Spectral sequences

Sven van Nigtevecht

v0.1 — 22nd October 2024

Overview of the mini-course

It is often said that one of the big goals of homotopy theory is to compute the stable homotopy groups of spheres, $\pi_*\mathbf{S}$. However, this is not what the professionals are trying to compute, nor is it what they want to compute: many would prefer to know the ∞-page of the Adams spectral sequence for the sphere instead. The difference is that this also records a *filtration* on $\pi_*\mathbf{S}$; practically, this cuts up $\pi_*\mathbf{S}$ into bite-sized pieces. Given the enormous complexity of $\pi_*\mathbf{S}$, the helpfulness cutting it up like this should not be ignored. The downside of the ∞-page is that it does not actually tell you $\pi_*\mathbf{S}$: there are extension problems that one would have to solve. Tagging these on as additional information would be a holy grail for computational homotopy theory.

In these notes, we will argue that what this approach is actually trying to compute is the bigraded homotopy groups of the *synthetic analogue* of the sphere spectrum:

$$\pi_{*,*}(\nu\mathbf{S}).$$

Roughly speaking, we will see that this captures both the Adams filtration on $\pi_*\mathbf{S}$, as well as all extension problems. This algebra has an action of a formal element, usually denoted by τ. The $\mathbf{Z}[\tau]$-module structure is the way in which these homotopy groups capture the filtration; one of the main goals of these lectures will be to understand the meaning of τ.

This begs the question of why $\nu\mathbf{S}$ knows about the Adams filtration on the sphere. The reason is the motto that

synthetic spectra are to Adams spectral sequence
as spectra are to homology theories.

What we mean by this is that synthetic spectra are homotopical objects of which Adams spectral sequences are but an algebraic shadow. Working at the homotpical level, and only at the very end looking at this underlying algebra, has many benefits, such as

- new techniques for deducing new differentials,

- ways to construct new spectral sequences,
- obstruction theories for building new spectra.

Usually, the parameter τ is introduced alongside synthetic spectra. In this course, we study them separately. We will begin by studying τ in its natural habitat: in the category of filtered abelian groups, and the ∞-category of filtered spectra. One can think of filtered spectra as a homotopical version of *all* spectral sequences. As a result, we will see that some of these tools find their origin in filtered spectra rather than synthetic spectra. With a firm understanding of these, we can then move on to synthetic spectra: there is an adjunction

$$\mathrm{FilSp} \rightleftarrows \mathrm{Syn}$$

that transports τ into synthetic spectra. It is there that the use of τ is most powerful, and also turns Syn into a tool that is all that one needs to interact with the Adams spectral sequence.

1 Filtered abelian groups

We will begin with an elementary, algebraic concept. Note that the following terminology is nonstandard;[1] we use the term *classical* to distinguish it from the later concept of Definition 1.7.

Definition 1.1. Let A be an abelian group.

(1) A **classical filtration** on A is a sequence of subgroups

$$\cdots \subseteq F^1 \subseteq F^0 \subseteq F^{-1} \subseteq \cdots \subseteq A.$$

Let $\{ F^s \}$ be a classical filtration on A.

(2) If $a \in A$ is an element, then the **filtration** of a is the integer s such that

$$a \in F^s \quad \text{but} \quad a \notin F^{s+1}.$$

We say that a has filtration ∞ if it lies in all the F^s, and that it has filtration $-\infty$ if it is in none of the F^s.

(3) The **associated graded** of $\{ F^s \}$ is the graded abelian group $\mathrm{Gr}\, F$ given by

$$\mathrm{Gr}^s F = F^s / F^{s+1}.$$

With the above definition of the filtration of an element, then by definition, the subgroup F^s is the subgroup of elements of filtration at least s. It might therefore be helpful to think of F^s as $F^{\geqslant s}$.

[1] And open to suggestions for improvements.

The idea behind a classical filtration is that it is a tool to better understand the group A. One can think it as starting with the elements of filtration $+\infty$, and at each step, the associated graded measures how many elements we 'add' as we move down in filtration. In the end, this procedure allows us to see all the elements that do not have filtration $-\infty$. As a result, we think of elements of filtration $\pm\infty$ as bad, and hope to find ourselves in situations where they do not exist.

An example of a result that formalises this idea is the following.

Theorem 1.2. *Let A and B be abelian groups with classical filtrations $\{ F^s A \}$ and $\{ F^s B \}$, respectively. Let $f\colon A \to B$ be a map that respects these filtrations. Suppose that*

(1) *f induces an isomorphism $F^\infty A \cong F^\infty B$;*

(2) *we have $\lim_s^1 F^s A = 0$;*

(3) *f induces an isomorphism on associated graded $\operatorname{Gr}^s A \cong \operatorname{Gr}^s B$ for all s;*

(4) *both A and B have no elements of filtration $-\infty$.*

Then f is an isomorphism of abelian groups, and moreover restricts to an isomorphism $F^s A \cong F^s B$ for every s.

