

An introduction to filtered and synthetic spectra

Sven van Nigtevecht

23 December 2025

Preface

These are expository notes about spectral sequences, filtered spectra, and synthetic spectra. Feedback is always appreciated, so please do not hesitate to contact me in case you find typos, corrections, or have other comments. I can make no claim to originality of the contents of these notes, except for the arrangement and presentation of the results. The baseline for these notes is a mini-course I gave at Utrecht University in the fall of 2024. The current version of these notes is essentially a lengthy excerpt from my PhD thesis, but unfortunately does not yet contain all the material I covered in the mini-course. The main subjects that are currently missing are: an in-depth example computation, Goerss–Hopkins obstruction theory, and a look at the construction of the ∞ -category of synthetic spectra (and variants thereof). I hope to include these in later versions. For the moment, I would direct readers to [CDvN24] for a sample computation, and to [PV22] for an introduction to obstruction theories.

Acknowledgements

First and foremost, I would like to thank my PhD advisor, Lennart Meier, for all of the extensive feedback he provided on everything I wrote, for the problems he suggested me to work on, and for the way he shaped my mathematical thinking. Much of the material in these notes was born out of my collaborations with Christian Carrick and Jack Davies. I would like to thank Shaul Barkan, Robert Burklund, Irakli Patchkoria, and Piotr Pstrągowski for teaching me so much synthetic mathematics, for answering my questions patiently, and for the interest they showed in my work, which was incredibly encouraging. Large portions of these notes grew out of my attempts to better understand the Omnibus Theorem of Burklund–Hahn–Senger; their work has been very influential. Many thanks to John Rognes for his many and detailed comments and corrections, and thank you as well to Miguel Barata, Maite Carli, Bastiaan Cnossen, Marius Nielsen, and Ryan Quinn for catching typos and feedback on earlier versions.

These notes were written as part of my PhD, which was supported by the NWO grant VI.Vidi.193.111. Parts of these notes were written during my stay at the Isaac Newton Institute for Mathematical Sciences, Cambridge, during the programme *Equivariant homotopy theory in context* (supported by EPSRC grant EP/Z000580/1).

Contents

1	Introduction	1
1.1	Filtered spectra	2
1.2	Synthetic spectra	4
1.3	History	7
1.4	Outline	8
1.5	Conventions	9
2	Filtered spectra and spectral sequences	11
2.1	Filtered abelian groups	11
2.2	Filtered spectra	16
2.3	Spectral sequences	22
2.3.1	Convergence	26
2.4	Digression: reflecting	31
2.5	The Adams spectral sequence	35
2.5.1	The Tot spectral sequence	35
2.5.2	The cosimplicial Adams spectral sequence	43
2.5.3	Convergence	49
3	The τ-formalism	52
3.1	Filtered abelian groups	53
3.1.1	Examples	56
3.2	Filtered spectra	58
3.2.1	Inverting τ	59
3.2.2	Modding out by τ	61
3.2.3	Completing at τ	62
3.3	Total differentials	63
3.3.1	Applications	67
3.4	Digression: the τ -Bockstein spectral sequence	69
3.4.1	The truncated τ -Bockstein spectral sequence	75
3.5	The Omnibus Theorem	78
3.5.1	The truncated Omnibus Theorem	81

3.6	Deformations	85
3.6.1	Monoidal deformations	88
4	Synthetic spectra	96
4.1	Categorical properties	97
4.2	The homological t-structure	102
4.3	Synthetic spectra as a deformation	104
4.3.1	The signature spectral sequence	106
4.3.2	Inverting τ	108
4.3.3	Modding out by τ	110
4.3.4	Synthetic lifts	113
4.4	The signature of a synthetic analogue	115
4.5	The synthetic Omnibus Theorem	122
5	Variants of synthetic spectra	126
5.1	Cellularity	126
5.2	Filtered models	128
5.3	Evenness	130
5.3.1	Even filtered spectra	131
5.3.2	Even deformations	133
5.3.3	Even synthetic spectra	134
5.4	Motivic spectra	137
A	Informal introduction to spectral sequences	141
A.1	The reconstruction problem	141
A.2	Differentials: obstructions to lifting	142
A.3	Differentials: kernels of transition maps	143
A.4	Graphical depiction of spectral sequences	145
A.5	Reformulation in terms of τ	146
	Bibliography	147

Chapter 1

Introduction

It is no exaggeration to say that the road to modern homotopy theory is paved with spectral sequences. Unfortunately, one big barrier to using spectral sequences is the abundance of indices, maps, and diagrams hiding in their definition. This might lead one to regard the construction of a spectral sequence as a black box, only to be opened in the most dire of circumstances. This is particularly unfortunate because there is much power to be had in working directly with the filtrations, much like how working with spectra has proved to be a better approach than working directly with homology theories.

We believe that this need not be so. With these notes, we set out to show that working directly with filtrations is not only possible, but practical. An important reason for this is the τ -formalism, which provides a way to off-load much of the notational headache. More than merely being convenient for working with filtrations, we show that taking τ seriously also makes it possible to understand more exotic categories of filtrations, such as *synthetic spectra*.

Our goal, then, is to give a mostly self-contained introduction to spectral sequences, filtered spectra and synthetic spectra, all through the lens of the τ -formalism. The essence of this part may be summarised by the following slogan.

*Filtered spectra are to spectral sequences,
as synthetic spectra are to Adams spectral sequences,
as spectra are to homology theories.*

We mean this in the following sense. In the second clause of each analogy, we are referring to the purely algebraic objects. These lack good categorical properties (e.g., the category of spectral sequences is not abelian, is not monoidal, etc.). Working with the objects of the first clause offers a way to remedy these problems, as these give rise to the corresponding algebraic objects, but constitute a good homotopy theory (having homotopy limits and colimits, a symmetric monoidal structure, etc.).

In the first analogy, there are some slight caveats. First, as spectra are inherently stable objects, one can only hope to capture spectral sequences arising in the stable setting,

so we ignore for these purposes spectral sequences coming from, e.g., towers of spaces. However, even within this more specialised setting, there are spectral sequences that do not arise from filtered spectra. In practise, most spectral sequences do come from a filtered spectrum, so we view this as more of a technicality. In this sense, their relationship is akin to the one between stable ∞ -categories and triangulated categories.

For synthetic spectra, the situation is the opposite: for every spectrum, its Adams spectral sequence comes from a preferred synthetic spectrum, but not every synthetic spectrum captures an Adams spectral sequence. However, far from being a technical point, this is a key feature of synthetic spectra. A general synthetic spectrum can be thought of as a *modified Adams spectral sequence*: the mere fact that it lives in the synthetic category means that it has a much closer relationship to Adams spectral sequences than a plain filtered spectrum would have.

Almost all results in this part are well known or folklore. The value, we believe, lies in having all of these results in one place. In this introduction, we give an overview of the main results, discuss some of the history of synthetic spectra, and end with a more detailed outline of these notes.

1.1 Filtered spectra

A *filtered spectrum* is a functor $X: \mathbf{Z}^{\text{op}} \rightarrow \text{Sp}$, where \mathbf{Z} denotes the poset of the integers under the usual ordering. We usually depict this as a sequence

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

We refer to the maps between the individual spectra as the *transition maps*. This comes with a notion of bigraded homotopy groups, being given by

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

We write FilSp for the ∞ -category of filtered spectra. Note, however, that this use of ∞ -categories is merely preferential. In fact, nearly all results on filtered spectra in these notes are not inherently modern in any way, and could have been obtained long ago, even right alongside the introduction of spectral sequences arising from filtered chain complexes.

Remark 1.1. Although we have written everything in terms of filtered spectra, readers who are more familiar with chain complexes can instead work with *filtered chain complexes*. The definition is the same: it is a diagram

$$\cdots \longrightarrow C^1 \longrightarrow C^0 \longrightarrow C^{-1} \longrightarrow \cdots$$

where each C^s is a chain complex, and each transition map is a map of chain complexes. (Beware, then, that the upper index is *not* the chain complex degree. Alternatively, one might write C_\bullet^s for the chain complex in position s .) There is some further change in terminology:

- ◆ Instead of working with homotopy groups, one should work with the homology groups of chain complexes.
- ◆ Instead of working with suspensions Σ , one should work with the shift operator [1].
- ◆ Instead of working with cofibres, one should work with the mapping cone of chain complexes. If the map of chain complexes is injective, then this is (quasi-isomorphic to) the quotient in the usual sense. Up to quasi-isomorphism, one can replace a filtered chain complex by one where all maps to be injective, so this is not much of a restriction.

If X is a filtered spectrum, then we let $X^{-\infty}$ denote its colimit. We think of X as a tool to understand its colimit. More precisely, the homotopy groups of $X^{-\infty}$ are given by the colimit over the transition maps:

$$\pi_n X^{-\infty} \cong \operatorname{colim}_s \pi_{n,s} X.$$

Knowing the homotopy groups $\pi_{n,s} X$ for all s is more information than knowing the group $\pi_n X^{-\infty}$. Even if one is only interested in knowing $\pi_n X^{-\infty}$, it is nevertheless a good idea to remember $\pi_{n,s} X$ for every s , along with all transition maps between them.

However, in practice, the difficulty of understanding the homotopy groups of the spectra X^s is on par with those of $X^{-\infty}$, even for cleverly chosen filtrations. Usually, the homotopy groups of the cofibres of the transition maps are much easier to compute; we refer to these cofibres as the *associated graded spectra*, and write

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

The (attempted) passage from the homotopy groups $\pi_n \operatorname{Gr}^s X$ to the homotopy groups $\pi_{n,s} X$ is exactly the structure of a *spectral sequence*.

We give an informal but in-depth introduction to how the spectral sequence arises from a filtration in Appendix A. For the moment, we give a rough indication. Fixing an integer n and applying π_n to the diagram X , we obtain a filtered abelian group. For a fixed s , certain elements in $\pi_n X^s$ may not be in the image of the map $\pi_n X^{s+1} \rightarrow \pi_n X^s$; we think of these as ‘being born’ at stage s . An element might be sent to zero after an application of a number of transition maps $\pi_n X^s \rightarrow \pi_n X^{s-r}$ (for some $r \geq 1$); we think of these as ‘dying’ at some later point. A spectral sequence encodes the event of an element being born and dying r steps to the right by a *differential* of length r .

At this point, it becomes useful to introduce some notation. We reserve the formal symbol τ , and let it act on the homotopy groups $\{\pi_{*,*} X\}_{n,s}$ via the transition maps: if $\alpha \in \pi_{n,s} X = \pi_n(X^s)$ is an element, then we define $\tau \cdot \alpha$ to be the image of α under $X^s \rightarrow X^{s-1}$. This turns $\pi_{*,*} X$ into a bigraded $\mathbf{Z}[\tau]$ -module, where τ has bidegree $(0, -1)$. By the previous discussion, this means that τ^r -torsion elements (i.e., elements that are annihilated by multiplication by τ^r) correspond to differentials of length r or shorter in

the spectral sequence. One can make the translation between the spectral sequence and the $\mathbf{Z}[\tau]$ -module $\pi_{*,*} X$ more precise; the resulting Rosetta stone is known as the *Omnibus Theorem*. Its proof essentially revolves around making the diagram chase of Appendix A very precise, and also taking into account potential convergence issues. We summarise this usage of τ and the surrounding comparison results by calling it the τ -formalism.

There are many advantages to remembering the filtration that gives rise to this spectral sequence, rather than only remembering the latter. The boundary maps $\mathrm{Gr}^s X \rightarrow \Sigma X^{s+1}$ encode information about differentials of *all* lengths. We refer to these as *total differentials*. Using these instead of the ordinary differentials leads, for instance, to a significantly strengthened version of the Leibniz rule, which we refer to as the *filtered Leibniz rule*. None of this would be possible when working with bare spectral sequences. The τ -formalism helps us in keeping track of total differentials, but is again not inherently necessary.

More than help us with internal computations, the τ -formalism also governs the structure of the ∞ -category of filtered spectra. This is particularly important for exporting the τ -formalism to other contexts, where it can become significantly more powerful. The map τ can be lifted to be a map in the ∞ -category FilSp . We can perform homotopical constructions with it, such as forming its cofibre $C\tau$. It turns out that $C\tau$ admits an \mathbf{E}_∞ -algebra structure in FilSp . Moreover, the associated graded functor

$$\mathrm{FilSp} \longrightarrow \mathrm{grSp}, \quad X \longmapsto \mathrm{Gr} X$$

can be identified with the functor

$$\mathrm{FilSp} \longrightarrow \mathrm{Mod}_{C\tau}(\mathrm{FilSp}), \quad X \longmapsto C\tau \otimes X.$$

In the main text, we discuss the general theory of exporting this formalism, through what has become known as *deformations*. For the purposes of this introduction, we will focus on our main example: the case of *synthetic spectra*. This will also allow us to discuss the differences between our account of the τ -formalism and the existing literature.

1.2 Synthetic spectra

The generality of filtered spectra makes them very useful: proving something about filtered spectra yields applications to all spectral sequences. However, their generality can also make them a little unwieldy. Say, for instance, that we are working with a particular type of spectral sequence where the first page has more structure than merely the homotopy groups of a spectrum. It would then be desirable to work in a modification of filtered spectra where this additional structure exists in the category itself. This additional structure should make the category easier to work with, making it a more powerful tool for studying that specific type of spectral sequence. In the case of Adams spectral sequences, this is exactly what the ∞ -category of *synthetic spectra* is.

Let E be a multiplicative homology theory. The E -based Adams spectral sequence tries to approximate maps between spectra by maps between their E -homology. Taking one spectrum to be a sphere, we obtain a spectral sequence that tries to compute homotopy groups. Under hypotheses on E , this is of the form

$$E_2^{n,s} \cong \text{Ext}_{E_*E}^{s,n+s}(E_*, E_*X) \implies \pi_n X,$$

where the Ext groups refer to Ext groups of E_*E -comodules, which roughly speaking is remembering E -homology (co)operations present on the E -homology of a spectrum. The original case introduced by Adams [Ada58] is the one where E is \mathbb{F}_p -homology, which to this day remains the main tool for computing stable homotopy groups of spheres. Another popular, more chromatic flavour is the case where $E = \text{MU}$, which is referred to as the *Adams–Novikov spectral sequence* (abbreviated ANSS). For general E , both the computation of the Ext groups as well as its differentials are highly nontrivial tasks. In the case $E = \mathbb{F}_2$, the state of the art in terms of computing the Ext groups is for $n + s \leq 200$ by Bruner–Rognes [BR22], and more recently for $n + s \leq 261$ by Lin [Lin23], and the state of the art in computing differentials is with almost complete information up until dimension 90 by Isaksen–Wang–Xu [IWX23], and with very recent further information going up until dimension 126 by Lin–Wang–Xu [LWX25]. In these notes, our focus is mostly on how to compute differentials.

Rather than thinking of Adams spectral sequences as filtered spectra arising in a particular way, synthetic spectra let us picture them as living in their own category. There is a stable ∞ -category $\text{Syn}_E(\text{Sp})$ of E -synthetic spectra, along with functors

$$\nu_E: \text{Sp} \longrightarrow \text{Syn}_E(\text{Sp}) \quad \text{and} \quad \sigma: \text{Syn}_E(\text{Sp}) \longrightarrow \text{FilSp},$$

called the E -synthetic analogue and signature functor, respectively. A series of computations in $\text{Syn}_E(\text{Sp})$ shows that the composite $\sigma \circ \nu_E$ is, in a precise sense, the E -Adams spectral sequence functor. We think of σ as a forgetful functor, sending a synthetic spectrum to its ‘underlying spectral sequence’.

Not only can σ be thought of as an underlying spectral sequence functor, it is also the mechanism through which we import the τ -formalism. Namely, σ is the right adjoint in an adjunction

$$\text{FilSp} \begin{array}{c} \xrightarrow{\rho} \\ \xleftarrow{\sigma} \end{array} \text{Syn}_E(\text{Sp}).$$

Synthetic spectra come with their own notion of bigraded spheres, and these are the image under ρ of filtered bigraded spheres. By adjunction therefore, if X is a synthetic spectrum, then $\pi_{*,*} X$ is captured by the filtered homotopy groups $\pi_{*,*}(\sigma X)$. Moreover, the functor ρ sends τ in FilSp to a map that is normally called τ in $\text{Syn}_E(\text{Sp})$. In this way, all results in the τ -formalism directly apply to $\text{Syn}_E(\text{Sp})$ as well. For instance, the total differentials and the Omnibus Theorem are available in this context too.

