbb{S}/2$ is a proper subcategory of Sp^ω, since it is contained in the full subcategory of compact spectra that become 0 when tensored with $\mathbb{S}[\frac{1}{2}]$ (and $\mathbb{S}$ is not in there). More generally, for any spectrum E, the full subcategory of Sp^ω of *E-acyclic spectra*, $X \in \mathrm{Sp}^\omega$ with $E \otimes X \simeq 0$, defines a thick subcategory of Sp^ω.

Determining thick subcategories generated by given objects can be subtle. For example, what thick subcategory is generated by $\mathbb{S}/\eta = \mathrm{cofib}(\eta : \mathbb{S}^1 \to \mathbb{S})$? Contrary to $\mathbb{S}/2$, looking at $\mathbb{S}[\eta^{-1}]$-acyclics is boring, since η is nilpotent: $\eta^4 = 0$ in $\pi_4(\mathbb{S})$, and so $\mathbb{S}[\eta^{-1}] = 0$, every spectrum is $\mathbb{S}[\eta^{-1}]$-acyclic. And in fact, we can use the same nilpotence fact to conclude that $\mathbb{S}/\eta$ generates Sp^ω as thick subcategory. Indeed, let us write $\mathrm{thick}(\mathbb{S}/\eta)$ for the thick subcategory generated by $\mathbb{S}/\eta$. We have cofiber sequences

$$\Sigma^1\mathbb{S}/\eta^n \to \mathbb{S}/\eta^{n+1} \to \mathbb{S}/\eta,$$

which inductively prove that $\mathbb{S}/\eta^n \in \mathrm{thick}(\mathbb{S}/\eta)$. But since $\eta^4 \simeq 0$, $\mathbb{S}/\eta^4 \simeq \mathbb{S} \oplus \mathbb{S}^1$, so $\mathbb{S} \in \mathrm{thick}(\mathbb{S}/\eta)$ (as retract of an object). But since $\mathbb{S}$ generates Sp^ω as thick subcategory, $\mathrm{thick}(\mathbb{S}/\eta) = \mathrm{Sp}^\omega$.

Remark 14.21. We could discuss thick subcategories of Sp (without the compactness), but this is a bit unnatural, since thick subcategories only involve finite sums. For "big" categories like Sp, more natural versions to consider are *thick ideals* or *localizing subcategories*, where one asks for closure under tensoring with arbitrary objects or closure under arbitrary colimits respectively. For Sp, these agree.

So there seems to be some relationship between thick subcategories and nilpotence. Inspired by this, let us make the following definition:

Definition 14.22. For a compact spectrum X, define its *support* to be the set of all $K(p,n)$ for which $K(p,n)_*X \neq 0$.

Theorem 14.23 (Weak thick subcategory theorem)**.** *Assume X and Y are compact spectra such that the support of Y is contained in the support of X. Then $Y \in \mathrm{thick}(X)$. In particular, two compact spectra with the same support generate the same thick subcategory.*

Proof. We deduce this from the $K(n)$ version of the smash nilpotence theorem. Observe that Künneth for $K(p,n)$ shows that $\mathbb{S} \to X^\vee \otimes X$ induces a nonzero map $K(p,n)_*\mathbb{S} \to K(p,n)_*(X^\vee \otimes X)$ for each $K(p,n)$ in the support of X. As $K(p,n)_*$ is a graded field, a nonzero map like this is automatically injective. So if we let $F \to \mathbb{S} \to X^\vee \otimes X$ be the fiber, we see that $F \to \mathbb{S}$ induces the zero map on $K(p,n)_*$ for each $K(p,n)$ in the support of X. As $K(p,n)_*Y = 0$ for $K(p,n)$ outside of the support of X by assumption, this means that the map

$$F \otimes Y \to Y$$

is zero on all $K(p,n)_*$. Since these are finite spectra, the $K(n)$ smash nilpotence theorem applies even without a condition on $H\mathbb{F}_p$-homology, and we see that $F \otimes Y \to Y$ is nilpotent. So, for some k, $F^{\otimes k} \otimes Y^{\otimes k} \to Y^{\otimes k}$ is zero.