Proof. See [Boa99, Theorem 2.6]. ■

Remark 1.3. In the above definition, we used a *decreasing indexing* on the filtration. One should think of this as *cohomological indexing*. For most of these lectures, we will index the filtration cohomologically. The reason is that most filtrations we consider here (for example, the Adams filtration) are of the form

$$\cdots \subseteq F^2 \subseteq F^1 \subseteq F^0 = A.$$

Remark 1.4. There is an obvious variant of Definition 1.1 for graded abelian groups. In this case, the associated graded is naturally a bigraded abelian group.

Remark 1.5. In [Boa99, Section 2], the following terminology is introduced.

- If a filtration has no elements of filtration $-\infty$ (i.e., every element of A appears in one of the F^s, or equivalently, if $\operatorname{colim}_s F^s = A$), then the filtration is said to be *exhaustive*.
- If there are no elements of filtration $+\infty$ (i.e., if the limit $\lim_s F^s$ vanishes), then the filtration is said to be *Hausdorff*.
- If the derived limit $\lim_s^1 F^s$ vanishes, then the filtration is said to be *complete*. (Note that a filtration can be complete without being Hausdorff. So, to further use this terminology, the limit of a "Cauchy sequence" need not be unique.)

Warning 1.6. In this document, we will deviate from Boardman's terminology from the previous remark: see Definition 1.9

By definition, a classical filtration only grows as we move down in filtration. It turns out to be useful to allow for a more general concept, one where we allow the groups to shrink as well.

Definition 1.7.

(1) A **filtered abelian group** is a functor $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Ab}$, where we view $\mathbf{Z}$ as a poset under the usual ordering. We write $\mathrm{FilAb} := \mathrm{Fun}(\mathbf{Z}^{\mathrm{op}}, \mathrm{Ab})$ for the category of filtered abelian groups.

(2) If $A\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{Ab}$ is a filtered abelian group, then we write A^{∞} and $A^{-\infty}$ for its limit and colimit, respectively.

(3) The tensor product of abelian groups induces a presentably symmetric monoidal structure on FilAb via Day convolution, viewing $\mathbf{Z}^{\mathrm{op}}$ as a symmetric monoidal category under addition. A **filtered commutative ring** is a commutative algebra object in FilAb.

(4) If A is a filtered abelian group, then its **associated graded** is the graded abelian group $\mathrm{Gr}\, A$ given by
$$\mathrm{Gr}^s A := \mathrm{coker}(A^{s+1} \to A^s).$$

In diagrams, a filtered abelian group A consists of abelian groups A^s for $s \in \mathbf{Z}$, together with maps
$$\cdots \longrightarrow A^1 \longrightarrow A^0 \longrightarrow A^{-1} \longrightarrow \cdots .$$

This generalises the notion of a classical filtration, as follows.

- A classical filtration determines a filtered abelian group whose transition maps are injective. The only difference is that the ambient abelian group from Definition 1.1 is no longer present. We will instead view the colimit $A^{-\infty}$ as the ambient abelian group. Said differently, giving a filtration on an abelian group B now consists of providing a filtered abelian group A, together with a map $A^{-\infty} \to B$.
- Conversely, a filtered abelian group A gives rise to an **induced classical filtration** $\{ F^s \}$ on its colimit $A^{-\infty}$, via
$$F^s := \mathrm{im}(A^s \to A^{-\infty}) \subseteq A^{-\infty}. \tag{1.8}$$
This filtration has, essentially by definition, no elements of filtration $-\infty$. Note however that the assignment $A \mapsto \{ F^s \}$ loses information: the transition maps in the filtered spectrum need not be injective.

We still need to deal with the potential presence of elements of filtration $+\infty$; we will use the following terminology. The reason for this terminology is that it matches up with the notion of completeness for filtered spectra later: see **????**.

Definition 1.9. We say a filtered abelian group A is **derived complete** if

$$\lim A = 0 \qquad \text{and} \qquad \lim^1 A = 0.$$

Remark 1.10 (Filtered tensor product). Concretely, the tensor product of $A, B \in \text{FilAb}$ is given levelwise by

$$(A \otimes B)^s = \operatorname*{colim}_{i+j \geqslant s} A^i \otimes B^j,$$

with the natural transition maps between them. A filtered commutative ring is a filtered abelian group A together with pairings

$$A^s \otimes A^t \longrightarrow A^{s+t}$$

for every $s, t \in \mathbf{Z}$, satisfying the obvious commutative ring diagrams. The unit for this monoidal structure is

$$\cdots \longrightarrow 0 \longrightarrow \mathbf{Z} = \mathbf{Z} = \cdots,$$

with the first $\mathbf{Z}$ appearing in filtration 0.

We leave it to the reader to verify that the associated graded assembles to a symmetric monoidal functor

$$\text{Gr}\colon \text{FilAb} \longrightarrow \text{grAb}.$$

Remark 1.11. If a filtered commutative ring describes a classical filtration $\{ F^s \}$, then this is the same as a commutative ring structure on the ambient abelian group, such that for all s, t, we have

$$F^s \cdot F^t \subseteq F^{s+t}.$$

Note that this means that the filtration might "jump": a product in F^{s+t} might land in the subgroup F^N for $N > s + t$. In other words, filtration is *subadditive* under multiplication. See Example 1.17 below for an example.