This does not mean that the use of filtered spectra does away with synthetic spectra. Rather, synthetic spectra are a natural home for Adams spectral sequences, and the τ -formalism therein is significantly more powerful. The main aspects in which the structure of $\text{Syn}_E(\text{Sp})$ is simpler than FilSp are the following.

- ◆ There is a t-structure on $\text{Syn}_E(\text{Sp})$, called the *homological t-structure*, whose heart is equivalent to the abelian category of (graded) E_*E -comodules. In this t-structure, for *every* spectrum X , the synthetic spectrum $\nu_E X$ is connective.
- ◆ The ∞ -category of $C\tau$ -modules in $\text{Syn}_E(\text{Sp})$ is equivalent to (a version of) the derived ∞ -category of (graded) E_*E -comodules.

These aspects are closely related, but not the same. Together, they can be regarded as the reason that $\text{Syn}_E(\text{Sp})$ is closely related to E -Adams spectral sequences: for instance, the Ext groups on the E_2 -page are precisely the mapping objects in the derived ∞ -category of comodules. More precisely, it is with these properties that we can compute $\sigma \circ \nu_E$ to be the Adams spectral sequence functor.

Neither of these features is present in FilSp . For instance, while $\nu_E X$ is connective in the homological t-structure, the filtered spectrum $\sigma(\nu_E X)$ is usually not connective in the standard t-structure on FilSp . This makes it easier to manipulate and study $\nu_E X$ in the synthetic context, again showing that $\text{Syn}_E(\text{Sp})$ is the natural home for Adams spectral sequences.

As another example, we have a commutative diagram

$$\begin{array}{ccc} \text{Syn}_E(\text{Sp}) & \xrightarrow{\sigma} & \text{FilSp} \\ C\tau \otimes - \downarrow & & \downarrow C\tau \otimes - \\ \text{Mod}_{C\tau}(\text{Syn}_E(\text{Sp})) & \xrightarrow{\sigma} & \text{Mod}_{C\tau}(\text{FilSp}). \end{array}$$

As explained before, the vertical functor on the right can be identified with the associated graded functor. The diagram now tells us that its composite with σ factors through $C\tau$ -modules in $\text{Syn}_E(\text{Sp})$. Because this is an ∞ -category of an algebraic nature, we can much more effectively compute there, making the associated graded of the signature more accessible.

One can use this additional structure in another way, which has become known as the *$C\tau$ -method*^[1] or *$C\tau$ -philosophy* of Gheorghe, Isaksen, Wang and Xu [GWX21; IWX23]. It has been one of the landmark advances in computational stable homotopy theory. It works as follows. Fixing a spectrum F , one can set up the $\nu_E F$ -Adams spectral sequence internal to

^[1]We warn that these terms should not be confused with what we call the τ -formalism. The latter is an overarching term, while $C\tau$ -method is a specific technique that is particularly powerful in the synthetic τ -formalism.

$\mathrm{Syn}_E(\mathrm{Sp})$ (where F can be different from E). We can push this spectral sequence along either one of the two functors

$$\mathrm{Sp} \longleftarrow \mathrm{Syn}_E(\mathrm{Sp}) \longrightarrow \mathrm{Mod}_{C\tau}(\mathrm{Syn}_E(\mathrm{Sp})).$$

Upon mapping it to spectra, we recover the ordinary F -Adams spectral sequence, while upon mapping it to $C\tau$ -modules, we obtain a purely algebraic spectral sequence, where we may essentially compute differentials by hand (or by computer). Because these two are now related to a spectral sequence in $\mathrm{Syn}_E(\mathrm{Sp})$, we can thus deduce differentials in spectra from differentials in the algebraic realm. Isaksen–Wang–Xu [IWX23] work in the case $E = \mathrm{MU}$ and $F = \mathbf{F}_2$; see Section 1.3 below for a further discussion.

A related application of this structure on synthetic $C\tau$ -modules, is to the computation of Toda brackets. Toda brackets in $\mathrm{Syn}_E(\mathrm{Sp})$ map to Toda brackets under the functor $C\tau \otimes -$, whereupon they become Massey products, which can again be computed algebraically by hand. This leads to a synthetic version of Moss’s Theorem; see [CDvN24, Section 3].

1.3 History

Our introduction of the τ -formalism is quite different from the way it arose historically. Whereas we presented it first of all as a notational device, its origins are much more complicated.

It first appeared when *C-motivic spectra* began to be used to study the stable homotopy groups of spheres [Isa19; GWX21; IWX20; IWX23]. In motivic spectra, one can run the motivic \mathbf{F}_2 -Adams spectral sequence for the motivic sphere. There is a functor $\mathrm{Sp}_{\mathbf{C}} \rightarrow \mathrm{Sp}$ from \mathbf{C} -motivic spectra to ordinary spectra called *Betti realisation*, and this turns the motivic \mathbf{F}_2 -ASS into the ordinary \mathbf{F}_2 -ASS. It turns out that the motivic version is remarkably similar to the ordinary one. First off, the motivic version is trigraded, due to motivic homotopy groups being bigraded. There is an endomorphism τ of the motivic sphere, and the motivic dual Steenrod algebra is almost the ordinary one tensored with $\mathbf{F}_2[\tau]$. Isaksen [Isa19] realised that this additional grading and the presence of τ introduce constraints, and used this to deduce motivic Adams differentials, which upon Betti realisation yield new differentials in the ordinary Adams spectral sequence. One can get further information out of this approach by combining it with the motivic Adams–Novikov spectral sequence.

Later, Gheorge–Wang–Xu [GWX21] showed that modules over $C\tau$ in (cellular, p -complete) \mathbf{C} -motivic spectra are equivalent to the derived ∞ -category of $\mathrm{BP}_* \mathrm{BP}$ -comodules. Combined with the observation that differentials in the motivic Adams–Novikov spectral sequence correspond to differentials in the ordinary Adams–Novikov spectral sequence, this led to the *C τ -method* (explained above in synthetic terms), and with it a great advancement in our understanding of the sphere spectrum; see [IWX23].

There is a certain unreasonable effectiveness of these motivic methods, since motivic spectra are inherently algebro-geometric objects, and it is not a priori clear why this algebraic geometry is connected to the Adams–Novikov spectral sequence. This was explained when different models were given for the subcategory of cellular motivic spectra (the subcategory in which these applications take place).

- ◆ Pstrągowski [Pst22] defined E -synthetic spectra, and proved that if $E = \text{MU}$, this gives a model for cellular motivic spectra (at least after p -completion).
- ◆ Gheorge–Isaksen–Krause–Ricka [GIKR21] gave a model for (p -complete) cellular \mathbf{C} -motivic spectra in terms of modules in filtered spectra.
- ◆ Burklund–Hahn–Senger [BHS22] directly compare these synthetic and filtered models.

This explains why \mathbf{C} -motivic homotopy theory was so successfully applied: synthetic spectra are by nature designed to be a good category of Adams filtrations. See also [Pst22, Remark 4.62] for a further discussion comparing synthetic spectra with [IWX23].

With the construction of synthetic spectra, it became possible to use the same type of techniques for other E as well, not just in the case $E = \text{MU}$. Burklund–Hahn–Senger [BHS23, Theorem 9.19] prove the *Omnibus Theorem* to formalise the idea that E -synthetic analogues capture the E -Adams spectral sequence. Later, Patchkoria–Pstrągowski [PP23] define a more general setting for defining synthetic categories, allowing for much more general E and for a much more general stable ∞ -category in the place of spectra.

If we only care about the applications to computations and spectral sequences rather than motivic homotopy theory as a whole, then synthetic spectra offer a more light-weight technical setup to do these computations with. In these notes, we further argue that most synthetic techniques originate in the even more light-weight context of filtered spectra. Not only does this make proofs and constructions more concrete, but it also allows for straightforward generalisations. For instance, while the Omnibus Theorem of Burklund–Hahn–Senger only applies to synthetic analogues, the version we will deduce (using the adjunction $\rho \dashv \sigma$ above) from our filtered version applies to *all* (convergent) synthetic spectra, and moreover applies in the same fashion to any (good enough) deformation. We give a more detailed comparison of the proofs of these two versions of the Omnibus Theorem in Remark 4.78.

1.4 Outline

We begin by reviewing the theory of filtered abelian groups, filtered spectra and spectral sequences without the use of τ . This is the content of Chapter 2. At this point, we do not yet introduce the τ -formalism: rather, this chapter is aimed at setting up the basic concepts and terminology to be used later on, and to make various conventions and subtleties explicit. In particular, we include a short discussion on the Adams spectral

sequence in a non-synthetic sense, to provide all the necessary background for the later chapters.

Next, in Chapter 3 we introduce the τ -formalism in the filtered setting, starting with filtered abelian groups, and afterwards in filtered spectra. Aside from discussing *total differentials*, the main goal of this chapter is to prove the *Omnibus Theorem* in the filtered setting. The device for proving this is the τ -Bockstein spectral sequence, which we introduce in this chapter as well. Finally, we end with a general discussion of *deformations*, showing how to export the τ -formalism to other ∞ -categories.

With this preparation in hand, in Chapter 4 we come to our other main topic, which is *synthetic spectra*. After reviewing the basic categorical properties, we show how the theory of deformations applies to synthetic spectra, and work this out in detail. The main goal of this chapter is to compute the signature of a synthetic analogue. After this, the synthetic Omnibus Theorem follows as a corollary. Finally, in Chapter 5 we discuss certain variants and properties of synthetic spectra, in particular the comparison between synthetic and motivic spectra.

1.5 Conventions

These notes are written mainly for a homotopical audience. We have in mind a reader who has seen spectral sequences before, but is not necessarily intimately familiar with their construction. In certain sections, familiarity with the Adams spectral sequence from a practical perspective will be useful, but is not strictly speaking required.

Throughout, we assume a working knowledge of ∞ -categories in the sense of Joyal and Lurie; the standard references are [HTT] and [HA]. We distinguish between categories and ∞ -categories: by the term *category*, we mean a $(1, 1)$ -category, while by the term *∞ -category*, we mean an $(\infty, 1)$ -category.

When a morphism in an ∞ -category admits a two-sided inverse up to homotopy, we refer to it as being an *isomorphism*, rather than an *equivalence*. This should not cause much confusion, as we do not compare the ∞ -categories we work in with a model category giving rise to them. The only exception is that we speak of an *equivalence* of ∞ -categories rather than an *isomorphism* of ∞ -categories (which, we admit, is not a fully consistent choice of terminology).

We use the term *space* for what is also referred to as an ∞ -groupoid or an *anima*. The ∞ -category of spaces is denoted by \mathcal{S} . The ∞ -category of spectra is denoted by Sp , and we write \otimes for the smash product of spectra. By the term *E_∞ -ring*, we mean an E_∞ ring spectrum.

We do not distinguish notationally between an abelian group and its corresponding Eilenberg–MacLane spectrum. For example, we do not write HZ , but simply write \mathbf{Z} . The context will clarify whether we are dealing with a spectrum or an abelian group.

If A is a graded abelian group and n is an integer, then we write $A[n]$ for the graded abelian group given by

$$(A[n])_k = A_{k-n}.$$

We will use the same formula for graded modules, comodules, etc.

We use Adams indexing for all of our spectral sequences, both in our formulas and in depicting spectral sequences. This means a d_r -differential has bidegree $(-1, r)$. Meanwhile, we use a homological-algebra indexing for Ext groups: for integers s and t , we write

$$\text{Ext}^{s,t}(M, N) = \text{Ext}^s(M[t], N).$$

In the case of Adams spectral sequences, this means we will often have the expression

$$E_2^{n,s} \cong \text{Ext}_{E_*E}^{s, n+s}(E_*, E_*X).$$

Our reasoning for this is that Adams indexing is most useful for working with spectral sequences, while the homological-algebra indexing on Ext is what one uses when computing these Ext groups.

Often, we refer to the E_r -page of a spectral sequence as its *r-th page* or as *page r*.

For ease of reference, we include here a list of the places where we make or clarify indexing conventions, sorted roughly by theme.

- ◆ Homological vs. cohomological indexing, decreasing vs. increasing filtrations, towers vs. filtrations: Remarks 2.3, 2.8, 2.9, 2.30, 2.41 and 2.54, Construction 2.31, and Definition 2.33.
- ◆ First- vs. second-page indexing: Remarks 2.40 and 3.55.
- ◆ Filtered τ -formalism: Definition 3.13, Notations 3.33, 3.35 and 3.50, Construction 3.47, and Remark 3.55.
- ◆ Synthetic indexing: Definition 4.13, Remarks 4.20, 4.26, 4.37 and 5.44, Variant 4.39, and Theorem 4.77.

Chapter 2

Filtered spectra and spectral sequences

As is well-known, filtrations give rise to spectral sequences. The goal of this chapter is to review the theory of filtrations in the stable setting and the resulting spectral sequences, as well as dealing with more subtle issues like convergence.

We begin by studying filtered abelian groups in Section 2.1. Partially we do this as a warm-up, but mainly because it is the natural structure on the homotopy groups of a filtered spectrum. After introducing and reviewing their basic category theory in Section 2.2, we discuss their relation to spectral sequences in Section 2.3. For a more relaxed introduction to how spectral sequences arise from filtered spectra, we refer to Appendix A. Next, our goal is to describe our main example: the Adams spectral sequence. This is the topic of Section 2.5, where we also include some background on spectral sequences arising from cosimplicial objects. To prepare for this, we include a short digression on a duality between filtrations and towers of spectra, which we refer to as *reflection*, in Section 2.4.

There is a myriad of sources on spectral sequences, which would be impossible to list here. We learned much of this chapter from [Ant24], [Boa99], [Hed20], [Rog12; Rog21].

2.1 Filtered abelian groups

We begin with an elementary algebraic concept. We will use the adjective *strict* (which we borrow from [Ant24]) to distinguish it from the later, more general concept of Definition 2.7.

Definition 2.1. Let A be an abelian group.

- (1) A **strict 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 strict filtration on A .

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

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

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

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

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

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 $F^{\geq s}$. Note that F^∞ is the limit $\lim_s F^s$, while $F^{-\infty}$ is the colimit $\text{colim}_s F^s$.

We regard a strict filtration on A is a tool to help us understand the group A . One can think of it as starting with the elements of filtration $+\infty$ and moving down in filtration, where at each step the associated graded is measuring how many elements we ‘add’. In the end, this procedure allows us to see all the elements that do not have filtration $-\infty$. In practice, the associated graded is what one has the most control over. 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. For a further discussion and other results in this direction, we refer to [Boa99, Section 2].

Proposition 2.2. *Let A and B be abelian groups equipped with strict filtrations $\{F^s A\}$ and $\{F^s B\}$, respectively. Let $f: A \rightarrow B$ be a map that respects these filtrations. Suppose that*

- (1) *the map f induces an isomorphism $F^\infty A \xrightarrow{\cong} F^\infty B$;*
- (2) *the first derived limit $\lim_s^1 F^s A$ vanishes;*
- (3) *the map f induces an isomorphism on associated graded $\text{Gr}^s A \xrightarrow{\cong} \text{Gr}^s B$ 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 $F^s A \xrightarrow{\cong} F^s B$ for every s .

Proof. See [Boa99, Theorem 2.6]. ■

Remark 2.3. In the above definition, we used a *decreasing indexing* on the filtration. One should think of this as *cohomological indexing* for filtrations. We follow this convention because most filtrations we consider (for example, the Adams filtration) are of the form

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

Remark 2.4. There is an obvious variant of Definition 2.1 for graded abelian groups. In this case, the associated graded is naturally a bigraded abelian group.

Remark 2.5. 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 first-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. In other words, the limit of a “Cauchy sequence” need not be unique.)

Warning 2.6. In these notes, we will deviate from Boardman’s terminology recalled in the previous remark: see Definition 2.11 below.

By definition, a strict 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 2.7.