An easy argument using that every compact spectrum is in thick($\mathbb{S}$) shows that any thick subcategory of Sp^ω is an ideal, i.e. closed under tensoring with arbitrary compact spectra. So by construction, the cofiber of $F \to \mathbb{S}$ is in thick(X), and tensoring with powers of F, the cofiber of any $F^{\otimes m} \to F^{\otimes m-1}$ is in thick(X). So also the cofiber of $F^{\otimes k} \to \mathbb{S}$ is in thick(X), and finally the cofiber of

$$F^{\otimes k} \otimes Y^{\otimes k} \to Y^{\otimes k}$$

is in thick(X). But since this map is 0, that cofiber has $Y^{\otimes k}$ as retract, so $Y^{\otimes k} \in \mathrm{thick}(X)$.

Next, dualizability of Y gives maps

$$Y \to Y \otimes Y^\vee \otimes Y \to Y$$

which exhibit Y as retract of $Y^{\otimes 2} \otimes Y^\vee$, tensoring with Y^{k-2} we see that $Y^{\otimes k-1} \in \mathrm{thick}(X)$ if $Y^{\otimes k} \in \mathrm{thick}(X)$. Inductively, we see $Y \in \mathrm{thick}(X)$. □

There is an analogous statement for thick subcategories of compact p-local spectra, where one considers only the $K(p,n)$ for a fixed p.

Theorem 14.23 already impressively restricts the thick subcategories of Sp^ω: It allows us to describe them purely in terms of a support condition, i.e. as all X which are simultaneously $K(p,n)_*$-acyclic for any subset of the spectra $K(p,n)$. (At least those thick subcategories which are generated by finitely many objects.) It turns out that $K(p,n)$-acyclicity and $K(p,m)$-acyclicity for a finite spectrum X are not independent: if $K(p,n)_*X = 0$, then $K(p,m)_*X = 0$ for any $m \leq n$, as we will prove below. So the supports of finite spectra are quite constrained.

To work towards the stronger form of the thick subcategory theorem which precisely classifies thick subcategories of Sp^ω, it will be necessary to have a ready supply of finite spectra with specific supports. In the original proof of these results, this relied on a subtle construction of a specific spectrum with $K(n)_*X \neq 0$ and $K(n-1)_*X = 0$ due to Mitchell. We will instead extract the existence of such X (and more) from the results on vanishing lines and self-maps in $\mathrm{Syn}_{\mathbb{F}_p}$.

For this, we will need a $\mathrm{Syn}_{\mathbb{F}_p}$ version of $K(n)$. Fix p, and write $k(n) = \bigotimes_i MU/x_i$, where the tensor product is over MU and ranges over all $i \neq p^n - 1$. In particular, $K(n) = k(n) \otimes_{MU} MU[v_n^{-1}$ and $k(n) = \tau_{\geq 0} K(n)$. From the cofiber sequence

$$\Sigma^{2(p^n-1)} k(n) \to k(n) \to H\mathbb{F}_p,$$

one can see that $H_*(k(n), \mathbb{F}_p) \to H_*(H\mathbb{F}_p; \mathbb{F}_p)$ is injective, with image the subalgebra of $\mathcal{A}_*$ generated by all ξ_i, and all τ_i with $i \neq n$ (for odd p), or all ζ_i for $i \neq n+1$ and ζ_{n+1}^2 (for $p = 2$). This is a description as comodule, and in fact of the form C_J of Proposition 10.15, with $J = \mathbb{Z}_{\geq 0} \setminus \{n\}$. As a special case of that Proposition, we see:

Proposition 14.24. *The first page of the Adams spectral sequence for $k(n)$ is $\mathbb{F}_p[a_n]$, with a_n in (Adams) bidegree $(2p^n - 1, 1)$.*

This clearly degenerates for degree reasons, and so we also understand the τ-Bockstein spectral sequence for $\nu(k(n))$:

Corollary 14.25.

$$\pi_{*,*}\nu(k(n))) = \mathbb{F}_p[\tau, \tilde{v_n}],$$

with $\widetilde{v_n} \in \pi_{2(p^n-1),2p^n-1}(\nu(k(n)))$.

Here, $\widetilde{v_n}$ van be thought of as $\frac{v_n}{\tau}$, and is a synthetic lift of v_n which witnesses that v_n is in positive Adams filtration. It is in the image of a corresponding element in $\nu(MU)$ under the map $\nu(MU) \to \nu(k(n))$, since also for $\nu(MU)$ the τ-Bockstein spectral sequence degenerates and $v_n \in \pi_{2(p^n-1)}(MU)$ is detected in positive Adams filtration.