We introduced a classical filtration as a tool to better understand the abelian group $A^{-\infty}$. From this perspective, a filtered abelian group as in Definition 1.7 only serves as a way to give rise to its induced classical filtration via (1.8). The kernels of the maps $A^s \to F^s$ would then be considered an anomaly, because they determine the zero element in A. This is *not* the perspective we will take: we will view the entire filtration as the object of interest. To some extent, many benefits from the synthetic perspective come from remembering these additional elements.

1.1 Introducing τ

Recall that a module over the polynomial ring $\mathbf{Z}[x]$ is the same as an abelian group with an endomorphism. Using this, we can give a different description of filtered abelian groups, as follows. We reserve the letter τ as a formal variable for the polynomial ring $\mathbf{Z}[\tau]$. We view this as a graded ring by giving τ degree -1. A filtered abelian group $A\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{Ab}$ has an underlying graded abelian group, and the transition maps can be viewed as a graded $\mathbf{Z}[\tau]$-module structure on this graded abelian group. This determines a functor $\mathrm{FilAb} \to \mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{grAb})$.

Proposition 1.12. *The functor*

$$\mathrm{FilAb} \xrightarrow{\simeq} \mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{grAb})$$

is a symmetric monoidal equivalence, where we regard grAb *as having the symmetric monoidal structure* without *any signs in the swap maps.*

Mathematically, there is nothing deep about this statement. The value is in the human aspect: it can be less mentally taxing to think in terms of algebraic equations involving τ, than it is to picture the diagram that is a filtered abelian group. Even for classical filtrations this is very helpful, particularly when recording filtration jumps and hidden relations. We explore these benefits through a number of examples.

Remark 1.13. Certain properties of filtered abelian groups can be rephrased using τ.

- A filtered abelian group is a classical filtration if and only if the corresponding $\mathbf{Z}[\tau]$-module is τ-power torsion free.
- The colimit of a filtered abelian group A can be viewed as inverting τ on A. More precisely, there is a symmetric monoidal equivalence

  $$\mathrm{Mod}_{\mathbf{Z}[\tau^{\pm}]}(\mathrm{grAb}) \simeq \mathrm{Ab}$$

  given by evaluation at degree 0. We write $(-)^{\tau=1}$ for the composite[2]

  $$\mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{grAb}) \xrightarrow{\tau^{-1}} \mathrm{Mod}_{\mathbf{Z}[\tau^{\pm}]}(\mathrm{grAb}) \simeq \mathrm{Ab}.$$

  Under the equivalence between FilAb with $\mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{grAb})$, the functor $A \mapsto A^{-\infty}$ becomes the functor $A \mapsto A^{\tau=1}$.
- A filtered abelian group is derived complete if and only if the corresponding $\mathbf{Z}[\tau]$-module is τ-adically complete. This is proved in the same way as Proposition 1.16 below, together with a similar identification as for the case of colimits.

[2] Alternatively, one can define $(-)^{\tau=1}$ as taking the quotient by $\tau - 1$. This explains the loss of the grading, because $\tau - 1$ is not a homogeneous element.

Definition 1.14. Let A be an abelian group. The ***p*-adic filtration** on A is the filtered abelian group

$$\cdots \xrightarrow{p} A \xrightarrow{p} A \xrightarrow{p} A = \cdots,$$

which we index to be constant from degree 0 onwards. If A is a commutative ring, then this is naturally a filtered commutative ring.

The induced classical filtration is

$$F^s = p^s A \subseteq A.$$

Note, however, that the maps in the filtration itself need not be injective, as A might contain p-torsion. The additional information in the p-adic filtration of Definition 1.14 is that it remembers all possible choices of p-divisions of elements. This filtration is a classical filtration precisely when the abelian group is p-torsion free, or in other words, when p-divisions are unique.

Example 1.15. The abelian group $\mathbf{Z}$ is p-torsion free for all p. The graded $\mathbf{Z}[\tau]$-algebra corresponding to the p-adic filtration on $\mathbf{Z}$ is

$$\mathbf{Z}[\tau, \widetilde{p}]/(\tau \cdot \widetilde{p} = p) \qquad \text{where } |\widetilde{p}| = 1.$$

We think of $\widetilde{p}$ as a refinement of $p \in \mathbf{Z}$ that records the fact that p has filtration 1.

The fact that the filtration is constant from filtration 0 onward translates to the fact that in the $\mathbf{Z}[\tau]$-module, multiplication by τ is an isomorphism in degrees zero and below. The element 1 is not τ-divisible however, reflecting the fact that the transition map from filtration 1 to filtration 0 is not surjective. The elements that are of filtration at least 1 correspond to the elements that are τ-divisible. ▲

It would have been more appropriate to consider $\mathbf{Z}_p$ instead of $\mathbf{Z}$ in the previous example, because the p-adic filtration on $\mathbf{Z}$ is not *derived* complete. In general, we have the following.

Proposition 1.16. *Let A be an abelian group. Then the p-adic filtration on A is derived complete if and only if A is p-complete as an abelian group, i.e., if the natural map*

$$A \longrightarrow A_p^\wedge := \lim_k A/p^k$$

is an isomorphism.