- (1) A **filtered abelian group** is a functor $\mathbf{Z}^{\text{op}} \rightarrow \text{Ab}$, where we view \mathbf{Z} as a poset under the usual ordering. We write

$$\text{FilAb} := \text{Fun}(\mathbf{Z}^{\text{op}}, \text{Ab})$$

for the (presentable, abelian) category of filtered abelian groups.

- (2) If $A : \mathbf{Z}^{\text{op}} \rightarrow \text{Ab}$ 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}^{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 $\text{Gr } A$ given by

$$\text{Gr}^s A := \operatorname{coker}(A^{s+1} \rightarrow A^s).$$

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

$$\cdots \rightarrow A^1 \rightarrow A^0 \rightarrow A^{-1} \rightarrow \cdots .$$

We refer to these maps as **transition maps**.

Let us explain our (perhaps slightly nonstandard) terminology.

Remark 2.8 (Filtrations vs. towers). We deliberately use the term *filtration* instead of *tower* in the above. Throughout, we use the word *filtration* to indicate that we think of the colimit as the underlying object, and the limit as an error term. When we use the word *tower*, we instead regard the limit as the underlying object and the colimit as the error term. For an example of the difference, see Example 2.65.

Remark 2.9 (Homological vs. cohomological grading). Fitting with Remark 2.3, we regard the usage of \mathbf{Z}^{op} to index filtered objects as *cohomological* indexing of filtered objects. This is also why we use superscripts to indicate the index. If we instead think of these objects as towers in the sense of the previous remark, then the usage of \mathbf{Z}^{op} is a *homological* indexing convention.

Next, let us compare the notion of a filtered abelian group with that of a strict filtration as in Definition 2.1.

- ◆ A strict filtration is a special case of a filtered abelian group, namely one whose transition maps are injective. The only difference is that the ambient abelian group from Definition 2.1 is no longer present in Definition 2.7. We will instead regard the colimit $A^{-\infty}$ as the ambient abelian group. Said differently, giving a strict filtration on an abelian group B in the sense of Definition 2.1 consists of providing a filtered abelian group A in the sense of Definition 2.7 with injective transition maps, together with a map $A^{-\infty} \rightarrow B$.

Going forward, we will usually use the term *strict filtration* to refer to a filtered abelian group with injective transition maps. When we use the version of Definition 2.1, we will always check that there are no elements of filtration $-\infty$, to be consistent with the previous story.

- ◆ Conversely, a filtered abelian group A gives rise to an **induced strict filtration** $\{F^s\}$ on its colimit $A^{-\infty}$, via

$$F^s := \text{im}(A^s \longrightarrow A^{-\infty}) \subseteq A^{-\infty}. \quad (2.10)$$

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 abelian group need not be injective.

We still need to deal with the potential presence of elements of filtration $+\infty$; we will use the following terminology.

Definition 2.11. We say a filtered abelian group A is **derived complete** if

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

We will revisit this later in Section 3.1, where it is called τ -*completeness*. For a discussion without the language of τ , see [Boa99, Definition 2.7, Proposition 2.8].

Remark 2.12. The reason we call this *derived complete* is that it matches the notion of completeness for filtered spectra to be introduced in Definition 2.28 below. More specifically, by post-composing with the inclusion, a filtered abelian group A determines a functor $\mathbf{Z}^{\text{op}} \rightarrow \mathcal{D}(\text{Ab})$ to the derived ∞ -category of abelian groups. Then A is derived complete if and only if the (derived) limit of $\mathbf{Z}^{\text{op}} \rightarrow \mathcal{D}(\text{Ab})$ is zero. Because the forgetful functor $\mathcal{D}(\text{Ab}) \rightarrow \text{Sp}$ preserves limits, this is equivalent to viewing A as a functor $\mathbf{Z}^{\text{op}} \rightarrow \text{Sp}$ landing in discrete (a.k.a. Eilenberg–MacLane) spectra, and asking for the (homotopy) limit of this functor to vanish. See Remark 2.51 for an elaboration on this point.

Remark 2.13 (Filtered tensor product). Concretely, the tensor product of $A, B \in \text{FilAb}$ is given levelwise by

$$(A \otimes B)^s = \text{colim}_{i+j \geq s} A^i \otimes B^j,$$

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 $s, t \in \mathbf{Z}^{\text{op}}$, being compatible with the transition maps and satisfying the obvious commutative ring diagrams. The unit for this monoidal structure is

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

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

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

Remark 2.14 (Filtration is subadditive). Suppose that we have a filtered ring structure on a strict filtration $\{F^s\}$. Then this structure is the same as a commutative ring structure on the ambient abelian group such that for all s and 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. Products that jump in filtration become zero in the associated graded.

We introduced a strict filtration as a tool to better understand its ambient abelian group. It is suggestive then that the only purpose of a filtered abelian group as in Definition 2.7 is to give rise to its induced strict filtration via (2.10). Said differently, one might consider the kernels of the maps $A^s \rightarrow F^s$ to be an anomaly, because they determine the zero element in $A^{-\infty}$. This is decidedly *not* the perspective we will take: the entire filtration is the object of interest. Many of the benefits from the synthetic perspective come from remembering the filtration as a whole.

However, there is one downside to working with filtered abelian groups: the associated graded cannot measure the kernels of the transition maps. In order to take these into account as well, we have to move to the derived setting.

2.2 Filtered spectra

Instead of moving to the derived setting by using derived abelian groups, we immediately go to the more universal case of spectra.

Definition 2.15.

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

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

for the (presentable, stable) ∞ -category of filtered spectra.

- (2) If $X: \mathbf{Z}^{\text{op}} \rightarrow \text{Sp}$ 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}_∞ -ring** is an \mathbf{E}_∞ -algebra object in FilSp .
- (4) A **graded spectrum** is a functor $\mathbf{Z}^{\text{discr}} \rightarrow \text{Sp}$, where $\mathbf{Z}^{\text{discr}}$ is the discrete category with objects \mathbf{Z} . We write

$$\text{grSp} := \prod_{\mathbf{Z}} \text{Sp} = \text{Fun}(\mathbf{Z}^{\text{discr}}, \text{Sp}).$$

for the (presentable, stable) ∞ -category of graded spectra.

- (5) If X is a filtered spectrum, then its **associated graded** is the graded spectrum $\text{Gr } X$ given by

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

We often depict a filtered spectrum $X: \mathbf{Z}^{\text{op}} \rightarrow \text{Sp}$ as a diagram

$$\cdots \rightarrow X^1 \rightarrow X^0 \rightarrow X^{-1} \rightarrow \cdots$$

The abuse of this notation is not great: a filtered spectrum is, up to contractible choice, determined by the spectra $\{X^n\}$ together with their transition maps; see [Ari21, Proposition 3.3, Corollary 3.4].

Remark 2.16. The role of spectra in the above definition is not special. Most of the results in this section apply to a suitable stable ∞ -category in the place of spectra. For concreteness, and to prevent this chapter from becoming needlessly long, we stick to the case of spectra. For a discussion in this greater generality, we refer to [Ant24], or parts of [Hed20, Part II].

Like with filtered abelian groups, if X is a filtered spectrum, then we think of $X^{-\infty}$ as the ‘underlying spectrum’ of X .

Using that cofibres are functorial, the associated graded assembles into a functor

$$\text{Gr}: \text{FilSp} \longrightarrow \text{grSp}.$$

We will see later that this can be upgraded to a symmetric monoidal functor; see Remark 3.26. For an alternative proof, see [Hed20, Section II.1.3].

The functor $\pi_*: \text{Sp} \rightarrow \text{grAb}$ induces a functor

$$\text{FilSp} \longrightarrow \text{Fil}(\text{grAb}).$$

We can find corepresenting objects for the individual abelian groups of this functor. By Yoneda, each transition map will then be corepresented by a map between these objects, but we defer a discussion of the resulting maps to the next chapter. Although we could define these corepresenting objects by hand, it will be convenient to introduce the following functor.

Definition 2.17. We write $i: \mathbf{Z} \rightarrow \text{FilSp}$ for the functor given by

$$s \longmapsto \Sigma_+^\infty \text{Hom}_{\mathbf{Z}}(-, s).$$

By the properties of Day convolution, the functor i is naturally symmetric monoidal.

Remark 2.18. The functor i gives the ∞ -category FilSp a universal property, which says that colimit-preserving functors out of FilSp correspond to functors out of \mathbf{Z} . We discuss this, as well as the structure that such a functor puts on the target category, in detail later in Section 3.6.

Definition 2.19 (Filtered bigraded spheres). Let n and s be integers.

- (1) The **filtered (n, s) -sphere** is

$$\mathbf{S}^{n,s} := \Sigma^n i(s).$$

We refer to n as the **stem**, and to s as the **filtration**.

- (2) We write $\Sigma^{n,s}: \text{FilSp} \rightarrow \text{FilSp}$ for the functor given by tensoring with $\mathbf{S}^{n,s}$ on the left.
 (3) We write $\pi_{n,s}: \text{FilSp} \rightarrow \text{Ab}$ for the functor

$$\pi_{n,s}(-) := [\mathbf{S}^{n,s}, -].$$

Unwinding definitions, we see that $\mathbf{S}^{n,s}$ is the filtered spectrum given by

$$\cdots \longrightarrow 0 \longrightarrow \mathbf{S}^n \longleftarrow \mathbf{S}^n \longleftarrow \cdots,$$

where the first \mathbf{S}^n appears in position s . In diagrams therefore, $\Sigma^{n,s}$ is given by applying Σ^n levelwise, and by shifting the filtered spectrum s units to the left. Likewise, we see that for any filtered spectrum X , we have a natural isomorphism

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

Remark 2.20. This definition of the filtered bigraded spheres is designed to be compatible with *first-page indexing* of filtered spectra, i.e., the indexing that makes the underlying spectral sequence start on the first page. It is possible to change this to a second-page indexing (or any page); see Remark 2.40 for a further discussion.

The filtered spheres are, in fact, generators for FilSp .

Proposition 2.21.

- (1) For every $s \in \mathbf{Z}$, the filtered sphere $\mathbf{S}^{0,s}$ is a compact and invertible object in FilSp , with inverse $\mathbf{S}^{0,-s}$. In particular, the monoidal unit of FilSp is compact.
- (2) The bigraded homotopy groups $\pi_{*,*}$ detect isomorphisms of filtered spectra.
- (3) As a stable ∞ -category, FilSp is generated under colimits by the spheres $\mathbf{S}^{0,s}$ for $s \in \mathbf{Z}$. That is, the objects $\mathbf{S}^{n,s}$ for $n, s \in \mathbf{Z}$ generate FilSp under colimits. In particular, FilSp is compactly generated by dualisables.
- (4) The monoidal ∞ -category FilSp is rigid in the sense that an object is compact if and only if it is dualisable.

Proof. The first property follows because $\text{Map}(\mathbf{S}^{0,s}, X) \cong \Omega^\infty X^s$, and the second is evident. Item (3) follows from item (2) using [Yan22, Corollary 2.5]. To prove item (4), first notice that the unit $\mathbf{S}^{0,0}$ is compact, so that all dualisable objects are compact. As it is also generated by compact dualisable objects, it follows that every compact object is dualisable too; see, e.g., (the footnote to) [NPR24, Terminology 4.8]. ■

By default, we will equip FilSp with the following t-structure. We borrow the name from [Bar23], though there appears to be no agreed-upon name for it.

Definition 2.22. The **diagonal t-structure** on FilSp is the t-structure where a filtered spectrum X is connective if and only if

$$\pi_{n,s} X = 0 \quad \text{whenever } n < s.$$

We write $\tau_{\geq n}^{\text{diag}}$ and $\tau_{\leq n}^{\text{diag}}$ for the n -connective cover and n -truncation functors with respect to this t-structure, respectively.

As this specified class of connective objects is presentable and closed under colimits, this determines a unique accessible t-structure on FilSp by [HA, Proposition 1.4.4.11]. It enjoys the following properties.

Proposition 2.23.

(a) A filtered spectrum X is connective if and only if

$$\pi_{n,s} X = 0 \quad \text{whenever } n < s.$$

(b) A filtered spectrum X is 0-truncated if and only if

$$\pi_{n,s} X = 0 \quad \text{whenever } n > s.$$

(c) The connective cover $\tau_{\geq 0}^{\text{diag}} X \rightarrow X$ induces an isomorphism

$$\pi_{n,s}(\tau_{\geq 0}^{\text{diag}} X) \xrightarrow{\cong} \pi_{n,s}(X) \quad \text{whenever } n \geq s.$$

Likewise, the 0-truncation $X \rightarrow \tau_{\leq 0}^{\text{diag}} X$ induces an isomorphism

$$\pi_{n,s}(X) \xrightarrow{\cong} \pi_{n,s}(\tau_{\leq 0}^{\text{diag}} X) \quad \text{whenever } n \leq s.$$

(d) The functor $\pi_{*,*}$ induces a symmetric monoidal equivalence

$$\text{FilSp}^{\heartsuit} \xrightarrow{\cong} \text{grAb}, \quad X \longmapsto (\pi_{n,n} X)_n.$$

(e) The diagonal t -structure is complete.

(f) The diagonal t -structure is compatible with filtered colimits.

(g) The diagonal t -structure is compatible with the monoidal structure.

Proof. See, e.g., [Hed20, Propositions II.1.22–II.1.24]. ■

This t -structure is a convenient device for giving functorial definitions of the Whitehead filtration and Postnikov tower.

Definition 2.24.

(1) The **constant filtration** is the functor $\text{Const}: \text{Sp} \rightarrow \text{FilSp}$ given by precomposition with $\mathbf{Z}^{\text{op}} \rightarrow \Delta^0$.

(2) The **Whitehead filtration** is the functor $\text{Wh}: \text{Sp} \rightarrow \text{FilSp}$ given by $\tau_{\geq 0}^{\text{diag}} \circ \text{Const}$.

(3) The **Postnikov tower** is the functor $\text{Post}: \text{Sp} \rightarrow \text{FilSp}$ given by $\tau_{\leq 0}^{\text{diag}} \circ \text{Const}$.

From this definition, the functors Wh and Post come with natural transformations

$$\text{Wh} \longrightarrow \text{Const} \quad \text{and} \quad \text{Const} \longrightarrow \text{Post}.$$

Remark 2.25. Note that Const is naturally a (strong) symmetric monoidal functor. Since the diagonal t-structure is monoidal, it follows that $\text{Wh}: \text{Sp} \rightarrow \text{FilSp}$ is naturally a lax symmetric monoidal functor.

We now check that these are the expected filtrations.

Lemma 2.26. *Let X be a spectrum. Then $\text{Wh } X$ is of the form*

$$\cdots \longrightarrow \tau_{\geq 1} X \longrightarrow \tau_{\geq 0} X \longrightarrow \tau_{\geq -1} X \longrightarrow \cdots$$

with the natural maps as transition maps, and $\text{Post } X$ is of the form

$$\cdots \longrightarrow \tau_{\leq 1} X \longrightarrow \tau_{\leq 0} X \longrightarrow \tau_{\leq -1} X \longrightarrow \cdots$$

with the natural maps as transition maps.

Proof. If F is a filtered spectrum that is t -truncated in the diagonal t-structure for some t , then it follows that $\pi_n F^{-\infty} = 0$ for all n , so that $F^{-\infty} = 0$. Consequently, using the cofibre sequence

$$\text{Wh } X \longrightarrow \text{Const } X \longrightarrow \tau_{\leq -1}^{\text{diag}}(\text{Const } X)$$

and that taking colimits is exact, we find that the natural map $\text{Wh } X \rightarrow \text{Const } X$ induces an isomorphism on colimits. In other words, the colimit of $\text{Wh } X$ is X . Because $\text{Wh } X$ is connective in the diagonal t-structure, the spectrum $(\text{Wh } X)^s$ is s -connective. The cofibre of $(\text{Wh } X)^s \rightarrow X$ is an $(s-1)$ -truncated spectrum, because the cofibre of $\text{Wh } X \rightarrow \text{Const } X$ is (-1) -truncated in the diagonal t-structure. This means that the map $(\text{Wh } X)^s \rightarrow X$ exhibits the source as the connective cover of the target. This applies to all s , implying that for every s , the connecting map $(\text{Wh } X)^{s+1} \rightarrow (\text{Wh } X)^s$ is given by the $(s+1)$ -connective cover.