Definition 14.26. We define $\widetilde{K(n)} \in \mathrm{Syn}_{\mathbb{F}_p}$ as

$$\widetilde{K(n)} = \nu(k(n))[\widetilde{v_n}] = \nu(k(n)) \otimes_{\nu(MU)} \nu(MU)[\widetilde{v_n}^{-1}]$$

This is a synthetic lift of $K(n)$ which has a vanishing line of slope $\frac{1}{2(p^n-1)}$, with $\widetilde{K(n)}/\tau$ having homotopy groups $\mathbb{F}_p[\widetilde{v_n}^{\pm 1}]$ which live exactly on that line.

Proposition 14.27. *Assume $X \in \mathrm{Syn}_{\mathbb{F}_p}^{\omega}$ has a vanishing line of slope $< \frac{1}{2(p^n-1)}$. Then $\widetilde{K(n)} \otimes X \simeq 0$. In particular, $K(n) \otimes \tau^{-1}X \simeq 0$.*

Proof. By Corollary 12.31, $\nu(k(n)) \otimes X$ has the same vanishing line as X. As we pass from $\nu(k(n))$ to $\widetilde{K(n)}$ by inverting an element of slope $\frac{1}{2(p^n-1)}$ which exceeds the slope of the vanishing line, $\widetilde{K(n)} \otimes X = 0$. □

Theorem 14.28. *For every $n \geq 1$, there exists a compact spectrum Y with the following properties:*

1. *$K(m)_*(Y) = 0$ for $m < n$.*
2. *$K(m)_*(Y) \neq 0$ for $m \geq n$.*
3. *There exists a map $\Sigma^k Y \to Y$ which induces an isomorphism on $K(n)_*(Y)$ and acts nilpotently on $K(m)_*(Y)$ for $m \neq n$.*

Proof. We will construct this by using vanishing lines and self-maps in $\mathrm{Syn}_{\mathbb{F}_p}$. To do so, assume that $X \in \mathrm{Syn}_{\mathbb{F}_p}^{\omega}$ has a minimal vanishing line of slope $\mathrm{slope}(X)$, and assume that $v : \Sigma^{|v|}X \to X$ is a self-map parallel to that vanishing line whose cofiber has a vanishing line of strictly lower slope. Observe the following:

1. For all n with $\frac{1}{2(p^n-1)} > \mathrm{slope}(X)$, $\widetilde{K(n)}_* X = 0$.

2. For all n with $\frac{1}{2(p^n-1)} < \mathrm{slope}(X)$, $\widetilde{K(n)} \otimes X$ is a finite extension of shifts of $\widetilde{K(n)}$. In particular, it has a vanishing line of slope $< \mathrm{slope}(X)$, and so $(\widetilde{K(n)} \otimes X)[v^{-1}] = 0$. By compactness, some $v^k : X \to \Sigma^{-k|v|}\widetilde{K(n)} \otimes X$ is already nullhomotopic, and in particular some power of v acts by 0 on $\widetilde{K(n)} \otimes X$.

3. If $\frac{1}{2(p^n-1)} = \mathrm{slope}(X)$ for some n, then as $\mathrm{slope}(X/v) < \mathrm{slope}(X)$, we have $\widetilde{K(n)} \otimes (X/v) \simeq 0$, and so v acts invertibly on $\widetilde{K(n)} \otimes X$.

We now start with $X = \mathbb{S}_p^{0,0}$ and repeatedly take the cofiber of a self-map parallel to the minimal vanishing line (starting with p for the degenerate vertical case), until we reach the first X' with a vanishing line of slope $\leq \frac{1}{2(p^n-1)}$. From the first observation, this will have $\widetilde{K(m)} \otimes X' \simeq 0$ for all $m < n$. From the second observation, taking the cofiber of a self-map of slope $> \frac{1}{2(p^m-1)}$ cannot make $\widetilde{K(m)} \otimes X'$ zero, so since $\widetilde{K(m)} \otimes \mathbb{S}_p^{0,0}$ isn't zero, $\widetilde{K(m)} \otimes X'$ is nonzero for every $m \geq n$. In particular, X' has a minimal vanishing line of slope exactly $\frac{1}{2(p^n-1)}$. Finally, the third observation guarantees that the self-map parallel to that vanishing line acts invertibly on $\widetilde{K(n)} \otimes X'$. Now, inverting τ to go to Sp gives Y with the desired properties: Since each of the self-maps we killed along the way acted nilpotently on $\widetilde{K(m)} \otimes X$ for $m \geq n$, we not only know that $\widetilde{K(m)} \otimes X$ stays nonzero, but also $(\widetilde{K(m)} \otimes X)[\tau^{-1}]$. ☐