Proof. Consider the diagram

$$\begin{array}{ccccc}
\vdots & & \vdots & & \vdots \\
\downarrow & & \downarrow & & \downarrow \\
A & \xrightarrow{p^3} & A & \longrightarrow & A/p^3 \\
{\scriptstyle p}\downarrow & & \| & & \downarrow \\
A & \xrightarrow{p^2} & A & \longrightarrow & A/p^2 \\
{\scriptstyle p}\downarrow & & \| & & \downarrow \\
A & \xrightarrow{p} & A & \longrightarrow & A/p.
\end{array}$$

Taking limits in the vertical direction, we get an exact sequence

$$0 \longrightarrow L \longrightarrow A \longrightarrow A_p^\wedge \longrightarrow K \longrightarrow 0$$

where

$$L = \lim(\cdots \xrightarrow{p} A \xrightarrow{p} A) \qquad \text{and} \qquad K = \lim^1(\cdots \xrightarrow{p} A \xrightarrow{p} A).$$

These are precisely the limit and first-derived limit of the p-adic filtration. In other words, we see that $A \to A_p^\wedge$ is an isomorphism if and only if the p-adic filtration on A is derived complete. ■

In Remark 1.11, we remarked that in a classically filtered ring, the filtration of elements is subadditive. The corresponding $\mathbf{Z}[\tau]$-algebra records this very elegantly.

Example 1.17. Consider the ring

$$A = \mathbf{Z}[\eta, \nu]/(2\eta = 0,\ 8\nu = 0,\ 4\nu = \eta^3).$$

We give A a classical filtration by letting both η and ν be of filtration 1, and all of $\mathbf{Z}$ be of filtration 0. The relation $4 \cdot \nu = \eta^3$ is then a jump in filtration: $4 \cdot \nu$ lands in F^1, but happens to land in the smaller subgroup F^3. In particular, we do not see this relation on the associated graded.

The corresponding $\mathbf{Z}[\tau]$-algebra keeps track of this more clearly. This is the graded algebra given by

$$\mathbf{Z}[\tau, \eta, \nu]/(2\eta,\ 8\nu,\ 4\nu = \tau^2\eta^3) \qquad \text{where } |\eta| = |\nu| = 1.$$

In some sense, we were forced to insert a τ^2-term in the last relation: unlike filtered rings, graded rings do not allow for a grading-jump under multiplication. Since τ has degree -1, the relation $4\nu = \tau^2\eta^3$ now respects this rule.

Note that, unlike in Example 1.15, we do not write $\widetilde{\eta}$ or $\widetilde{\nu}$, but instead use the symbols η and ν to directly record the filtration of the elements in the ring A. We did not do this in Example 1.15, because the symbol p is usually reserved for $1 + \cdots + 1$, and it is a bad idea to break this convention. ▲

Elements that are in the kernel of the transition map now translate to elements that are τ-torsion. This will become especially important when dealing with spectral sequences: there, τ-power torsion will encode the presence of *differentials*.

Variant 1.18. There is an obvious variant of all of the above for graded abelian groups. A **filtered graded abelian group** is a functor $\mathbf{Z}^{\mathrm{op}} \to \mathrm{grAb}$. Note that this means that the transition maps preserve degrees.

We give grAb the symmetric monoidal structure with the Koszul sign rule. Again using Day convolution, we get a symmetric monoidal structure on FilgrAb.

We view $\mathbf{Z}[\tau]$ as a bigraded ring by giving τ bidegree $(0, -1)$. We give bigrAb the symmetric monoidal structure with the Koszul sign rule according to the *second* variable. This results in a symmetric monoidal equivalence

$$\mathrm{FilgrAb} \xrightarrow{\simeq} \mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{bigrAb}).$$

Here the first grading is the internal grading, and the second grading is the grading arising from the filtration. There is a sign rule for swapping elements according to their *internal grading*; the filtration does not play a role in these signs. For more motivation for this indexing convention, see. refs to later stuff

2 Filtered spectra

In a classical filtration, the associated graded tell us how the groups "grow" as we move down in filtration. In a filtered abelian group as in Definition 1.7 however, these groups can also "shrink", but this is no longer captured by the associated graded. If we want to take both of these behaviours into account at once, we need to work in the *derived* setting instead.

Remark 2.1. Although what follows is written in the language of the ∞-category of spectra, this is not strictly necessary. Readers more comfortable with chain complexes can work with these instead, replacing the cofibre sequences of spectra with short exact sequences of chain complexes (or better, with the mapping cone of a map of chain complexes), and replacing homotopy groups of spectra with homology groups of chain complexes. Short exact sequences of chain complexes induce long exact sequences on homology, and this is all we need.

Definition 2.2.

(1) A **filtered spectrum** is a functor $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$, where we view $\mathbf{Z}$ as a poset under the usual ordering. We write

$$\mathrm{FilSp} := \mathrm{Fun}(\mathbf{Z}^{\mathrm{op}}, \mathrm{Sp})$$

for the ∞-category of filtered spectra.

(2) If $X \colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$ is a filtered spectrum, then we write X^{∞} and $X^{-\infty}$ for its limit and colimit, respectively.

(3) The smash product of spectra induces a presentably symmetric monoidal structure on FilSp via Day convolution. A **filtered $\mathbf{E}_\infty$-ring** is an $\mathbf{E}_\infty$-algebra object in FilSp.