The same reasoning applies to $\text{Post } X$, using the limit instead of the colimit. ■

The functor Const restricts to an equivalence from Sp to the full subcategory of FilSp on the *constant* filtered spectra, i.e., those for which all transition maps are isomorphisms. We will revisit this later under the guise of τ -invertible filtered spectra in Section 3.2.1.

This definition of the Whitehead tower functor is also helpful to prove that Wh is fully faithful, and to characterise when a filtered spectrum is isomorphic to the Whitehead filtration on a spectrum. We follow the proof given by [AP25, Theorem A.1].

Theorem 2.27. *The functor $\text{Wh}: \text{Sp} \rightarrow \text{FilSp}$ is fully faithful. Its essential image consists of those filtered spectra $F: \mathbf{Z}^{\text{op}} \rightarrow \text{Sp}$ such that for every $s \in \mathbf{Z}^{\text{op}}$, the natural map $F^s \rightarrow F^{-\infty}$ exhibits the source as the s -connective cover of the target.*

Proof. Let X and Y be spectra. The natural transformation $\text{Wh} \rightarrow \text{Const}$ induces the second map in the composite

$$\text{Map}_{\text{Sp}}(X, Y) \longrightarrow \text{Map}_{\text{FilSp}}(\text{Wh } X, \text{Wh } Y) \longrightarrow \text{Map}_{\text{FilSp}}(\text{Const } X, \text{Const } Y).$$

This composite is an isomorphism because Const is fully faithful. It therefore suffices to show that the second map in this composite is an isomorphism. We use that the natural transformation $\text{Wh} \rightarrow \text{Const}$ has two characterisations: first, it exhibits the source as the connective cover of the target; and second, it exhibits the target as the colimit (i.e., underlying spectrum) of the source. The first is by definition of the functor Wh , and the second follows from completeness of the t-structure on spectra. The map under consideration therefore factors as

$$\begin{array}{ccc} \text{Map}_{\text{FilSp}}(\text{Wh } X, \text{Wh } Y) & \xrightarrow{\cong} & \text{Map}_{\text{FilSp}}(\text{Wh } X, \text{Const } Y) \\ & \searrow & \downarrow \cong \\ & & \text{Map}_{\text{FilSp}}(\text{Const } X, \text{Const } Y), \end{array}$$

where the horizontal map is an isomorphism by the universal property of the connective cover, and the second since Const is right adjoint to taking the colimit.

To prove the essential image is of the claimed form, suppose F is a filtered spectrum of the described form. Write X for the spectrum $F^{-\infty}$. We claim that $F \rightarrow \text{Const } X$ exhibits F as the connective cover (in the diagonal t-structure) of the target. Since F is connective in the diagonal t-structure, the natural map $F \rightarrow \text{Const } X$ lifts to $F \rightarrow \text{Wh } X$. Now both $F^s \rightarrow X$ and $(\text{Wh } X)^s \rightarrow X$ exhibit the source as the s -connective cover of the target, so the map between them is an isomorphism. ■

Definition 2.28. Let X be a filtered spectrum. We say that X is **complete** if the limit X^∞ of X vanishes. We write $\widehat{\text{FilSp}}$ for the full subcategory of FilSp on the complete filtered spectra.

The inclusion $\widehat{\text{FilSp}} \subseteq \text{FilSp}$ admits a left adjoint, called the *completion functor*

$$\text{FilSp} \longrightarrow \widehat{\text{FilSp}}, \quad X \longmapsto \widehat{X} := \text{cofib}(X^\infty \rightarrow X)$$

We will explore this concept in more detail later under the name of τ -*completion* of filtered spectra in Section 3.2.3. In particular, we will see that a map of filtered spectra is an isomorphism after completion if and only if it is an isomorphism on associated graded; see Proposition 3.30. For an alternative discussion without the language of τ , see [Hed20, Section II.1.2].

Finally, we turn to the relationship between filtered spectra and filtered abelian groups.

Definition 2.29. Let X be a filtered spectrum. The **induced strict filtration** on the abelian group $\pi_n X^{-\infty}$ is the (strict) filtration given by

$$F^s \pi_n X^{-\infty} := \text{im}(\pi_n X^s \longrightarrow \pi_n X^{-\infty}).$$

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

In practice, 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 this 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.

2.3 Spectral sequences

So far, we have not yet delived on our promise that the derived setting is able to measure both the kernel and cokernel of transition maps. The notion of a spectral sequence makes this precise. Its purpose is to reconstruct the homotopy groups $\pi_* X^{-\infty}$ from the associated graded of X . The non-injectivity of the transition maps on homotopy groups leads to the *differentials* in a spectral sequence. In these notes, we will further argue that instead of only reconstructing $\pi_* X^{-\infty}$ from the associated graded, the better approach is to reconstruct *all* of the homotopy groups in the filtered spectrum, i.e., the bigraded homotopy groups $\pi_{*,*} X$.

We do not include an in-depth review or motivation of the setup of a spectral sequence here. We include a detailed but informal introduction to this in Appendix A.

Remark 2.30. The indexing for exact couples used below is not the most common. Usually, Serre indexing is used in formulas and definitions, while Adams indexing is used for displaying the spectral sequence. We prefer to work with one indexing system throughout, and we prefer Adams indexing as it is the most practical one and arises naturally from the diagram chase of Appendix A. Moreover, this indexing is more straightforward to generalise to contexts where homotopy groups have a more complicated indexing (such as filtered or synthetic spectra later in these notes).

For a reference on exact couples, see, e.g., [McC00, page 37 and following].

Construction 2.31. Let X be a filtered spectrum.

- (1) The **associated exact couple** of X is the exact couple of bigraded abelian groups defined by

$$A^{n,s}(X) := \pi_n X^s \quad \text{and} \quad E^{n,s}(X) := \pi_n \text{Gr}^s X,$$

with the natural maps from the long exact sequence between them, fitting into the following diagram, where each map is annotated by its (n, s) -bidegree.

$$\begin{array}{ccc} \pi_n X^s & \xrightarrow{(0, -1)} & \pi_n X^s \\ & \swarrow (-1, 1) & \searrow (0, 0) \\ & \pi_n \text{Gr}^s X & \end{array}$$

- (2) The **underlying spectral sequence** of X , denoted by $\{E_r^{n,s}(X), d_r\}_{r \geq 1}$, is the spectral sequence in bigraded abelian groups arising from the exact couple associated with X . By definition, this spectral sequence is of the form

$$E_1^{n,s}(X) := \pi_n \text{Gr}^s X \implies \pi_n X^{-\infty}$$

and the differential d_r has bidegree $(-1, r)$ for $r \geq 1$. We refer to $\pi_* X^{-\infty}$ as the **abutment** of the spectral sequence.

The term *abutment* and the above notation are not meant to indicate any form of convergence; rather, one should think of this as the object to which the spectral sequence is trying to converge.

Example 2.32. The associated graded of $\mathbf{S}^{0,0}$ is given by \mathbf{S} in degree 0, and zero elsewhere. In particular, the first page of the resulting spectral sequence is

$$E_1^{n,s} \cong \begin{cases} \pi_n \mathbf{S} & \text{if } s = 0, \\ 0 & \text{if } s \neq 0. \end{cases}$$

The spectral sequence of a general filtered bigraded sphere is a shift of the one for $\mathbf{S}^{0,0}$. These spectral sequences are rather uninteresting: they do not decompose $\pi_* \mathbf{S}$ in a new, meaningful way. As a result, we mostly think of the filtered bigraded spheres as useful formal objects, and not out of interest in their underlying spectral sequence. \blacktriangle

To avoid confusion, let us make some indexing conventions explicit.

Definition 2.33. Let X be a filtered spectrum. Let $r \geq 0$, and let $n, s \in \mathbf{Z}$.

- (1) Write $Z_r^{n,s} \subseteq E_1^{n,s}$ for the subgroup of the **r -cycles**, i.e., those x such that the differentials $d_1(x), \dots, d_r(x)$ vanish. This leads to a sequence of inclusions

$$\dots \subseteq Z_2^{n,s} \subseteq Z_1^{n,s} \subseteq Z_0^{n,s} = E_1^{n,s}.$$

We define

$$Z_\infty^{n,s} := \lim_r Z_r^{n,s} = \bigcap_r Z_r^{n,s}.$$

An element of $Z_\infty^{n,s}$ is called a **permanent cycle**.

- (2) Write $B_r^{n,s} \subseteq E_1^{n,s}$ for the subgroup of the **r -boundaries**, i.e., the image of the first r differentials. This leads to a sequence of inclusions

$$0 = B_0^{n,s} \subseteq B_1^{n,s} \subseteq B_2^{n,s} \subseteq \dots \subseteq E_1^{n,s}.$$

We define

$$B_\infty^{n,s} := \text{colim}_r B_r^{n,s} = \bigcup_r B_r^{n,s}.$$

Note that there is an inclusion $B_\infty^{n,s} \subseteq Z_\infty^{n,s}$.

(3) If $r \geq 1$, we define

$$E_r^{n,s} := Z_{r-1}^{n,s} / B_{r-1}^{n,s}.$$

(4) We define the ∞ -**term** and the **derived ∞ -term**, respectively, to be

$$\begin{aligned} E_\infty^{n,s} &:= Z_\infty^{n,s} / B_\infty^{n,s}, \\ RE_\infty^{n,s} &:= \lim^1_r Z_r^{n,s}. \end{aligned}$$

Warning 2.34. In [Boa99], Boardman writes Z_r for what we would write as Z_{r-1} , and likewise writes B_r for what we would write as B_{r-1} .

Warning 2.35. In general, the group $E_\infty^{n,s}$ is not a (co)limit of the groups $E_r^{n,s}$. In fact, as $E_{r+1}^{n,s}$ is a subquotient of $E_r^{n,s}$, there is in general no sensible map from $E_r^{n,s}$ to or from $E_{r+1}^{n,s}$. In certain cases, such a map does exist. For instance, suppose that $E_1^{*,s} = 0$ for $s \ll 0$. For fixed s , then for $r \gg 0$, we have that $E_{r+1}^{*,s}$ is a subgroup of $E_r^{*,s}$, and it follows that $E_\infty^{*,s}$ is the limit along the resulting sequence of inclusion maps.

Definition 2.36. Let X be a filtered spectrum, let $x \in E_1^{n,s}$ be an element, and let $\theta \in \pi_n X$.

- (1) Let $r \geq 1$. Suppose that x is an r -cycle, so that it defines an element in $Z_r^{n,s}$. We say that x **survives to page r** if its image in $E_r^{n,s}$ is nonzero. If x is a permanent cycle, then we say that x **survives to page ∞** if its image in $E_\infty^{n,s}$ is nonzero.
- (2) Suppose that x is a permanent cycle that survives to page ∞ . We then say that x **detects θ** (or that θ is *detected by* x) if there exists a lift $\alpha \in \pi_n X^s$ of x that maps to θ under $X^s \rightarrow X^{-\infty}$.

Beware that, in the definition of detection, the lift α above need not be unique if it exists.

Remark 2.37. The underlying spectral sequence from Construction 2.31 is functorial in the filtered spectrum. As exact couples form a 1-category, it suffices to lift this construction to a functor from hFilSp to exact couples. This is easily checked. Post-composing this with the functor from exact couples to spectral sequences, we obtain the desired functor

$$\text{FilSp} \longrightarrow \text{SSeq}(\text{grAb}), \quad X \longmapsto \{E_r^{n,s}(X), d_r\}.$$

Roughly speaking, this functor forgets the homotopy of the filtered spectrum and only remembers the homotopy of the associated graded, together with the induced differentials. The goal of the spectral sequence, then, is to reconstruct the homotopy of the filtered spectrum from this data. Whether this is even possible (aside from extension problems) is the question of *convergence*, which we discuss in Section 2.3.1.

We think of the ∞ -category FilSp as an ∞ -categorical enhancement of the category $\text{SSeq}(\text{grAb})$. Although not every (stable) spectral sequence arises from a filtered spectrum, in practice, they do.^[1] The additional homotopical structure on filtered spectra

has many advantages. For example, it allows us to talk about coherently multiplicative filtrations; see [Hed20, Part II].

For completeness, we compare the above construction of the underlying spectral sequence with some other ones appearing in the literature.

Remark 2.38 (Alternative constructions). We used exact couples in our definition of the underlying spectral sequence, as this fits most closely with the explanation given in Appendix A. There are a number of alternative approaches in the literature, including at least the following.

- ◆ Cartan–Eilenberg systems [CE56, Section XV.7], which are also used by Lurie in [HA, Section 1.2.2].
- ◆ Coherent cochain complexes in spectra, introduced by Ariotta [Ari21].
- ◆ The décalage functor on the level of filtered spectra. Décalage was introduced by Deligne for chain complexes [Del71], and later generalised to cosimplicial objects by Levine [Lev15], and generalised to filtered spectra by Antieau, as first written down by Hedenlund [Hed20, Part II] and later by Antieau [Ant24].

All of these approaches agree, as a consequence of the following results.

- ◆ Ariotta [Ari21, Theorem 3.19] constructs an equivalence of ∞ -categories from $\widehat{\text{FilSp}}$ to coherent cochain complexes in spectra.
- ◆ Let X be a filtered spectrum. This gives rise to a Cartan–Eilenberg system via [HA, Definition 1.2.2.9]. Iterating the décalage functor on filtered spectra also results in a spectral sequence. Antieau [Ant24, Theorem 4.13] showed that the spectral sequence arising from this Cartan–Eilenberg system is isomorphic to the one arising from iterating the décalage functor. More generally, he showed this when working with filtered objects of a stable ∞ -category with sequential limits and colimits that is equipped with a t-structure.
- ◆ Let X be a filtered spectrum. This gives rise to both a Cartan–Eilenberg system as before, as well as to an exact couple via Construction 2.31. The two resulting spectral sequences are isomorphic; see the discussion in [CE56, page 336], or, e.g., Creemers [Cre25, Theorem 5.1] for a detailed argument.

Remark 2.39. We use the variant of a spectral sequence that computes homotopy classes out of a compact object. Alternatively, as in [HA, Section 1.2.2], one can also start with a

^[1]An example of a spectral sequence that does not arise from a filtration is the p -Bockstein spectral sequence in the way that it is set up in [Bro61, Section 1] or [McC00, Chapter 10]. However, this is not an essential issue: there is an alternative way to set up the p -Bockstein spectral sequence that does come from a filtered spectrum (namely, from the p -adic filtration on a spectrum or chain complex; see Example 2.65 and Example 3.93). Arguably, this latter version even has a nicer abutment, converging to $\pi_*(X_p^\wedge)$, rather than to $\mathbf{F}_p \otimes F$, where F denotes the torsion-free quotient of π_*X .

stable ∞ -category equipped with a t-structure, and set up a spectral sequence to compute the heart-valued homotopy groups. One should assume that heart-valued homotopy groups preserve sequential colimits for this to work in a reasonable generality.

Finally, we make a few remarks regarding indexing and notation.

Remark 2.40 (First- vs. second-page indexing). We choose to index spectral sequences arising from filtered spectra to start on the first page; let us call this *first-page indexing*. For various reasons (such as aesthetics, or to better fit alternative definitions of a spectral sequence), it can be useful to reindex this to start on the second (or any later) page. One can achieve *second-page indexing* via the reindexing

$$\tilde{E}_{r+1}^{n,s} := E_r^{n,n+s}.$$

It is straightforward to check that this turns d_r -differentials into \tilde{d}_{r+1} -differentials; in particular, this makes the spectral sequence start on the second page. Using this second-page indexing, the filtered bigraded spheres take the form (using the functor i from Definition 2.17)

$$\tilde{\mathbf{S}}^{n,s} := \Sigma^n i(n+s).$$

Concretely, this is the filtered spectrum that is \mathbf{S}^n in positions $n+s$ and below (with identities between them), and zero elsewhere. In this grading, categorical suspension takes the form $\Sigma^{1,-1}$.