Next, we will prove that for a finite spectrum X, $K(n) \otimes X = 0$ implies $K(n-1) \otimes X = 0$. As main ingredient, we need the following, which is an easy variant of the arguments used for Propositions 14.19 and 14.20:

Lemma 14.29. *Fix p and n, and let $R = E(n)/v_{n-1} = E(n) \otimes_{MU} MU/v_{n-1}$. Then for any spectrum X, the following are equivalent*

1. $R \otimes X = 0$.

2. $K(n) \otimes X = 0$ *and* $K(m) \otimes X = 0$ *for each* $m \leq n-2$.

Proof. Tensoring R over MU with $MU_{(p)}[v_{n-2}^{-1}]$, we obtain a Landweber flat MU-algebra with height $n-2$ at p. In particular, $R \otimes X = 0$ implies $E(n-2) \otimes X = 0$ and hence $K(m) \otimes X = 0$ for every $m \leq n-2$. Tensoring R over MU with all $MU_{(p)}/v_m$ for $m < n-1$, we obtain $K(n)$, so $R \otimes X = 0$ also implies $K(n) \otimes X = 0$.

In the backward direction, for every function $\sigma : \{0, \dots, n-2\} \to \{0, 1\}$, write R_σ for

$$R_\sigma = R \otimes_{MU} \bigotimes_{i=0}^{n-2} \left\{ \begin{array}{l} MU/v_i \text{ if } \sigma(i) = 0 \\ MU[v_i^{-1}] \text{ if } \sigma(i) = 1 \end{array} \right\}.$$

Then $R \otimes X$ is equivalent to $R_\sigma \otimes X$ for every σ. Those R_σ with $\sigma(m) = 1$ for some $m \le n-2$ are modules over $MU_{(p)}[v_m^{-1}]$, their vanishing is thus implies by the vanishing of $K(i) \otimes X$ for all $i \le n-2$. The only other R_σ is the one where all $\sigma(i) = 0$, but that is just $K(n)$ itself. □

Theorem 14.30. *If X is a finite spectrum with $K(n) \otimes X = 0$, then $K(n-1) \otimes X = 0$.*

Proof. From Theorem 14.28, there exists a finite spectrum Y with $K(m) \otimes Y = 0$ for $m \le n-2$, and $K(m) \otimes Y \neq 0$ for $m \ge n-1$. In particular, by Künneth it suffices to prove $K(n-1) \otimes (X \otimes Y) = 0$ to show $K(n-1) \otimes X = 0$, and so we may assume without loss of generality not only that $K(n) \otimes X = 0$, but also $K(m) \otimes X = 0$ for $m \le n-2$. By Lemma 14.29, this means we may assume $R \otimes X = 0$ where $R = E(n)/v_{n-1}$. Equivalently, this means that v_{n-1} acts invertibly on $E(n)_* X$.

We now use the Landweber filtration theorem to see that $MU_* X$ admits a finite filtration with associated graded terms of the form MU_*/I_m for some m. By Proposition 14.10, $E(n)_* X$ inherits a finite filtration with associated graded terms of the form $E(n)_* \otimes MU_*/I_m$. These are zero for $m > n$, and nonzero for $m \le n$. If $E(n)_* X$ is nonzero, it means that there is a surjective map (of $E(n)_*$-modules)

$$E(n)_* X \to E(n)_* \otimes MU_*/I_m$$

for some $m \le n$, which we may compose further with the surjective map

$$E(n)_* \otimes MU_*/I_m \to E(n)_* \otimes MU_*/I_n$$

to assume $m = n$. But now v_{n-1} acts by 0 on $E(n)_* \otimes MU_*/I_n$ and invertibly on $E(n)_* X$, contradiction. So $E(n)_* X$ has to be zero, and in particular $K(n-1) \otimes X = 0$. □

So the supports of finite spectra are very constrained: If $K(p, n) \otimes X \neq 0$, then $K(p, m) \otimes X \neq 0$ for all $m \ge n$. The support is therefore determined completely by just recording for each p the smallest n for which $K(p, n) \otimes X \neq 0$.