Remark 2.3. Indexed in this way, it might be better to refer to a functor $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$ as a *tower* of spectra. We will be stubborn however and continue calling it a filtration, thinking of it as a *cohomologically-indexed filtration*. This is also the reason why we use superscripts instead of subscripts for the indices.

Remark 2.4. Informally, a filtered spectrum is a diagram of spectra

$$\cdots \longrightarrow X^1 \longrightarrow X^0 \longrightarrow X^{-1} \longrightarrow \cdots .$$

In fact, this description is not far from the formal one: a filtered spectrum is uniquely determined by the spectra and the maps between them. Conversely, any collection of spectra $\{ X^n \}_n$ together with maps $\{ X^n \to X^{n-1} \}_n$ corresponds to filtered spectrum, uniquely up to equivalence. Informally, this means there are no higher coherences between the transition maps. See [Ari21, Proposition 3.3, Corollary 3.4] for a proof. As a result, it is easy to construct filtered spectra by hand. When constructing functors into filtered spectra however, it might be easier (or even necessary) to write down coherent definitions, rather than describing the values by hand.

Definition 2.5. Recall that homotopy groups of spectra form a lax symmetric monoidal functor $\pi_* \colon \mathrm{Sp} \to \mathrm{grAb}$, where grAb has the Koszul sign rule for twist maps. We therefore obtain a lax symmetric monoidal functor

$$\mathrm{FilSp} \longrightarrow \mathrm{FilgrAb}, \quad X \longmapsto \pi_* X.$$

We refer to this functor as the **(filtered) bigraded homotopy groups**, and write

$$\pi_{n,s} X = \pi_n(X^s).$$

It is more convenient to denote this using the formal parameter τ.

Variant 2.6. Using the equivalence of Variant 1.18, we can rewrite this as a lax symmetric monoidal functor

$$\pi_{*,*} \colon \mathrm{FilSp} \longrightarrow \mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{bigrAb}).$$

The action of τ on $\pi_n(X^s)$ is the map given by the effect of the map $X^s \to X^{s-1}$ on π_n.

Later we will lift τ on $\pi_{*,*} X$ to a self-map on the filtered spectrum X itself. For now, we only think of it as a way to work with the bigraded homotopy groups.

Definition 2.7. Let X be a filtered spectrum. The **associated graded** of X is the graded spectrum $\mathrm{Gr}\, X$ given by

$$\mathrm{Gr}^s X := \mathrm{cofib}(X^{s+1} \to X^s).$$

Using that cofibres are functorial, this results in a functor

$$\mathrm{Gr}\colon \mathrm{FilSp} \longrightarrow \mathrm{grSp}.$$

Notation 2.8. Generally, we will use Greek letters to denote elements in the homotopy groups of a filtered spectrum, and use Roman letters to denote elements in the homotopy groups of the associated graded.

maybe move somewhere else

As the associated graded $\mathrm{Gr}^s X$ sits, on homotopy groups, in a long exact sequence with X^{s+1} and X^s, it measures both how the transition map $X^{s+1} \to X^s$ grows *and* shrinks the homotopy groups. This makes it even more useful for understanding the homotopy groups of the colimit.

The first step is to cut this task up into steps, by introducing a classical filtration on these homotopy groups.

somewhere: motivation for filtered spectra, saying that map of filtered spectra is an iso iff it's an iso on the limit and the associated graded

Definition 2.9. Let X be a filtered spectrum. The **induced (classical) filtration** on $\pi_n X^{-\infty}$ is the filtration given by

$$F^s\, \pi_n X^{-\infty} := \mathrm{im}(\pi_n X^s \to \pi_n X^{-\infty}).$$

Note that, because $\pi_*\colon \mathrm{Sp} \to \mathrm{grAb}$ preserves filtered colimits, this is the same as the classical filtration on $\pi_n X^{-\infty}$ induced by the filtered abelian group $\pi_n X$. In particular, the classical filtration on $\pi_n X^{-\infty}$ has no elements of filtration $-\infty$.

In practise, it is not easy to compute the homotopy groups $\pi_n X^s$ directly, so we should not compute this filtration from the definition. What is usually much more accessible is the associated graded of the filtered spectrum, but carries considerably less information. One might try and invest the associated graded with as much structure as possible, so that it starts to remember the homotopy of the filtered spectrum itself. This is precisely what a *spectral sequence* does.

3 Spectral sequences

Before giving a formal definition, we begin with a more informal diagram chase to illustrate the inner workings of a spectral sequence. This is meant simultaneously as an introduction to spectral sequences from the perspective of filtered spectra, as well as to explain and motivate our indexing conventions. In particular, we will see that the natural choice of indexing results in what is commonly known as *Adams grading*.

you make a similar remark down the line. merge or do something else to them

Remark 3.1. Not every spectral sequence arises from a filtered spectrum, but in practise they do. In these notes, we will for all intents and purposes equate the notion of a filtered spectrum and a spectral sequence.

Remark 3.2. It is also possible to orient the situation so that a filtration lets us compute the homotopy groups of the limit rather than the colimit. This present a few more technical hurdles to overcome. To some extent, one can flip between the two situations; we discuss this in **??**.

boardman ref?

remark about generality. it should all generalise to working in a general stable cat, with htpy replaced by mapping out of a compact object (to ensure that filtered colimits are fine), or more generally something where sequential colimits are fine. Then also mention Lurie's setup using a t-structure, which I guess is slightly different but related.