Remark 2.41. In accordance with our previous indexing conventions, n here is indexed homologically, whereas s is indexed cohomologically. As such, it would be more honest to write E_n^s and A_n^s , but we do not do so, as the current notation is well established (and would leave little room for the page-index r). Depending on the context, it might be more natural to alter either of these conventions; see, e.g., Remark 2.54.

2.3.1 Convergence

Convergence of a spectral sequence is the question whether one can reconstruct $\pi_* X^{-\infty}$ from the spectral sequence.^[2] More precisely, convergence concerns reconstructing $\pi_* X^{-\infty}$ by reconstructing the induced strict filtration on $\pi_* X^{-\infty}$. We can however only hope to reconstruct the associated graded of this strict filtration, and other methods are necessary for solving the extension problems.

In most accounts of spectral sequences, convergence is additional structure on the spectral sequence. For a spectral sequence arising from a filtered spectrum, all of this structure is supplied by the filtered spectrum, so that convergence becomes a property.

^[2]Other times, the convergence issue is to prove that the colimit of the filtration is isomorphic to a desired spectrum. This is, of course, a question that is more specific to the situation at hand, so in this section we regard the colimit as the desired object to study.

We closely follow Boardman's account [Boa99]. He works with (unrolled) exact couples, but we specialise everything to the setting of filtered spectra. We include a few detailed remarks, both for the curious reader and for use later in the more technical parts of these notes. An alternative introduction is given by Hedenlund in [Hed20, Section 1.2.2].

Recall from Definition 2.29 that a filtered spectrum X gives rise to an *induced strict filtration* on $\pi_n X^{-\infty}$, denoted by $F^s \pi_n X^{-\infty}$.

Construction 2.42. Let X be a filtered spectrum. Write $\partial_{n,s}: \pi_n \text{Gr}^s X \rightarrow \pi_{n-1} X^{s+1}$ for the boundary map. By a diagram chase (see, e.g., [Boa99, Lemma 5.6] or [Rog21, Lemma 2.5.10]), there is a natural isomorphism of graded abelian groups

$$\frac{F^s \pi_n X^{-\infty}}{F^{s+1} \pi_n X^{-\infty}} \cong \frac{\ker \partial_{n,s}}{B_\infty^{n,s}}.$$

Using that $\ker \partial_{n,s}/B_\infty^{n,s}$ naturally injects into $E_\infty^{n,s}$, this leads to a natural injective map

$$\text{Gr}^s \pi_n X^{-\infty} = \frac{F^s \pi_n X^{-\infty}}{F^{s+1} \pi_n X^{-\infty}} \hookrightarrow E_\infty^{n,s}. \quad (2.43)$$

Definition 2.44. Let X be a filtered spectrum. We say that the underlying spectral sequence **converges strongly** to $\pi_* X^{-\infty}$ if

- (a) the induced strict filtration $\{F^s \pi_* X^{-\infty}\}$ is *derived complete* in the sense of Definition 2.11, i.e.,

$$\lim_s F^s \pi_* X^{-\infty} = 0 \quad \text{and} \quad \lim_s^1 F^s \pi_* X^{-\infty} = 0;$$

- (b) the natural map (2.43) is an isomorphism for all n and s .

Warning 2.45. This terminology is abusive: the above definition of strong convergence is not a condition on the spectral sequence, but rather on the filtered spectrum. In fact, the above conditions do not make sense if we do not specify which filtered spectrum gives rise to the spectral sequence.

In words, condition (a) says that we can (up to extension problems) reconstruct $\pi_* X^{-\infty}$ from the induced strict filtration, and condition (b) says that the spectral sequence is able to recover the associated graded of this filtration.

It is more accurate to speak of convergence to the induced strict filtration on $\pi_* X^{-\infty}$, but we will usually not do this. If the filtered spectrum is clear from the context, we may also be brief and simply say that the spectral sequence *converges strongly*, which should always be understood as convergence to $\pi_* X^{-\infty}$.

Remark 2.46. Because the map (2.43) is always injective, we learn the following, even in the absence of any of the above convergence criteria. If $x \in E_1^{n,s}$ is a permanent cycle

that survives to page ∞ , then for any lift $\alpha \in \pi_n X^s$, the image of α in $\pi_n X^{-\infty}$ is nonzero. Indeed, by injectivity of (2.43), the image of α defines a nonzero map in $\text{Gr}^s \pi_n X^{-\infty}$. In other words, any element that x detects is nonzero. (Without convergence hypotheses however, a lift of x to $\pi_n X^s$ may not exist.)

Remark 2.47. Item (a) in particular implies that every nonzero element of $\pi_* X^{-\infty}$ is detected by a permanent cycle that survives to page ∞ . Meanwhile, item (b) implies that every permanent cycle that survives to page ∞ detects an element in $\pi_* X^{-\infty}$ (which is nonzero by Remark 2.46).

Remark 2.48. Cartan–Eilenberg [CE56, page 321] also gives names to other notions of convergence. In the setting of filtered spectra, these notions are the following. The spectral sequence is said to *converge weakly* if (b) in Definition 2.44 holds, and said to *converge* if (b) holds and $\lim_s F^s \pi_* X^{-\infty} = 0$. We will not use this terminology.

As the name suggests, strong convergence is the strongest type of convergence one can hope for. Our goal then is to find conditions that guarantee strong convergence. In practice, we only have limited knowledge about the homotopy groups $\pi_{n,s} X = \pi_n X^s$, so we would prefer convergence criteria that involve mostly the spectral sequence rather than the filtered spectrum itself.

Boardman’s notion of *conditional convergence* does exactly this. It splits the problem up into two parts. First, one needs to establish conditional convergence, which is a structural (and often mild) condition on the filtered spectrum. Second, once this is established, there are conditions phrased entirely in terms of the spectral sequence that guarantee strong convergence. These conditions are more computational in nature, and need to be checked on a case-by-case basis, but are often met. The second step is the reason for using the word ‘conditional’: the spectral sequence converges strongly, conditional on these requirements being met.

Definition 2.49. Let X be a filtered spectrum. We say that the underlying spectral sequence **converges conditionally** to $\pi_* X^{-\infty}$ if the filtered spectrum X is complete, i.e., if the limit X^∞ vanishes.

Warning 2.50. As with strong convergence, this terminology is abusive: conditional convergence is a condition on the filtration rather than the underlying spectral sequence.

Remark 2.51. Using the Milnor short exact sequence (see, e.g., [Boa99, Theorem 4.9])

$$0 \longrightarrow \lim_s^1 \pi_{n+1} X^s \longrightarrow \pi_n X^\infty \longrightarrow \lim_s \pi_n X^s \longrightarrow 0,$$

we see that X^∞ vanishes if and only if

$$\lim_s \pi_* X^s = 0 \quad \text{and} \quad \lim_s^1 \pi_* X^s = 0.$$

Warning 2.52. The previous remark may appear to suggest that the vanishing of X^∞ implies that the induced strict filtration on $\pi_* X^{-\infty}$ is derived complete. This is *not* true in general, and this is exactly what leads to Boardman's convergence criteria. To explain why, we introduce the following notation:

$$F^s \pi_{n,w} X = \text{im}(\pi_{n,w+s} X \longrightarrow \pi_{n,w} X).$$

There are two problems. The first is that for any w , the natural map

$$\pi_n X^\infty \longrightarrow F^\infty \pi_{n,w} X := \lim_s F^s \pi_{n,w} X$$

need not be surjective. This does happen if RE_∞ vanishes; see [Boa99, Lemma 5.9]. The second problem is that the natural map

$$\text{colim}_w F^\infty \pi_{n,w} X = \text{colim}_w \lim_s F^s \pi_{n,w} X \longrightarrow \lim_s \text{colim}_w F^s \pi_{n,w} X = F^\infty \pi_n X^{-\infty}$$

also need not be surjective. (There is an analogous version of this map with the first derived limit in the place of the limit, but this is always surjective; see [Boa99, Lemma 8.11].) It is surjective in certain cases, such as when the filtered spectrum is left or right concentrated in the sense of Definition 2.53 below. In general, Boardman's *whole-plane obstruction* is the obstruction to this implication; see Remark 2.57 below for a further discussion, and [Boa99, Lemma 8.11] for the precise result alluded to here. Clearly, if all three maps above are surjective, then the vanishing of X^∞ does imply derived completeness of the induced strict filtration on $\pi_* X^{-\infty}$.

The 'conditional' part of conditional convergence becomes easier if the filtered spectrum is of a special form. The following terminology is nonstandard. For the general case, see Remark 2.57.

Definition 2.53. Let X be a filtered spectrum.

- (1) We say that X is **right concentrated** if the transition maps $X^{s+1} \rightarrow X^s$ are isomorphisms for $s \gg 0$.
- (2) We say that X is **left concentrated** if the transition maps $X^{s+1} \rightarrow X^s$ are isomorphisms for $s \ll 0$:

These conditions can be checked entirely in terms of the associated graded: being right concentrated means that $E_1^{*,s}$ vanishes for $s \gg 0$, and being left concentrated means that it vanishes for $s \ll 0$.

Remark 2.54. In practice, we usually reindex a right-concentrated filtered spectrum to be of the form

$$\dots \xrightarrow{\cong} X^1 \longrightarrow X^0 \longrightarrow X^{-1} \longrightarrow \dots,$$

in which case $X^\infty = X^1$ and $E_1^{*,s} = 0$ for $s > 0$. Boardman calls the resulting spectral sequence a *half-plane spectral sequence with exiting differentials*. Likewise, we index a left-concentrated filtered spectrum to be of the form

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

in which case $X^{-\infty} = X^0$ and $E_1^{*,s} = 0$ for $s < 0$. Boardman calls the resulting spectral sequence a *half-plane spectral sequence with entering differentials*.

Roughly speaking, for right-concentrated filtered spectra, the ‘conditional’ part is vacuous, while for left-concentrated filtered spectra, the only thing one has to check is the vanishing of a derived limit term. This can be checked in terms of the spectral sequence, without requiring any further knowledge of the filtration.

Theorem 2.55 (Conditional convergence, Boardman). *Let X be a filtered spectrum. Suppose that X is right concentrated. Then the following are equivalent.*

- (1a) *The spectral sequence underlying X converges conditionally.*
- (1b) *The spectral sequence underlying X converges strongly.*

Suppose instead that X is left concentrated. Then any two of the following imply the third.

- (2a) *The spectral sequence underlying X converges conditionally.*
- (2b) *The derived ∞ -term $RE_\infty^{n,s}$ from Definition 2.33 (4) vanishes for all n and s .*
- (2c) *The spectral sequence underlying X converges strongly.*

Proof. This is [Boa99, Theorem 6.1 and Theorem 7.3], respectively. The translation between his and our notation is the following (using notation from Warning 2.52):

$$\begin{aligned} A^s &= \pi_{*,s} X & Q^s &= F^\infty \pi_{*,s} X \\ A^\infty &= \lim_s \pi_{*,s} X & RQ^s &= \lim_t^1 F^t \pi_{*,s} X \\ RA^\infty &= \lim_s^1 \pi_{*,s} X. & & \blacksquare \end{aligned}$$

As explained in [Boa99, Section 7], it is often easy to verify that the derived ∞ -term vanishes. This happens if the spectral sequence collapses at a finite page (i.e., $d_r = 0$ for $r > r_0$), or if for every n and s , only finitely many differentials leaving bidegree (n, s) are nonzero (but where this bound is allowed to depend on n and s). In general, as with any first-derived limit of abelian groups, one can use the Mittag-Leffler condition to check its vanishing.

In certain cases, the following variant will be useful as well. It is slightly more general than working with right-concentrated filtered spectra that are complete: we only ask that every filtered abelian group $\pi_n X$ becomes zero for $s \gg 0$, but not necessarily that they become zero at the same point.

Proposition 2.56. *Let X be a filtered spectrum. Suppose that for every n , the group $\pi_n X^s$ vanishes for $s \gg 0$ (where the bound on s is allowed to depend on n). Then X is complete and the spectral sequence underlying X converges strongly.*

Proof. From Construction 2.42, it is clear that the map (2.43) is an isomorphism if $Z_\infty^{n,s} = \ker \partial_{n,s}$ for all n, s . The inclusion $\ker \partial_{n,s} \subseteq Z_\infty^{n,s}$ always holds. The reverse inclusion now follows from a diagram chase involving the definition of the differential: see, e.g., [Rog21, Lemma 2.5.8]. Finally, the completeness of X follows from Remark 2.51. ■

Remark 2.57 (Whole-plane spectral sequences). In the case where the filtered spectrum does not become constant in either direction, the situation becomes more difficult. Boardman [Boa99, Section 8, Equation (8.7)] defines a group W he calls the *whole-plane obstruction*. For any filtered spectrum X , Theorem 8.10 of op. cit. implies that any two of the following imply the third.

- (a) The spectral sequence underlying X converges conditionally.
- (b) The derived ∞ -term $RE_\infty^{n,s}$ vanishes for all n and s , and W vanishes.
- (c) The spectral sequence underlying X converges strongly.

He gives a criterion saying that if there is no infinite family of differentials that all cross each other in their interior, then $W = 0$; see [Boa99, Lemma 8.1]. For an alternative description of the whole-plane obstruction using Cartan–Eilenberg systems (see Remark 2.38 for more on Cartan–Eilenberg systems), see [HR19]. They moreover give an alternative proof for Boardman’s criterion: see [HR19, Proposition 5.3].

2.4 Digression: reflecting

So far, we have been thinking of a filtered spectrum (and consequently its underlying spectral sequence) as a tool to understand its colimit. It is also possible to orient things the other way around, using a filtered spectrum to understand its limit instead. In that case, the colimit is the object that should vanish to guarantee convergence properties. The convergence discussion of spectral sequences does become slightly more involved in this context, because the functor $\pi_*: \mathbb{S}p \rightarrow \mathbf{grAb}$ does not preserve sequential limits. Nevertheless, for every convergence result we discussed above, Boardman [Boa99] also gives the limit-oriented version.

Rather than working with this limit-oriented version, we will use a duality to move back to the colimit-oriented version. We refer to this as *reflection*. This is well known (being used, for instance, in [Bou79, Section 5]), but we thought it would be helpful to make this translation and its basic properties explicit, especially its interaction with completion of filtered spectra. It will not play a big role in the rest of these notes. Let us also point out that this duality is a feature specific to the stable setting.

Remark 2.58. Reflection can also be set up in the setting of exact couples. In that setting, when passing from a tower to a filtration, there will be a *Rees system* in the sense of [EH66, Section 7] witnessing that the two resulting spectral sequences are isomorphic. We do not detail how to specialise this argument to the setting of filtered spectra, focussing instead on the behaviour of the duality at a more categorical level.

We remind the reader of the terminology we introduced in Remark 2.8: we will refer to objects of FilSp as *filtrations* when working in the colimit-oriented context, and as *towers* when working to the limit-oriented context. Accordingly, if a notion of passing back and forth between towers and filtrations is to make sense, it should interchange the limit and colimit. A natural candidate then is to reflect a filtration in its colimit, and to reflect a tower in its limit.

Definition 2.59. The **associated tower functor**, respectively the **associated filtration functor**, are the functors defined by

$$\begin{aligned} (-)^{\text{tow}} : \text{FilSp} &\longrightarrow \text{FilSp}, & X &\longmapsto \text{cofib}(\Sigma^{0,-1}X \rightarrow \text{Const } X^{-\infty}), \\ (-)^{\text{fil}} : \text{FilSp} &\longrightarrow \text{FilSp}, & X &\longmapsto \text{fib}(\text{Const } X^{\infty} \rightarrow \Sigma^{0,1}X). \end{aligned}$$

Concretely, if X is a filtered spectrum, then

$$(X^{\text{tow}})^s = \text{cofib}(X^{s+1} \longrightarrow X^{-\infty}) \quad \text{and} \quad (X^{\text{fil}})^s = \text{fib}(X^{\infty} \longrightarrow X^{s-1}).$$

We include a shift in these definitions because when working with the spectral sequence associated to a tower, one usually lets the first page consist of the fibres of the transition maps, not the cofibres.