Definition 14.31. For a finite spectrum X, say X *is of type n at p* if $K(p, n) \otimes X \neq 0$ but $K(p, n-1) \otimes X = 0$ (or $n = 0$). If there is no n with $K(p, n) \otimes X \neq 0$, say X has type ∞ at p.

As $K(p, 0) = H\mathbb{Q}$ for each p, if X has type 0 at some p it has type 0 at all p, and we just say X is of type 0. Since $\mathbb{S}$ also is of type 0, Theorem 14.23 shows that any X of type 0 already generates Sp^ω as thick subcategory. If X has positive type at any p, $H\mathbb{Q} \otimes X = 0$ and X has torsion homotopy groups, implying that also $[X, X]$ is finite torsion. So $n \cdot \mathrm{id}_X$ is nullhomotopic for some n, and $X_{(p)} = 0$ for p coprime to n. In particular, X has infinite height at almost all primes.

Theorem 14.32 (Thick subcategory theorem)**.** *Let P be the set of primes, and $h : P \to \mathbb{Z}_{\geq 0} \cup \{\infty\}$ a function with the property that $h(p)$ is either constant 0, or $h(p) \geq 1$ for all p.*

Then

$$\mathrm{Sp}^{\omega}_{\mathrm{ht} \geq h} := \{X \in \mathrm{Sp}^{\omega} \mid X \textit{ has height} \geq h(p) \textit{ at } p\}$$

defines a thick subcategory of Sp^{ω}, each thick subcategory is of this form, and all of them are different.

Proof. It is clear that these define thick subcategories. Now let $\mathcal{C} \subseteq \mathrm{Sp}^{\omega}$ be any thick subcategory. Define $h : P \to \mathbb{Z}_{\geq 0} \cup \{\infty\}$ by letting $h(p)$ be the minimum of heights of X at p for all $X \in \mathcal{C}$. If $h(p) = 0$ for any p, then there is X with $H\mathbb{Q} \otimes X \neq 0$, and then $\mathbb{S}$ and X generate the same thick subcategory (the entire Sp^{ω}).

Otherwise, $h(p) \geq 1$ for each p. For each p, Theorem 14.28 gives us a finite spectrum Y_p of height exactly $h(p)$. By taking the p-local part of Y_p (which, since Y_p is p-power torsion, is still finite), we may assume that the height of Y_p is ∞ at each other prime. If $X \in \mathcal{C}$ is any object of height $h(p)$ at p, the support of Y_p is contained in the support of X, and so $Y_p \in \mathrm{thick}(X) \subseteq \mathcal{C}$.

We claim $\mathcal{C}$ agrees with the thick subcategory generated by all Y_p. To see this, let $X \in \mathcal{C}$ be arbitrary. Its height is ∞ at all but finitely many primes, and $\geq h(p)$ at those other primes. In particular, the support of X is contained in $\bigoplus Y_p$, where we take a finite sum over all primes where X has finite support. It follows that X is contained in $\mathrm{thick}(\bigoplus Y_p)$, and in particular in the thick subcategory generated by all Y_p.

Since the Y_p only depended on $h(p)$, this shows that $\mathcal{C}$ is completely characterized by h. Also, every h occurs as "height function" of a thick subcategory, namely the thick subcategory generated by the corresponding objects Y_p. □

14.3 The periodicity theorem

The final part of the nilpotence story is the *periodicity theorem*, which can be regarded as a generalisation of Adams periodicity.

Definition 14.33. Fix p, and let X be a finite p-local spectrum. A *v_n-self map* is a map $f : \Sigma^k X \to X$ such that $K(n)_*(f)$ is an isomorphism, and $K(m)_*(f)$ is zero for all $m \neq n$.

Remark 14.34. Instead of asking $K(m)_*(f)$ to be zero, we could just ask $K(m)_*(f)$ to be nilpotent, and then every such f would still admit a power f^k where $K(m)_*(f^k) = 0$ for all m already. The key observation is again that $K(m)_*(X) \cong (H\mathbb{F}_p)_*(X)[v_m^{\pm 1}]$ for large enough m, so there is a uniform bound on the nilpotence order of $K(m)_*(f)$ (characterized for large m as the nilpotence order of $(H\mathbb{F}_p)_*(f)$).