Before formalising this approach and making it functorial, we give a more hands-on description of what a spectral sequence is trying to do. For simplicity, and since this covers most of our use cases, throughout this section we only consider the case where the filtered spectrum is constant after degree 0:

$$\cdots \longrightarrow X^2 \longrightarrow X^1 \longrightarrow X^0 \xrightarrow{\cong} X^{-1} \xrightarrow{\cong} \cdots .$$

We will as a result simply ignore the spectra in negative filtration. Our goal then is to understand $\pi_* X^0$.

We may do this one degree at a time, so henceforth we fix an integer n. Hitting the above diagram with the functor π_n, we obtain a diagram of abelian groups:

$$\cdots \longrightarrow \pi_n X^2 \longrightarrow \pi_n X^1 \longrightarrow \pi_n X^0.$$

This filtered abelian group induces a classical filtration on $\pi_n X^0$, and this is what we aim to understand. Our first job then should be to understand when an element in $\pi_n X^0$ is in the image of $\pi_n X^1$; in other words, determine which elements have filtration at least 1.

3.1 Using the associated graded

We have a cofibre sequence

$$X^1 \longrightarrow X^0 \longrightarrow \operatorname{Gr}^0 X,$$

leading to a long exact sequence

$$\cdots \longrightarrow \pi_n X^1 \longrightarrow \pi_n X^0 \longrightarrow \pi_n \operatorname{Gr}^0 X \longrightarrow \cdots .$$

This allows us to test whether $\alpha \in \pi_n X^0$ has filtration at least 1: this happens if and only if it goes to zero under the map $\pi_n X^0 \to \pi_n \operatorname{Gr}^0 X$. This pattern continues: if $\alpha \in \pi_n X$ has filtration at least 1, we then ask if it filtration is at least 2. Choosing a lift to $\pi_n X^1$, we look at the associated graded $\operatorname{Gr}^1 X$, whose homotopy sits in a long exact sequence

$$\cdots \longrightarrow \pi_n X^2 \longrightarrow \pi_n X^1 \longrightarrow \pi_n \operatorname{Gr}^1 X \longrightarrow \cdots ,$$

and we can iterate this procedure until α does not lift further, at which point we have determined the filtration of α.

This way of thinking only goes so far: it presupposes that we understand the elements of $\pi_n X^s$, which we usually do not. In practise, what is more understandable is the homotopy of the associated graded. Instead of starting with the $\pi_n X^s$, we will start with the groups $\pi_n \operatorname{Gr}^s X$ for all s, and then try to piece the $\pi_n X^s$ back together from this data. This presents two issues:

(1) not every element in $\pi_n \operatorname{Gr}^s X$ comes from $\pi_n X^s$ (in other words, there are "fake elements"),

(2) even if an element in $\pi_n \operatorname{Gr}^s X$ lifts to $\pi_n X^s$ (in other words, it is not "fake"), then it may map to zero in $\pi_n X^0$.

We can solve both of these issues using the same mechanism. We equip the homotopy of the associated graded with more information that will make it "remember" the homotopy of the filtered spectrum. This additional information comes in the form of self-maps on the associated graded, known as the *differentials*. Concretely, a differential will connect a "fake" element to an element that maps to zero under (a composite of) the transition maps. As a result, we see that the purpose of these "fake" elements is to introduce *relations* in the homotopy of $\pi_* X^0$.

3.2 Differentials: obstructions to lifting

First, let us address issue (1). For this, we use the long exact sequence

$$\cdots \longrightarrow \pi_n X^{s+1} \longrightarrow \pi_n X^s \longrightarrow \pi_n \operatorname{Gr}^s X \longrightarrow \pi_{n-1} X^{s+1} \longrightarrow \cdots,$$

which tells us that an element in $\pi_n \operatorname{Gr}^s X$ comes from $\pi_n X^s$ if and only if it maps to zero in $\pi_{n-1} X^{s+1}$. The question, then, is how explicit we can make this condition, where 'explicit' refers to describing it in terms of the associated graded as much as possible. It would also be helpful to organise this information in a digestible way.

To make notation easier, we will focus on the case $s = 0$. Our situation is summarised by the diagram

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X^2 & \longrightarrow & X^1 & \longrightarrow & X^0 \\ & & & & & \overset{\partial}{\nwarrow} & \downarrow \\ & & & & & & \operatorname{Gr}^0 X, \end{array}$$

where the dashed arrow indicates that the map is of degree 1: it is the boundary map $\partial\colon \operatorname{Gr}^0 X \to \Sigma X^1$ of a cofibre sequence. By exactness, an element $a \in \pi_n \operatorname{Gr}^0 X$ comes from $\pi_n X^0$ if and only if its image in $\pi_{n-1} X^1$ is zero. However, like we said before, we usually do not know much about the homotopy groups of X^1, so this is not a helpful description. To approximate the question of the image $\partial a \in \pi_{n-1}X^1$ being zero, we first ask if its image in the associated graded $X^1 \to \operatorname{Gr}^1 X$ is zero:

$$
\begin{array}{ccccccc}
\cdots & \longrightarrow & X^2 & \longrightarrow & X^1 & \longrightarrow & X^0 \\
 & & & & \downarrow & \overset{\partial}{\nwarrow} & \downarrow \\
 & & & & \operatorname{Gr}^1 X & & \operatorname{Gr}^0 X.
\end{array}
$$