Definition 2.60. Let X be a filtered spectrum. The **fibre-associated graded** of X is the graded spectrum $\text{fibGr } X$ given by

$$\text{fibGr}^s X := \text{fib}(X^s \longrightarrow X^{s-1}).$$

In a picture, a tower X together with its fibre-associated graded looks as follows:

$$\begin{array}{ccccccc} & \text{fibGr}^1 X & & \text{fibGr}^0 X & & \text{fibGr}^{-1} & \\ & \downarrow & & \downarrow & & \downarrow & \\ \dots & \longrightarrow & X^1 & \longrightarrow & X^0 & \longrightarrow & X^{-1} \longrightarrow \dots \end{array}$$

The following is a straightforward diagram chase.

Proposition 2.61. *Let X be a filtered spectrum. Then there are natural isomorphisms of graded spectra*

$$\text{fibGr}(X^{\text{tow}}) \cong \text{Gr } X \quad \text{and} \quad \text{Gr}(X^{\text{fil}}) \cong \text{fibGr } X.$$

The reflection functors from Definition 2.59 destroy some of the information contained in the original object: the associated tower only depends on the completion of the original object. An analogous statement is true for the associated filtration, for which we introduce the following terminology.

Definition 2.62. We call a filtered spectrum **cocomplete** if its colimit vanishes. We write $\widetilde{\text{FilSp}}$ for the full subcategory of FilSp on the cocomplete filtered spectra.

Analogously to the case of completion, one can check that the inclusion $\widetilde{\text{FilSp}} \subseteq \text{FilSp}$ admits a right adjoint given by

$$\text{FilSp} \longrightarrow \widetilde{\text{FilSp}}, \quad X \longmapsto \check{X} := \text{fib}(X \rightarrow X^{-\infty}).$$

We call this functor **cocompletion**, which then features in a colocalisation

$$\widetilde{\text{FilSp}} \xrightleftharpoons[\widetilde{(-)}]{} \text{FilSp}.$$

We think of $\widehat{\text{FilSp}}$ as the conditionally convergent filtrations, and of $\widetilde{\text{FilSp}}$ as the conditionally convergent towers. The reflection functors of Definition 2.59 translate between these in the following way.

Proposition 2.63.

(1) *We have commutative diagrams*

$$\begin{array}{ccc} \text{FilSp} & \xrightarrow{(-)^{\text{tow}}} & \text{FilSp} \\ \widehat{(-)} \downarrow & & \uparrow \\ \widehat{\text{FilSp}} & \xrightarrow{(-)^{\text{tow}}} & \widetilde{\text{FilSp}} \end{array} \quad \text{and} \quad \begin{array}{ccc} \text{FilSp} & \xrightarrow{(-)^{\text{fil}}} & \text{FilSp} \\ \widetilde{(-)} \downarrow & & \uparrow \\ \widetilde{\text{FilSp}} & \xrightarrow{(-)^{\text{fil}}} & \widehat{\text{FilSp}}. \end{array}$$

(2) *If X is a filtered spectrum, then we have natural isomorphisms*

$$(X^{\text{tow}})^{\infty} \cong (\widehat{X})^{-\infty} \quad \text{and} \quad (X^{\text{fil}})^{-\infty} \cong (\check{X})^{\infty}.$$

In particular, if X is complete, then we have a natural isomorphism

$$(X^{\text{tow}})^{\infty} \cong X^{-\infty},$$

while if X is cocomplete, then we have a natural isomorphism

$$(X^{\text{fil}})^{-\infty} \cong X^{\infty}.$$

(3) *The reflection functors restrict to inverse equivalences*

$$\widehat{\text{FilSp}} \begin{array}{c} \xrightarrow{(-)^{\text{tow}}} \\ \xleftarrow{(-)^{\text{fil}}} \end{array} \widetilde{\text{FilSp}}.$$

Proof. For item (1), we only show the first diagram, as the argument for the second is dual to that for the first. Using that cofibres preserve colimits, it follows immediately from the definition that X^{tow} is cocomplete for all X . It remains to be verify that the map $X \rightarrow \widehat{X}$ becomes an isomorphism after taking associated towers. Write T for X^{tow} and R for $(\widehat{X})^{\text{tow}}$. Then for every s , we have a commutative diagram where all rows and columns are cofibre sequences

$$\begin{array}{ccccc} X^\infty & \longrightarrow & X^{s+1} & \longrightarrow & (\widehat{X})^{s+1} \\ \parallel & & \downarrow & & \downarrow \\ X^\infty & \longrightarrow & X^{-\infty} & \longrightarrow & (\widehat{X})^{-\infty} \\ \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & T^s & \longrightarrow & R^s. \end{array}$$

It follows that $T^s \rightarrow R^s$ is an isomorphism for all s , proving the claim.

For item (2), we again only show the statement about associated towers. For a general filtered spectrum X , we have

$$(X^{\text{tow}})^\infty = \lim_s \text{cofib}(X^{s+1} \longrightarrow X^{-\infty}) \cong \text{cofib}(X^\infty \longrightarrow X^{-\infty}).$$

By item (1), it suffices to consider the case where X is complete. In this case, the latter term is naturally isomorphic to $X^{-\infty}$, proving the claim.

Finally, for item (3), we check that the composite $((-)^{\text{tow}})^{\text{fil}}$ is isomorphic to the identity on complete objects; the argument for the other composite on cocomplete is analogous. Let X be complete, and write $T = X^{\text{tow}}$, and $F = T^{\text{fil}}$. Then using that $X^\infty = 0$, we have a commutative diagram where all rows and columns are cofibre sequences

$$\begin{array}{ccccc} \Omega X^s & \longrightarrow & 0 & \longrightarrow & X^s \\ \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & X^{-\infty} & \xlongequal{\quad} & X^{-\infty} \\ \downarrow & & \downarrow & & \downarrow \\ F^s & \longrightarrow & T^\infty & \longrightarrow & T^{s-1}. \end{array}$$

This supplies a natural cofibre sequence $\Sigma^{-1,0}X \rightarrow 0 \rightarrow F$, proving the claim. ■

Example 2.64. Let X be a spectrum. It is a straightforward exercise to see that the reflection functors switch the Whitehead filtration and Postnikov tower of X : we have natural isomorphisms

$$(\mathrm{Wh} X)^{\mathrm{tow}} \cong \mathrm{Post} X \quad \text{and} \quad (\mathrm{Post} X)^{\mathrm{fil}} \cong \mathrm{Wh} X.$$

Note that this uses that the standard t-structure on spectra is complete. If we were to work in $\mathrm{Fil}(\mathcal{C})$ for a stable ∞ -category \mathcal{C} equipped with a t-structure, then in order to interchange $\mathrm{Wh} X$ and $\mathrm{Post} X$, we would need to reflect them in X , rather than in their (co)limit. \blacktriangle

Example 2.65. Let X be a spectrum. The p -Bockstein filtration of X is the filtered spectrum

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

indexed to be constant from filtration 0 onwards. Observe that this filtration is complete if and only if X is p -complete. Its associated tower is

$$\cdots \longrightarrow X/p^3 \longrightarrow X/p^2 \longrightarrow X/p \longrightarrow 0 \longrightarrow \cdots,$$

where X/p^n appears in filtration $n - 1$. The limit of this tower is X_p^\wedge . If we take the associated filtration of this tower, then we obtain the p -adic filtration on X_p^\wedge :

$$\cdots \xrightarrow{p} X_p^\wedge \xrightarrow{p} X_p^\wedge \xrightarrow{p} X_p^\wedge = \cdots. \quad \blacktriangle$$

2.5 The Adams spectral sequence

We give a brief introduction to the Adams spectral sequence. While this is not intended as a first introduction, we take some care to explain some of the subtleties in its construction. Our model for the Adams spectral sequence will be based on a cosimplicial object. While this has some small downsides, it leads to an easier expression for the resulting filtered spectrum, because the cosimplicial décalage has an easy expression (Definition 2.78). This limitation is not essential by any means, and can be avoided by working with filtered resolutions (Remark 2.90). Later, we will use synthetic spectra to set up this improved version; see Section 4.4.

2.5.1 The Tot spectral sequence

We write Δ for the simplex category, and write $\Delta_{\leq n}$ for the full subcategory on the objects $[k]$ where $k \leq n$.

Definition 2.66. Let \mathcal{C} be a pointed ∞ -category, and let $X^\bullet: \Delta \rightarrow \mathcal{C}$ be a cosimplicial object of \mathcal{C} .

(1) The **totalisation** of X^\bullet is the limit

$$\mathrm{Tot} X^\bullet = \lim_{\Delta} X^\bullet.$$

(2) For $n \geq 0$, the n -th **partial totalisation** is the limit

$$\mathrm{Tot}_n X^\bullet = \lim_{\Delta_{\leq n}} X^\bullet.$$

(3) Suppose that \mathcal{C} admits (partial) totalisations. Using the filtration

$$\cdots = 0 \subseteq \Delta_{\leq 0} \subseteq \Delta_{\leq 1} \subseteq \Delta_{\leq 2} \subseteq \cdots \subseteq \Delta,$$

we obtain a tower

$$\mathrm{Tot} X^\bullet \longrightarrow \cdots \longrightarrow \mathrm{Tot}_2 X^\bullet \longrightarrow \mathrm{Tot}_1 X^\bullet \longrightarrow \mathrm{Tot}_0 X^\bullet \longrightarrow 0 \longrightarrow \cdots$$

with limit $\mathrm{Tot} X^\bullet$. We call this tower the **totalisation tower** (or *Tot tower*) of X^\bullet . This is natural in X , resulting in a functor $\mathrm{towTot}: \mathcal{C}^\Delta \rightarrow \mathrm{Fil}(\mathcal{C})$.

We call this a *tower* in accordance with Remark 2.8; below in Definition 2.72, we will turn this tower into a filtration. Note that the Tot tower is in particular a cocomplete object of $\mathrm{Fil}(\mathcal{C})$ in the sense of Definition 2.62.

Remark 2.67 (Cubes). The n -th partial totalisation can be computed as follows. We write

$$\Delta_{/[n]}^{\mathrm{inj}} \subseteq \Delta_{/[n]}$$

for the full subcategory of the slice on those objects given by injective maps to $[n]$. Note that this subcategory is a punctured $(n+1)$ -cube, and in particular is a finite category. We have a forgetful functor

$$\Delta_{/[n]}^{\mathrm{inj}} \longrightarrow \Delta_{\leq n}$$

and this functor is homotopy initial; see, e.g., [HA, Lemma 1.2.4.17]. In words, the n -th partial totalisation can be computed as the total fibre of a punctured $(n+1)$ -cube.

Notation 2.68. Let $\Delta_{\mathrm{inj}} \subseteq \Delta$ denote the wide subcategory on those maps that are injective. We will write j for the inclusion functor.

Recall that a **semicosimplicial object** of \mathcal{C} is a functor $\Delta_{\mathrm{inj}} \rightarrow \mathcal{C}$. The definition of the Tot tower can be mimicked for a semicosimplicial object, instead taking the limit over Δ_{inj} or over $(\Delta_{\mathrm{inj}})_{\leq n}$. There is a subtle difference between these two.

Remark 2.69. The inclusion $j: \Delta_{\mathrm{inj}} \subseteq \Delta$ is homotopy initial; see, e.g., [Dug17, Example 21.2]. As a result, we will also write Tot for the limit-functor

$$\mathrm{Sp}^{\Delta_{\mathrm{inj}}} \longrightarrow \mathrm{Sp},$$

and we therefore have a natural factorisation

$$\begin{array}{ccc} \mathrm{Sp}^\Delta & \xrightarrow{\mathrm{Tot}} & \mathrm{Sp} \\ & \searrow^{j^*} & \nearrow^{\mathrm{Tot}} \\ & \mathrm{Sp}^{\Delta_{\mathrm{inj}}} & \end{array}$$

However, the same is not true for the partial totalisations: the inclusion $(\Delta_{\mathrm{inj}})_{\leq n} \subseteq \Delta_{\leq n}$ is not homotopy initial; see, e.g., [Dug17, Section 21.6]. Limits over $(\Delta_{\mathrm{inj}})_{\leq n}$ are bigger in some sense; see Remark 2.70 and Proposition 2.76.

The semicosimplicial Tot tower is in fact an example of a cosimplicial Tot tower.

Remark 2.70. Right Kan extension along $j: \Delta_{\mathrm{inj}} \subseteq \Delta$ is a right adjoint to the forgetful functor $\mathrm{Sp}^\Delta \rightarrow \mathrm{Sp}^{\Delta_{\mathrm{inj}}}$, and likewise for Ab in place of Sp . If X^\bullet is a semicosimplicial spectrum, then we can compute this right Kan extension to be the cosimplicial object given by

$$(j_* X)^n = \prod_{[n] \rightarrow [k]} X^k$$

with the product ranging over all surjective maps in Δ . A similar formula holds for semicosimplicial abelian groups. We learn a number of things from this computation.

- (1) The functor π_* preserves right Kan extension along j .
- (2) The restriction of $j_* X^\bullet$ to $\Delta_{\leq n}$ is right Kan extended from the restriction of X^\bullet to $(\Delta_{\mathrm{inj}})_{\leq n}$. In particular, the Tot tower of $j_* X$ is the semicosimplicial Tot tower of X^\bullet .

When working with cosimplicial rather than semicosimplicial objects, the Tot tower sets up a one-to-one correspondence between cosimplicial objects and certain towers. In the following, let us write $\mathrm{Fil}^{\geq 0}(\mathcal{C})$ for the full subcategory of $\mathrm{Fil}(\mathcal{C})$ on those filtered objects that vanish in negative filtration. (This subcategory should not be confused with the connective part of the diagonal t-structure on $\mathrm{Fil}(\mathcal{C})$.) Then by definition, the Tot tower functor of Definition 2.66 lands in $\mathrm{Fil}^{\geq 0}(\mathcal{C})$.

Theorem 2.71 (Stable Dold–Kan correspondence, Lurie). *Let \mathcal{C} be a stable ∞ -category. Then the Tot tower functor restricts to an equivalence*

$$\mathrm{towTot}: \mathcal{C}^\Delta \xrightarrow{\simeq} \mathrm{Fil}^{\geq 0}(\mathcal{C}).$$

Proof. This is a reformulation of [HA, Theorem 1.2.4.1]. Indeed, the notion of stability is self-dual, so the equivalence proved there dualises to an equivalence

$$\mathrm{towTot}: \mathcal{C}^\Delta \xrightarrow{\simeq} \mathrm{Fun}(\mathbf{Z}_{\geq 0}^{\mathrm{op}}, \mathcal{C}).$$

Finally, right Kan extension along the inclusion $\mathbf{Z}_{\geq 0}^{\mathrm{op}} \rightarrow \mathbf{Z}^{\mathrm{op}}$ (informally, putting zeroes in negative filtrations) results in a functor $\mathrm{Fun}(\mathbf{Z}_{\geq 0}^{\mathrm{op}}, \mathcal{C}) \rightarrow \mathrm{Fil}(\mathcal{C})$ that is fully faithful with essential image $\mathrm{Fil}^{\geq 0}(\mathcal{C})$. ■

The Tot tower leads to a spectral sequence. We prefer to work with spectral sequences arising from filtrations rather than towers, so we apply the reflection duality from Section 2.4 to land in this situation.

Definition 2.72. Let X^\bullet be a cosimplicial spectrum. By reflecting the Tot tower of X^\bullet via the functor $(-)^{\text{fil}}$ from Definition 2.59, we obtain a filtered spectrum that we call the **totalisation filtration** (or *Tot filtration*) of X^\bullet . We denote its terms by

$$\text{Tot}^n X^\bullet := \text{fib}(\text{Tot} X^\bullet \longrightarrow \text{Tot}_{n-1} X^\bullet).$$

The resulting filtration $\text{filTot} X^\bullet$ is of the form

$$\dots \longrightarrow \text{Tot}^2 X^\bullet \longrightarrow \text{Tot}^1 X^\bullet \longrightarrow \text{Tot} X^\bullet \longleftarrow \text{Tot} X^\bullet \longleftarrow \dots,$$

which is constant from filtration 0 onwards. We call the spectral sequence associated to this filtered spectrum the **totalisation spectral sequence** (or *Tot spectral sequence*, or *Bousfield–Kan spectral sequence*).