Lemma 14.35. *If X has a v_n-self map, then $K(n-1) \otimes X = 0$ and X has type $\geq n$.*

Proof. If $f : \Sigma^k X \to X$ is a v_n self-map, $K(n)_* \operatorname{cofib}(f) = 0$, so $K(n-1)_* \operatorname{cofib}(f) = 0$. Since $K(n-1)_*(f) = 0$, this means that $K(n-1)_*(X) = 0$. □

From Theorem 14.28, we also know that there exists a compact p-local spectrum Y of type n which admits a v_n self-map $f : \Sigma^k Y \to Y$ (a priori only acting nilpotently on $K(m)_* Y$ for $m \neq n$, but by the above remark a power is actually a v_n self-map). We now want to show that these exist for all type n spectra, and are essentially unique. For the uniqueness, we first need an easy algebraic lemma:

Lemma 14.36. *Let M be a finitely generated graded $K(n)_*$-module, and $f : M(k) \to M$ a $K(n)_*$-module map (with some shift). Then some power f^a is just given by multiplication with a power of v_n.*

Proof. By replacing f with some power, we may assume that the degree of f is divisible by $|v_n|$. Then $v_n^{-a} f$ for suitable a is of degree 0, and it suffices to deal with that case.

Since M is a finitely generated (automatically free) $K(n)_*$-module, $\operatorname{Hom}_{K(n)_*}(M, M)$ is as well, and in particular its degree 0 part is a finite-dimensional $\mathbb{F}_p$-module. The action of f on this is through an automorphism, which has finite order. Some power of f thus acts as the identity, finishing the proof. □

Proposition 14.37. *Let X and Y be finite p-local spectra, f_X and f_Y v_n self-maps of X and Y for some $n \geq 1$. Then there exist k and l such that*

$$\begin{array}{ccc} \Sigma^{k|f_X|} X & \xrightarrow{g} & \Sigma^{l|f_Y|} Y \\ \downarrow f_X^k & & \downarrow f_Y^l \\ X & \xrightarrow{g} & Y \end{array}$$

commutes up to homotopy for all $g : X \to Y$.

Proof. Using dualizability, we can form $X^\vee \otimes Y$. Both f_X and f_Y induce self-maps on $X^\vee \otimes Y$, by $(f_X)^\vee \otimes \mathrm{id}_Y =: f_L$ and $\mathrm{id}_X \otimes (f_Y) := f_R$), and by Künneth, they are v_n self-maps of $X^\vee \otimes Y$.

By replacing them with suitable powers, we may assume that both act on $K(n)_*(X^\vee \otimes Y)$ through multiplication with the same power of v_n. Now,

$$h = f_L - f_R$$

acts by 0 on $K(m)_*(X^\vee \otimes Y)$ for each m, so it is nilpotent. We also have that f_L and f_R commute, so f_R and h commute. For any e, we then have

$$f_L^e = (f_R + h)^e = \sum_{i=0}^{e} \binom{e}{i} f_R^{e-i} h^i.$$

As X and Y have type at ≥ 1, $X^\vee \otimes Y$ is torsion, so $p^a h = 0$ for some a. Since we also have $h^b = 0$ for some b, for e a large enough power of p all terms on the right hand side are 0 except for f_R^e. So we find e with $f_L^e = f_R^e$.

Now any $g : X \to Y$ represents a map $\mathbb{S} \to X^\vee \otimes Y$, and the fact that f_L^e and f_R^e act in the same way on this translates back to commutativity of the square. □

A special case of this is "asymptotic uniqueness" of these self-maps:

Corollary 14.38. *If X is a finite p-local spectrum and f, g are two v_n self-maps, some powers of f and g agree.*

Proof. Apply the previous result to $X = Y$ and $g = \mathrm{id}_X$. □

As another corollary, we have:

Corollary 14.39. *The full subcategory of $\mathrm{Sp}^\omega_{(p)}$ on compact p-local spectra which admit a v_n self-map is a thick subcategory.*

Proof. Let $\mathcal{C}$ denote that category. We first show that $\mathcal{C}$ is closed under cofibers. Assume X and Y admit v_n self-maps and $g : X \to Y$ is any map. Then by Proposition 14.37, X and Y admit v_n self-maps f_X and f_Y which fit into a commutative square with g. The induced map $\Sigma^k \operatorname{cofib}(g) \to \operatorname{cofib}(g)$ acts by an isomorphism on $K(n)_* \operatorname{cofib}(g)$ and at least squares to 0 on $K(m)_* \operatorname{cofib}(g)$ for each $m \neq n$, so a power of it is a v_n self-map.