Write $d_1(a)$ for the image of ∂a in $\pi_{n-1}\operatorname{Gr}^1 X$. If this is nonzero, then $\partial a \neq 0$ as well, so in particular we learn that a is not in the image of $\pi_n X^0$. If $d_1(a)$ is zero however, then we are still not done: all we know is that $\partial a \in \pi_{n-1}X^1$ lifts to $\pi_{n-1}X^2$. Choosing a lift, we can ask the same question, testing whether this element is zero by looking at its image in $\pi_{n-1}\operatorname{Gr}^2 X$:

$$
\begin{array}{ccccccc}
\cdots & \longrightarrow & X^2 & \longrightarrow & X^1 & \longrightarrow & X^0 \\
 & & \downarrow & & & \overset{\partial}{\nwarrow} & \downarrow \\
 & & \operatorname{Gr}^2 X & & & & \operatorname{Gr}^0 X.
\end{array}
$$

This choice will not be unique, and neither will the resulting class in $\pi_{n-1}\operatorname{Gr}^2 X$; the class in $\pi_{n-1}\operatorname{Gr}^2 X$ is only well defined up to the image of d_1. We write $d_2(a)$ for this element in $(\pi_{n-1}\operatorname{Gr}^2 X)/d_1$. If $d_2(a)$ is nonzero, then ∂a is also nonzero. If $d_2(a)$ is zero, then we continue the story, and can define $d_3(a)$ in $\pi_{n-1}\operatorname{Gr}^3 X$ (only well defined up to d_1 and d_2), et cetera.

We obtain inductively defined elements $d_r(a)$ for $r \geqslant 1$. If they all vanish, then our class a lifts to an element of the limit $\lim_s X^s$ of the filtered spectrum. This gets us into the tricky waters of convergence issues. In good situations, this limit vanishes; let us assume this is the case. This is good news: it means that we can detect if $\partial a \in \pi_{n-1}X^1$ is zero by checking if the $d_r(a)$ are zero for all $r \geqslant 1$. This, in turn, means that we can answer the question whether $a \in \pi_n \operatorname{Gr}^0 X$ comes from $\pi_n X$.

In summary then: we have an inductively defined list of *differentials* $d_r(a)$, which (in good cases) vanish if and only if a comes from an element in $\pi_n X$. While so far we only started with classes in $\pi_n \operatorname{Gr}^0$, the same applies when starting with an element of $\pi_n \operatorname{Gr}^s$, which lifts to $\pi_n X^s$ if and only if all differentials on it vanish.

3.3 Kernels of transition maps

On to issue (2), which is asking what the kernel of $\pi_n X^s \to \pi_n X^0$ is. Since the map $X^s \to X^0$ is a composite of s maps, we can focus on the map $X^s \to X^{s-1}$ and iterate this procedure. Here we will encounter some of the awkwardness of working solely in terms of the associated graded. To illustrate this, we start with an element $\alpha \in \pi_n X^s$, and write a for its image in $\pi_n \operatorname{Gr}^s X$. Our aim is to understand whether α maps to zero in $\pi_n X^{s-1}$. We have a long exact sequence

$$\cdots \longrightarrow \pi_{n+1}\operatorname{Gr}^{s-1} X \longrightarrow \pi_n X^s \longrightarrow \pi_n X^{s-1} \longrightarrow \cdots,$$

so by exactness, α maps to zero in $\pi_n X^{s-1}$ if and only if it is in the image of the map $\pi_{n+1} \operatorname{Gr}^s X \to \pi_n X^s$. Notice that in terms of a, this is equivalent to the existence of an element $b \in \pi_{n+1} \operatorname{Gr}^{s-1}$ such that $d_1(b) = a$. Iterating this procedure, we find that if the element $\alpha \in \pi_n X^s$ maps to a nonzero element in $\pi_n X$ if and only if a is not in the image of $d_1, \ldots, d_s$.

Remark 3.3. It might appear there is an asymmetry in the above: to resolve issue (1), we had to check that infinitely many differentials on a vanish, whereas for issue (2) we only have to check a condition involving finitely many differentials. This is due to the simplifying assumption we made earlier that the filtered spectrum is constant after filtration 0. This is equivalent to the associated graded being zero in negative filtrations. In effect, this means that differentials originating in filtration below 0 automatically vanish, so that the condition of not being hit by them is vacuous.