Note that the Tot filtration is, by construction, a complete filtration of $\text{Tot} X^\bullet$. In general however, the conditional convergence of the Tot spectral sequence need not be strong.

Remark 2.73. From Theorem 2.71 and Proposition 2.63, it follows that the Tot filtration functor induces an equivalence from \mathcal{C}^Δ to the full subcategory of $\widehat{\text{Fil}}(\mathcal{C})$ on those objects that are constant from filtration 0 onwards.

A nice feature of this spectral sequence is that there is a formula for its first page and its first differential purely in terms of the cosimplicial abelian groups $\pi_t X^\bullet$.

Notation 2.74.

- (1) Let A^\bullet be a semicosimplicial abelian group. The **unnormalised cochain complex** $C(A^\bullet)$ of A^\bullet is the cochain complex

$$\dots \longrightarrow 0 \longrightarrow C^0(A^\bullet) \longrightarrow C^1(A^\bullet) \longrightarrow C^2(A^\bullet) \longrightarrow \dots$$

given by, for $m \geq 0$,

$$C^m(A^\bullet) = A^m$$

and with differential $C^m(A^\bullet) \rightarrow C^{m+1}(A^\bullet)$ given by the alternating sum $\sum (-1)^i d^i$.

- (2) If A^\bullet is a cosimplicial group, we define its unnormalised cochain complex to be the unnormalised cochain complex of the semicosimplicial group $j^* A^\bullet$.
- (3) Let A^\bullet be a cosimplicial abelian group. The **normalised cochain complex** $N(A^\bullet)$ of A^\bullet is the cochain complex

$$\dots \longrightarrow 0 \longrightarrow N^0(A^\bullet) \longrightarrow N^1(A^\bullet) \longrightarrow N^2(A^\bullet) \longrightarrow \dots$$

given by, for $m \geq 0$,

$$N^m(A^\bullet) = \bigcap_{i=0}^{m-1} \ker(s^i: A^m \longrightarrow A^{m-1})$$

and with differential $N^m(A^\bullet) \rightarrow N^{m+1}(A^\bullet)$ given by the alternating sum $\sum (-1)^i d^i$.

If A^\bullet is a cosimplicial abelian group, then we have an evident map of cochain complexes

$$N(A^\bullet) \longrightarrow C(A^\bullet).$$

This turns out to be a quasi-isomorphism, and to even admit a quasi-inverse; see the dual^[3] of [HA, Proposition 1.2.3.17].

Notation 2.75. Let A^\bullet be a semicosimplicial abelian group, and let $s \geq 0$. We write $\pi^s A^\bullet$ for the s -th cohomology group of $C(A^\bullet)$.

Proposition 2.76.

- (1) Let X^\bullet be a cosimplicial spectrum, and let $E_r^{*,*}(X^\bullet)$ denote the resulting Tot spectral sequence. For all integers n and s , there is a natural isomorphism

$$E_1^{n,s}(X^\bullet) \cong N^s(\pi_{n+s} X^\bullet)$$

that identifies the d_1 -differential with the differential of $N(\pi_{n+s} X^\bullet)$.

- (2) Let X^\bullet be a semicosimplicial spectrum, and let $j_* X^\bullet$ denote the right Kan extension of it to a cosimplicial object (see Remark 2.70). For all integers n and s , there is a natural isomorphism

$$E_1^{n,s}(j_* X^\bullet) \cong C^s(\pi_{n+s} X^\bullet)$$

that identifies the d_1 -differential with the differential of $C(\pi_{n+s} X^\bullet)$.

- (3) Let X^\bullet be a cosimplicial spectrum. Then under the above identifications, the unit $X^\bullet \rightarrow j_* j^* X^\bullet$ induces, on the first page of the Tot spectral sequence, the natural map

$$N^s(\pi_{n+s} X^\bullet) \longrightarrow C^s(\pi_{n+s} X^\bullet).$$

In particular, this map of spectral sequences is an isomorphism from the second page onward.

Proof. This follows by dualising (i.e., applying it to $\mathcal{C} = \text{Sp}^{\text{op}}$) the discussion of [HA, Remark 1.2.4.4, Variant 1.2.4.9]. ■

Remark 2.77. One can phrase item (3) at the level of filtered spectra using the décalage functor. Let $\text{Déc}: \text{FilSp} \rightarrow \text{FilSp}$ denote Antieau's décalage functor; see [Ant24] or [Hed20,

^[3]To see that the definition of the normalised chain complex of a simplicial object in \mathcal{A}^{op} used therein indeed dualises to the normalised cochain complex of a cosimplicial object in \mathcal{A} , use [GJ09, Theorem III.2.1].

Part II]. Then if X^\bullet is a cosimplicial spectrum, the above implies that the unit $X^\bullet \rightarrow j_*j^*X^\bullet$ induces an isomorphism of filtered spectra

$$\text{Déc}(\text{filTot}(X^\bullet)) \longrightarrow \text{Déc}(\text{filTot}(j_*j^*X^\bullet)).$$

Indeed, the filtrations before applying décalage are complete, and since décalage preserves complete filtrations ([Ant24, Lemma 4.18]), it is enough to check this on associated graded; see Proposition 3.31. After applying décalage, the associated graded of this map is an isomorphism by item (3) above.

There is a variant of the Tot filtration where we have ‘turned a page’ in the spectral sequence, and thereby starts on the second page of the Tot spectral sequence. While there is a more general notion of such a page-turning operation (see Remark 2.38), the cosimplicial version has a simpler expression. These two definitions of décalage actually coincide: see Proposition 2.82.

The terminology is taken from Deligne’s operation of the same name for chain complexes from [Del71, Definition (1.3.3)]. The generalisation to cosimplicial spectra is a result of Levine [Lev15, Section 6]. Note that he also studies other (for instance, unstable) settings; we specialise his results to spectra. Moreover, Levine compares his décalage operation to that of Deligne; see [Lev15, Remark 6.4].

Definition 2.78. Let $X^\bullet: \Delta \rightarrow \text{Sp}$ be a semicosimplicial spectrum. We define the **décalage** of X^\bullet as the filtered spectrum given by

$$\text{Déc}^\Delta X^\bullet = \text{Tot}(\text{Wh } X^\bullet).$$

If X^\bullet is a cosimplicial spectrum, then we let $\text{Déc}^\Delta X^\bullet$ denote the décalage of the underlying semicosimplicial object j^*X^\bullet ; cf. Remark 2.69.

Concretely, the value at filtration s of $\text{Déc}^\Delta X^\bullet$ is given by $\text{Tot}(\tau_{\geq s} X^\bullet)$, the totalisation of the (semi)cosimplicial spectrum obtained by applying $\tau_{\geq s}$ levelwise to X^\bullet .

Remark 2.79. The functor $\text{Déc}^\Delta: \text{Sp}^{\Delta_{\text{inj}}} \rightarrow \text{FilSp}$ is naturally lax symmetric monoidal. Indeed, it is the composite

$$\text{Sp}^{\Delta_{\text{inj}}} \xrightarrow{\text{Wh}} \text{Fil}(\text{Sp})^{\Delta_{\text{inj}}} \xrightarrow{\text{Tot}} \text{FilSp}.$$

The first functor is lax symmetric monoidal (for the levelwise symmetric monoidal structure on cosimplicial objects) by Remark 2.25, and the second is lax symmetric monoidal because it is the right adjoint to the constant functor $\text{FilSp} \rightarrow \text{Fil}(\text{Sp})^{\Delta_{\text{inj}}}$ which is symmetric monoidal. The same applies for cosimplicial objects (or alternatively by noting that the forgetful functor to semicosimplicial objects is symmetric monoidal).

The namesake of the décalage construction is that it has turned the page once compared to the Tot spectral sequence. To state the comparison between these spectral sequences then, we require a reindexing; see Remark 2.40.

Theorem 2.80 (Levine). *Let X^\bullet be a (semi)cosimplicial spectrum, and let $\{E_r^{*,*}(X^\bullet)\}_{r \geq 1}$ denote the Tot spectral sequence associated to X^\bullet . Then there is an isomorphism of spectral sequences (where $r \geq 1$)*

$$E_r^{n,s}(\text{Déc}^\Delta X^\bullet) \cong E_{r+1}^{n,s-n}(X^\bullet).$$

Proof. This follows from [Lev15, Proposition 6.3], but note that Levine uses a different indexing from the Adams indexing used above. ■

Accordingly, it often makes sense to use *second-page indexing* for the spectral sequence underlying $\text{Déc}^\Delta X^\bullet$. Phrased like this, the above theorem says that this second-page indexed $\text{Déc}^\Delta X^\bullet$ is isomorphic to the second page onwards of $E_r^{*,*}(X^\bullet)$.

For later use, we record the following more basic property of the cosimplicial décalage construction.

Proposition 2.81. *Let X^\bullet be a (semi)cosimplicial spectrum. Then the filtered spectrum $\text{Déc}^\Delta X^\bullet$ is naturally a complete filtration of $\text{Tot } X^\bullet$, meaning that its limit vanishes and its colimit is naturally isomorphic to $\text{Tot } X^\bullet$.*

Proof. Complete filtered spectra are closed under limits, so completeness follows from the fact that Whitehead filtrations of spectra are complete. To compute the colimit of this filtration, note that, for every integer s , we have a fibre sequence

$$\tau_{\geq s+1} X^\bullet \longrightarrow X^\bullet \longrightarrow \tau_{\leq s} X^\bullet.$$

Taking totalisations, one therefore has a natural fibre sequence

$$\text{Tot}(\tau_{\geq s+1} X^\bullet) \longrightarrow \text{Tot } X^\bullet \longrightarrow \text{Tot}(\tau_{\leq s} X^\bullet).$$

Taking colimits over s , one has a natural cofibre sequence

$$\text{colim}_s \text{Tot}(\tau_{\geq s+1} X^\bullet) \longrightarrow \text{Tot } X^\bullet \longrightarrow \text{colim}_s \text{Tot}(\tau_{\leq s} X^\bullet).$$

By definition, the left-hand term gives the colimit of $\text{Déc}^\Delta(X^\bullet)$, so it suffices to show that the right-hand term vanishes. Since coconnectivity is preserved by limits, we see that for all s , the spectrum $\text{Tot } \tau_{\leq s} X^\bullet$ is s -truncated. As homotopy groups of spectra preserve filtered colimits, the colimit is therefore $(-\infty)$ -truncated, and hence vanishes. ■

Finally, we end by comparing Levine's cosimplicial décalage to Antieau's décalage. This appears to be folklore, but for lack of a citeable reference, we give a proof here. Note, however, that while Levine's décalage also applies to unstable settings, Antieau's décalage lives entirely in the stable world. Let us write $\text{Déc}: \text{FilSp} \rightarrow \text{FilSp}$ for Antieau's décalage functor.

Proposition 2.82. *Let X^\bullet be a (semi)cosimplicial spectrum. Then there is an isomorphism, natural in X^\bullet , of filtered spectra*

$$\text{Déc}^\Delta X^\bullet \cong \text{Déc}(\text{filTot } X^\bullet).$$

The proof uses the exact same argumentation as Antieau’s argument for the Atiyah–Hirzebruch spectral sequence in [Ant24, Propostion 9.2].

Proof. We assume familiarity with the definition of décalage via connective covers in the Beilinson t-structure, as explained in detail in [Ant24] or [Hed20, Part II].

Applying filTot to the diagram $\text{Wh } X^\bullet \rightarrow X^\bullet$ in $\text{Sp}^{\Delta_{\text{inj}}}$ yields a natural diagram in FilSp

$$\cdots \longrightarrow \text{filTot}(\tau_{\geq 1} X^\bullet) \longrightarrow \text{filTot}(\tau_{\geq 0} X^\bullet) \longrightarrow \cdots \longrightarrow \text{filTot}(X^\bullet). \quad (2.83)$$

We claim that this is the Whitehead filtration of $\text{filTot}(X^\bullet)$ in the Beilinson t-structure on FilSp . First, we show that for every integer n , the filtered spectrum $\text{filTot}(\tau_{\geq n} X^\bullet)$ is Beilinson n -connective. For this, we ought to show that its associated graded in filtration s is an $(n - s)$ -connective spectrum. Using Proposition 2.76, we compute

$$\pi_k(\text{Gr}^s \text{filTot}(\tau_{\geq n} X^\bullet)) = \text{C}^s(\pi_{k+s}(\tau_{\geq n} X^\bullet)),$$

which evidently vanishes if $k + s < n$, i.e., if $k < n - s$, so that the s -th associated graded is indeed $(n - s)$ -connective. When working with cosimplicial objects and the cosimplicial Tot filtration, the same applies, using the normalised cochain complex instead.

It follows that the natural map $\text{filTot}(\tau_{\geq n} X^\bullet) \rightarrow \text{filTot}(X^\bullet)$ factors through a map

$$\text{filTot}(\tau_{\geq n} X^\bullet) \longrightarrow \tau_{\geq n}^{\text{Bei}} \text{filTot}(X^\bullet).$$

We claim this is an isomorphism. As both filtrations are complete (the second one by [Ant24, Lemma 4.18]), it is enough to show this on associated graded; see Proposition 3.31. There, it follows using Proposition 2.76 and [Hed20, Proposition II.2.4].

We conclude that (2.83) is indeed the Beilinson Whitehead filtration of $\text{filTot } X^\bullet$. From the definition of décalage, it follows that

$$\text{Déc}(\text{filTot } X^\bullet) = \text{colim}_s \text{Tot}^s(\tau_{\geq *s} X^\bullet) \cong \text{Tot}(\tau_{\geq *} X^\bullet) = \text{Déc}^\Delta X^\bullet. \quad \blacksquare$$

Using this comparison result, the properties of décalage from Theorem 2.80 and Proposition 2.81 also follow from general properties of Antieau’s décalage functor from [Ant24] and [Hed20, Part II].

Remark 2.84. The cosimplicial décalage functor Déc^Δ still has an advantage over the composite functor $\text{Déc} \circ \text{filTot}$: the former is naturally lax symmetric monoidal (see Remark 2.79). Although the functor $\text{Déc}: \text{FilSp} \rightarrow \text{FilSp}$ is also lax symmetric monoidal, the functor filTot is not,^[4] so the lax monoidal structure only arises through the version of Definition 2.78.

^[4]If it were, then it would send \mathbf{E}_∞ -rings in Sp^Δ (for the levelwise monoidal structure) to filtered \mathbf{E}_∞ -rings. Instead, what we see is that it sends \mathbf{E}_∞ -rings in Sp^Δ to filtered objects in \mathbf{E}_∞ -rings.

Remark 2.85 (Beilinson vs. levelwise t-structures). We can deduce more than stated in the proposition above: we learn that under the stable Dold–Kan correspondence, the levelwise t-structure on cosimplicial spectra corresponds to the Beilinson t-structure on a subcategory of FilSp . This is remarked by Lawson in [Law24a, Remark 3.16]. In more detail: recall from Remark 2.73 that the functor $\text{filTot}: \text{Sp}^\Delta \rightarrow \widehat{\text{FilSp}}$ defines an equivalence onto those complete filtered spectra that are constant from filtration 0 onwards. In the proof above, we showed that filTot sends the levelwise Whitehead tower to the Beilinson Whitehead tower. Since it is an equivalence, the result follows.

2.5.2 The cosimplicial Adams spectral sequence

The most general definition of the E -Adams spectral sequence we will use is the following, though often we will work with more specific (and more structured) spectra E .

Definition 2.86. Let E be a spectrum with a map $\mathbf{S} \rightarrow E$, and let X be a spectrum. The map $\mathbf{S} \rightarrow E$ gives rise to an augmented semicosimplicial spectrum $\Delta_{\text{inj},+} \rightarrow \text{Sp}$ of the form

$$\mathbf{S} \longrightarrow E \rightrightarrows E \otimes E \rightrightarrows \cdots .$$

The **semicosimplicial E -based Adams resolution** for X is the semicosimplicial spectrum

$$\text{ASS}_E^\Delta(X) := E^{[\bullet]} \otimes X = \left(E \otimes X \rightrightarrows E \otimes E \otimes X \rightrightarrows \cdots \right).$$

Define the filtered spectrum

$$\text{ASS}_E(X) := \text{filTot}(\text{ASS}_E^\Delta(X)).$$

The **E -based Adams spectral sequence** for X is the spectral sequence underlying this filtered spectrum. More generally, if Y and X are spectra, then we define the semicosimplicial spectrum

$$\text{ASS}_E^\Delta(Y, X) := \text{map}(Y, E^{[\bullet]} \otimes X)$$

and the filtered spectrum

$$\text{ASS}_E(Y, X) := \text{filTot}(\text{ASS}_E^\Delta(Y, X)),$$

and define the E -based Adams spectral sequence for $[Y, X]$ to be the spectral sequence underlying this filtered spectrum.