$\mathcal{C}$ is clearly closed under shifts, so it is also closed under fibers. It remains to show that it is closed under retracts. Assume $X \oplus Y$ has a v_n self-map, and let $e : X \oplus Y$ be the idempotent obtained as composite $X \oplus Y \to X \to X \oplus Y$. Again by 14.37, $X \oplus Y$ actually has a v_n self-map commuting with e. Taking a colimit

$$X \oplus Y \xrightarrow{e} X \oplus Y \xrightarrow{e} \dots,$$

we get a self-map on X, which is a v_n self-map. □

Now, we are ready to prove the periodicity theorem:

Theorem 14.40 (Periodicity theorem)**.** *Let X be a compact p-local spectrum. Then X admits a v_n self-map if and only if X is of type $\geq n$, which is unique up to powers.*

Proof. We have seen in Lemma 14.35 that admitting a v_n self-map implies type $\geq n$, so the thick subcategory of all compact p-local spectra which admit a v_n self-map is contained in the thick subcategory of type $\geq n$ spectra. By Theorem 14.28 there exists a Y of type exactly n with a v_n self-map (see also Remark 14.34). By the thick subcategory theorem, Y generates the thick subcategory of type $\geq n$ spectra, so we see that conversely the thick subcategory of type $\geq n$ spectra is contained in the thick subcategory of spectra which admit a v_n self-map, so they are equal. The uniqueness part is Corollary 14.38. □

References

[1] John F Adams and HR Margolis. Modules over the steenrod algebra. *Topology*, 10(4):271–282, 1971.

[2] John Frank Adams. *Stable homotopy and generalised homology*. University of Chicago press, 1974.

[3] Vigleik Angeltveit. Uniqueness of morava k-theory. *Compositio Mathematica*, 147(2):633–648, 2011.

[4] Benjamin Antieau. Spectral sequences, Décalage, and the Beilinson t-structure. *arXiv preprint arXiv:2411.09115*, 2024.

[5] Stefano Ariotta. Coherent cochain complexes and Beilinson t-structures, with an appendix by Achim Krause. *arXiv preprint arXiv:2109.01017*, 2021.

[6] J Michael Boardman. Conditionally convergent spectral sequences. *Contemporary Mathematics*, 239:49–84, 1999.

[7] Ethan S Devinatz, Michael J Hopkins, and Jeffrey H Smith. Nilpotence and stable homotopy theory i. *Annals of Mathematics*, 128(2):207–241, 1988.

[8] Rune Haugseng. Yet another introduction to ∞-categories. `https://runegha.folk.ntnu.no/naivecat_web.pdf`.

[9] Fabian Hebestreit. Algebraic and Hermitian K-theory. `https://florianadler.github.io/AlgebraBonn/KTheory.pdf`.

[10] Michael J Hopkins and Jeffrey H Smith. Nilpotence and stable homotopy theory ii. *Annals of mathematics*, 148(1):1–49, 1998.

[11] Achim Krause. *Periodicity in motivic homotopy theory and over BP_*BP*. PhD thesis, Universitäts-und Landesbibliothek Bonn, 2018.

[12] Jacob Lurie. Chromatic homotopy theory. `https://people.math.harvard.edu/~lurie/252x.html`.

[13] Jacob Lurie. Higher Algebra. `https://www.math.ias.edu/~lurie/papers/HA.pdf`, September 2017.

[14] Jacob Lurie. Higher Topos Theory, September 2017.

[15] Jacob Lurie. Kerodon. `https://kerodon.net`, 2018.

[16] John Milnor. The Steenrod algebra and its dual. *Annals of Mathematics*, 67(1):150–171, 1958.

[17] John Milnor and James Stasheff. *Characteristic classes*. Number 76. Princeton university press, 1974.

[18] Goro Nishida. The nilpotency of elements of the stable homotopy groups of spheres. *Journal of the Mathematical Society of Japan*, 25(4):707–732, 1973.

[19] John Harold Palmieri. *Stable homotopy over the Steenrod algebra*, volume 716. American Mathematical Soc., 2001.

[20] Graeme Segal. Operations in stable homotopy theory. In *New developments in topology (Proc. Sympos. Algebraic Topology, Oxford, 1972)*, volume No. 11 of *London Math. Soc. Lecture Note Ser.*, pages 105–110. Cambridge Univ. Press, London-New York, 1974.

[21] Tashi Walde. Homotopy coherent theorems of dold–kan type. *Advances in Mathematics*, 398:108175, 2022.