3.4 Choosing lifts

Phrasing this solely in terms of the associated graded runs into some slightly delicate matters. By this we mean that we do not start with a class $\alpha \in \pi_n X^s$, but only with a class $a \in \pi_n \operatorname{Gr}^s X$. If $d_1(a) = \cdots = d_s(a) = 0$, then any lift of a to $\pi_n X^s$ will map to a nonzero element in $\pi_n X^0$. However, if $d_r(a) = 0$ for some $r \leqslant s$, then we only learn that *there exists* a lift of a to $\pi_n X^s$ that will map to zero in $\pi_n X^{s-r}$. It is not guaranteed that every lift will satisfy this: if $\alpha \in \pi_n X^s$ is a lift of a, then for any $\beta \in \pi_n X^s$ that comes from $\pi_n X^{s+1}$, the element $\alpha + \beta$ also lifts a. But the associated graded has no control over β: it maps to zero in $\pi_n \operatorname{Gr}^s X$, so from the point of view of a, it is invisible. In fact, β need not even map to zero in $\pi_n X^0$. The summary then is that the associated graded $\pi_n \operatorname{Gr}^s X$ only sees phenomena *up to higher filtration*.

This is a problem that we simply have to live with if all we understand is the associated graded. It can be delicate matter to check that a lift of an element hit by a d_r-differential is the lift that dies r filtrations down. In practise, one might be able to bootstrap this together by comparing different spectral sequences: in one spectral sequences, there might be no elements of higher filtration, so that there are no problems choosing a lift. This choice can then be transported to different spectral sequence where it is not clear how to choose this lift. We will see examples of this later.

3.5 Graphical depiction of spectral sequences

At this point, we need a way to organise all of this information in a way to make it more approachable for humans. Define

$$\mathrm{E}_1^{n,s} := \pi_n \operatorname{Gr}^s X,$$

and define, for every n, s, the first differential

$$d_1 \colon \mathrm{E}_1^{n,s} \longrightarrow \mathrm{E}_1^{n-1,s+1}$$

as the boundary map $E_1^{n,s} \to \pi_{n-1} X^{s+1}$ followed by the projection $\pi_{n-1} X^{s+1} \to E_1^{n-1,s+1}$.

We depict these by letting the horizontal axis correspond to the stem n, and the vertical axis correspond to the filtration s. The differential d_1 goes one to the left, and one up.

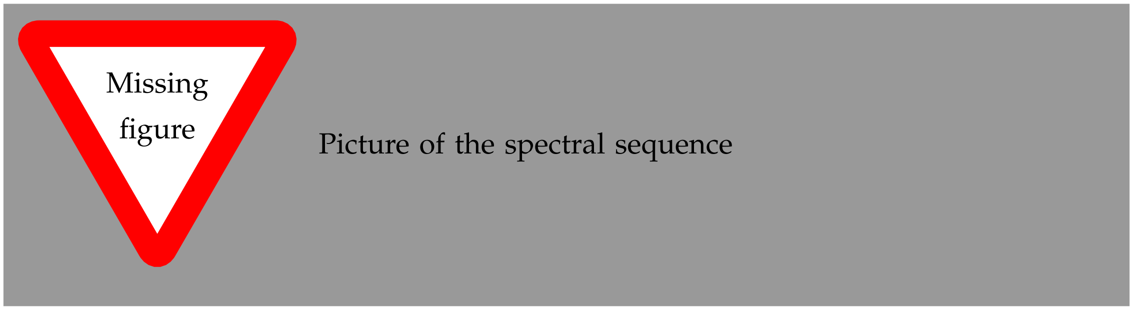

finish. probably some story about how you depict filtration vertically, and the stems you're interested in horizontally.

The differential d_r goes one to the left, and r units up. This map is however only well defined after taking homologies for the preceding differentials $d_1, \ldots, d_{r-1}$. We therefore inductively define, for $r \geqslant 2$,

$$E_r^{n,s} := H^{n,s}(E_{r-1}^{*,*}, d_{r-1}) = \frac{\ker(d_{r-1} \colon E_{r-1}^{n,s} \to E_{r-1}^{n-1,s+r-1})}{\operatorname{im}(d_{r-1} \colon E_{r-1}^{n+1,s-r+1} \to E_{r-1}^{n,s})}.$$

Roughly speaking, doing this process infinitely many times results in page ∞, denoted $E_\infty^{n,s}$. In good cases, this is isomorphic to the associated graded of the induced filtration on $\pi_n X^0$:

$$E_\infty^{n,s} \cong \frac{F^s \, \pi_n X^0}{F^{s+1} \, \pi_n X^0}.$$

In summary then: by passing from the filtered spectrum to the associated graded, we introduced "fake" elements. These elements are responsible for recording which elements die under the transition maps $\pi_n X^s \to \pi_n X^{s-1}$. Taking homology for a d_r-differential removes both the fake elements, and kills elements that die under $X^s \to X^{s-r}$. Letting all differentials run brings us to the associated graded of the filtration we were trying to understand.

3.6 Reformulation in terms of τ

introduce τ-module htpy group

The diagrams we have been looking at up till now can be difficult to keep in one's head. We can rephrase everything in terms of τ, which can help alleviate some of this difficulty. bad sentence

References

[Ari21] S. Ariotta. *Coherent Cochain Complexes and Beilinson T-Structures, with an Appendix by Achim Krause*. Comment: 66 pages, any comments are welcome! 2nd Sept. 2021. arXiv: 2109.01017v1 [math]. Preprint.

[Boa99] J. M. Boardman. 'Conditionally Convergent Spectral Sequences'. In: *Contemporary Mathematics*. Vol. 239. American Mathematical Society, 1999, pp. 49–84. DOI: 10.1090/conm/239/03597.