We are careful to work with semicosimplicial objects rather than cosimplicial objects, because upgrading $E^{[\bullet]}$ to a cosimplicial object requires an E_1 -structure on E ; see [MNN17, Construction 2.7]. Not all E of interest may admit this structure, and the Adams spectral sequence should not depend on it either. If E does admit this structure, then this upgrades $\text{ASS}_E^\Delta(Y, X)$ to a cosimplicial spectrum, which we will denote by the same notation. We

would then define $\text{ASS}_E(Y, X)$ using the cosimplicial Tot tower, which has a more efficient first page compared to the semicosimplicial approach; see Proposition 2.76.

This brings us to a potentially confusing point about the Adams spectral sequence: usually, one is interested in it only from the *second page* onward. One should view the first page as ‘ill-defined’ in some sense; it may admit many models, and the one from the above definition is a rather inefficient one at that. What we call the cosimplicial Adams resolution should be regarded as only one of many potential resolutions; the one we chose is convenient as it is functorial.

Our definition above is made to align with standard conventions. Alternatively, one can do away with the first page entirely, as follows.

Variation 2.87. Alternatively, we could have defined the Adams spectral sequence as arising from the filtered spectrum

$$\text{Déc}^\Delta(\text{ASS}_E^\Delta(Y, X)).$$

This filtered spectrum should be viewed as the ‘true’ incarnation of the Adams spectral sequence. It turns out to have better monoidality properties, though this is only possible to prove in general using more modern machinery; see Section 4.4, particularly Remarks 4.62 and 4.74. To align with standard indexing, one should index the resulting spectral sequence using second-page indexing; see Remark 2.40.

Remark 2.88. The Adams spectral sequence is not specific to spectra, but as this is the main case of interest, we will stick to it for our discussion. Baker–Lazarev [BL01] set up the Adams spectral sequence in modules over an \mathbf{E}_∞ -ring, and Mathew–Naumann–Noel [MNN17, Part 1] set up the Adams spectral sequence in a presentably symmetric monoidal ∞ -category. For an even more general setup, see the next remark.

Remark 2.89. Miller [Mil81; Mil12] showed that the E -based Adams spectral sequence depends on much less information than the ring spectrum E : it only depends on the class of morphisms that become nullhomotopic after tensoring with E . This is similar to how E -localisation of spectra depends on much less information than E . For a further discussion and extension of these ideas, see [PP23, Sections 2 and 3].

For the interested reader, we compare the cosimplicial approach to a filtered approach.

Remark 2.90 (Cosimplicial approach and completion). The downside to using the (semi)cosimplicial approach is that it is only able to retrieve the *completion* of the filtered spectrum giving rise to the Adams spectral sequence. Let us explain this by comparing it with the other approach. If E is a spectrum with a map $\mathbf{S} \rightarrow E$, let \bar{E} denote the fibre of this map. The *filtered E -Adams resolution*^[5] of the sphere is the filtered spectrum given by

$$\cdots \longrightarrow \bar{E} \otimes \bar{E} \longrightarrow \bar{E} \longrightarrow \mathbf{S} \longleftarrow \mathbf{S} \longleftarrow \cdots$$

^[5]Various people refer to this as the *E -Adams tower*, but our use of the words *tower* and *filtration* (see Remark 2.8) prevents us from using that terminology here.

indexed to be constant from filtration 0 onwards. If X is a spectrum, then the E -Adams filtration of X is by definition obtained by tensoring this with X levelwise. Clearly the resulting filtered spectrum has colimit X , i.e., it is a filtration of X . It need not be complete however, and this results in a convergence problem. If E is an \mathbf{E}_1 -ring, then the semicosimplicial object $E^{[\bullet]}$ naturally upgrades to a cosimplicial diagram; see [MNN17, Construction 2.7]. Mathew–Naumann–Noel [MNN17, Proposition 2.14] show that, under the stable Dold–Kan correspondence of Theorem 2.71, the Tot tower of the cosimplicial spectrum $E^{[\bullet]} \otimes X$ is matched up with the associated tower (in the sense of Definition 2.59) of the filtered E -Adams resolution for X . Because the associated tower functor factors through completion (Proposition 2.63), this shows that the cosimplicial approach recovers only the completion of the approach based on filtered Adams resolutions in the above sense.

In the generality of Definition 2.86, it is very hard to say much about the second page of the resulting spectral sequence. Things improve if we impose conditions on E . The following is a rather restrictive one, but luckily covers some of the main cases of interest.

Definition 2.91 ([Ada95], Condition III.13.3). Let E be a homotopy associative ring spectrum.

- (1) A finite spectrum P is called **finite E -projective** if E_*P is a projective E_* -module.^[6]
- (2) We say that E is of **Adams type** if it can be written as a filtered colimit of finite E -projective spectra E_α such that for every α , the natural map

$$E^*E_\alpha \longrightarrow \mathrm{Hom}_{E_*}(E_*E_\alpha, E_*) \quad (2.92)$$

is an isomorphism.

Remark 2.93. If E is an \mathbf{E}_1 -ring, then the condition on a finite-projective E_α that the map (2.92) is an isomorphism is automatic. Indeed, if E is \mathbf{E}_1 , then we have a good ∞ -category $\mathrm{Mod}_E(\mathrm{Sp})$ of E -modules, which we can use to set up an Ext spectral sequence as in [EKMM, Chapter IV]. The fact that E_*E_α is projective implies that the resulting Künneth spectral sequence computing E^*E_α is concentrated in filtration 0, implying that (2.92) is an isomorphism. (In fact, the definition of Adams type is engineered to be able to set up a Künneth and universal coefficient spectral sequence; see [Ada95, Chapter III.13].)

Remark 2.94. If E is of Adams type, then E_*E is in particular a flat (left and right) E_* -module. Indeed, homotopy groups preserve filtered colimits, and projective modules are flat.

Example 2.95.

- (1) The sphere is of Adams type: it is the colimit of the one-point diagram $\{ \mathbf{S} \}$.

^[6]This should not be confused with what one might call an *E -finite projective* spectrum, meaning a spectrum P such that E_*P is a finite projective E_* -module. This notion need not imply that the spectrum P is itself finite, but this is a condition we very intentionally require on P .

- (2) If E is \mathbf{F}_p , or more generally if π_*E is a graded field (e.g., if E is a Morava K-theory), then every finite spectrum is finite E -projective. Since every spectrum is a filtered colimit of finite spectra, it follows that E is of Adams type if π_*E is a graded field.

The case $E = \mathbf{F}_p$ is the one originally considered by Adams in [Ada58], and is often simply referred to as the *Adams spectral sequence*. (Strictly speaking, the original version of [Ada58] is based on *cohomology* rather than homology, but for $E = \mathbf{F}_p$, this difference is less material.)

- (3) The ring spectrum \mathbf{MU} is of Adams type, being witnessed by it being the colimit of Thom spectra of Grassmannians. Moreover, every Landweber-exact homotopy-associative ring spectrum is of Adams type; see [Dev97, Proposition 1.3]. In particular, Morava E-theories are of Adams type.

The cases $E = \mathbf{MU}$ and $E = \mathbf{BP}$ are both referred to as the *Adams–Novikov spectral sequence* (ANSS). We refer to [Rav78] for an introduction to the Adams–Novikov spectral sequence and its interplay with the \mathbf{F}_p -Adams spectral sequence.

- (4) A non-example is \mathbf{Z} : this follows since $\pi_*(\mathbf{Z} \otimes \mathbf{Z})$ contains p -torsion for every prime p . In particular, it is not flat over the integers, so Remark 2.94 shows it cannot be of Adams type. Likewise, for every prime p , the ring $\mathbf{Z}_{(p)}$ is not of Adams type.

However, one can modify the notion of Adams type and work with \mathbf{Z} and $\mathbf{Z}_{(p)}$ as if they were of Adams type; see [BP25].

- (5) Another non-example is real K-theory, both in its connective and periodic variants. Mahowald [Mah81] nevertheless computes the second page of the ko -based Adams spectral sequence, leading to a proof of the telescope conjecture at height 1 and at the prime 2; see [Mah82]. ▲

The reason for imposing these restrictions is to obtain an abelian category that computes the second page of the Adams spectral sequence. This abelian category is defined for any Hopf algebroid, not just ones arising from a ring spectrum. For a further and more detailed treatment of (the category of) comodules over a Hopf algebroid, we refer to [Hov04, Section 1] or [Rav04, Appendix A.1].

Definition 2.96. Let (A, Γ) be a graded Hopf algebroid. We write $\text{grComod}_{(A, \Gamma)}$ for the category of comodules over (A, Γ) in graded abelian groups. If n is an integer, then we write $[n]$ for the n -fold shift operator,

$$(M[n])_i = M_{i-n}.$$

The category $\text{grComod}_{(A, \Gamma)}$ is naturally symmetric monoidal, with the underlying A -module being given by the tensor product over A .

To ensure that we can do homological algebra in this setting, we need an assumption on Γ . We say that a (graded) Hopf algebroid (A, Γ) is **flat** if Γ is flat as an A -module via

either the left or right unit; note that it does not matter which unit we choose, as they differ by an automorphism of Γ .

Proposition 2.97. *Let (A, Γ) be a graded Hopf algebroid. Suppose that (A, Γ) is flat. Then the category $\text{grComod}_{(A, \Gamma)}$ is a Grothendieck abelian category, and the forgetful functor to grAb preserves small colimits and finite limits (in particular, it is exact). In particular, $\text{grComod}_{(A, \Gamma)}$ has enough injectives.*

Proof. Graded comodules over (A, Γ) are equivalent to comodules over the comonad

$$\text{grMod}_A \longrightarrow \text{grMod}_A, \quad M \longmapsto \Gamma \otimes_A M.$$

If (A, Γ) is flat, then this comonad is exact, from which it follows that the category of comodules is abelian and that the forgetful functor to grMod_A is exact; see, e.g., [PP23, Proposition A.2]. Consequently, the forgetful functor also detects exactness (since it is conservative). Combining this with the fact that grMod_A is a Grothendieck abelian category and that the comonad $\Gamma \otimes_A -$ preserves colimits, one can check directly that $\text{grComod}_{(A, \Gamma)}$ is Grothendieck abelian also. ■

The forgetful functor $\text{grComod}_{(A, \Gamma)} \rightarrow \text{grMod}_A$ has a right adjoint, given on underlying modules by $M \mapsto \Gamma \otimes_A M$. It follows immediately that the image of an injective A -module under this right adjoint is an injective comodule.

In general, if Γ is not flat over A , then $\text{grComod}_{(A, \Gamma)}$ is merely an additive category.

Notation 2.98. Let (A, Γ) be a graded Hopf algebroid, and let M and N be comodules over it. If s and t are integers, then we write

$$\text{Ext}_{(A, \Gamma)}^{s, t}(M, N) := \text{Ext}_{(A, \Gamma)}^s(M[t], N).$$

Note that this is a *homological* indexing convention for t , and a *cohomological* indexing convention for s .

Construction 2.99. If E is a homotopy-associative ring spectrum, then the pair (E_*, E_*E) is naturally a Hopf algebroid. If E is moreover of Adams type, then E_*E is in particular flat over E_* ; see Remark 2.94. This gives us a good symmetric monoidal abelian category of comodules over (E_*, E_*E) . We will abbreviate this category by

$$\text{grComod}_{E_*E}.$$

The fact that E is of Adams type tells us more about this category: it tells us that it is generated under colimits by dualisable objects; see [Pst22, Section 3.1].

Remark 2.100. Let E be a homotopy-associative ring spectrum. Then E being of Adams type implies that E_*E is flat; see Remark 2.94. It is not known whether the converse is true; see [BLP22] for a discussion (specifically Question (1) following Remark 7 in op. cit.). As we will need the stronger assumption of Adams type, and because of the lack of a known counterexample to flatness implying Adams type, we content ourselves with assuming this potentially stronger condition.

One potentially confusing point of the Adams spectral sequence is that there are two variants of what one might guess the induced strict filtration on $[Y, X]_*$ could be. One of them is more accessible algebraically, while the other is what comes out of the definition of Definition 2.86. In some cases, such as when Y is a sphere, these agree; see Proposition 2.104 below. We use the following (nonstandard) names to distinguish between them.

Definition 2.101. Let X, Y and E be spectra.

- (1) The **algebraic E -Adams filtration** on $[Y, X]$ is the strict filtration where, for every $s \geq 1$, a map f is in $F^s[Y, X]$ when it can be written as a composite

$$f = f_1 \circ \cdots \circ f_s$$

where each map f_i induces the zero map on E_* -homology. We put $F^0[Y, X]$ equal to $[Y, X]$ by definition.

- (2) The **topological E -Adams filtration** on $[Y, X]$ is the strict filtration where, for every $s \geq 1$, a map f is in $F^s[Y, X]$ when it can be written as a composite

$$f = f_1 \circ \cdots \circ f_s$$

where each map f_i becomes a nullhomotopic of spectra after tensoring with E . We put $F^0[Y, X]$ equal to $[Y, X]$ by definition.

Taking $Y = \mathbf{S}^n$ results in filtrations on $\pi_n X$, which is the case we will be interested in most of the time.

In general, the filtration captured by the spectral sequence is the topological one.

Proposition 2.102. *Let E be a spectrum with a homotopy-unital multiplication, and let X and Y be spectra. The induced strict filtration on $[Y, X]_*$ by the filtered spectrum $\text{ASS}_E(Y, X)$ is the topological E -Adams filtration.*

Proof. We freely use the notation introduced in Remark 2.90. By loc. cit., we may compute the induced strict filtration on $[Y, X]_*$ using the Adams resolution

$$\cdots \longrightarrow \bar{E} \otimes \bar{E} \otimes X \longrightarrow \bar{E} \otimes X \longrightarrow X. \quad (2.103)$$

Note that the map $\bar{E} \rightarrow \mathbf{S}$ becomes nullhomotopic after tensoring with E : indeed, the cofibre sequence defining \bar{E} becomes, after tensoring with E , the cofibre sequence

$$E \otimes \bar{E} \longrightarrow E \longrightarrow E \otimes E.$$

The unital multiplication on E provides a retraction of the second map, so indeed the first map is nullhomotopic.

We see therefore that each map in the Adams resolution (2.103) becomes null after tensoring with E . As a result, a map $\Sigma^n Y \rightarrow X$ that lifts to $\bar{E}^{\otimes s} \otimes X$ for $s \geq 1$ is of topological

Adams filtration $\geq s$. For the converse, it is enough to consider the case $s = 1$ and $n = 0$, i.e., we have a map $f: Y \rightarrow X$ that becomes nullhomotopic after tensoring with E . Then also the composite $Y \rightarrow E \otimes Y \rightarrow E \otimes X$ is nullhomotopic, so it follows that there exists a dashed factorisation

$$\begin{array}{ccccc} & & Y & & \\ & \swarrow \text{---} & \downarrow f & & \\ \bar{E} \otimes X & \longrightarrow & X & \longrightarrow & E \otimes X, \end{array}$$

so that f lifts to filtration 1 in the induced strict filtration. \blacksquare

Meanwhile, the algebraic filtration is preferable computationally. Clearly, the topological Adams filtration of a map is a lower bound for its algebraic Adams filtration. In general, the converse need not be true. If E is of Adams type, then it is true so long as E_*Y is projective over E_* .

Proposition 2.104. *Let E be a homotopy-associative ring spectrum of Adams type, and let X and Y be spectra. Suppose that E_*Y is projective over E_* (e.g., if π_*E is a graded field, or if $Y = \mathbf{S}^n$). Then the following hold.*

(1) *There is an isomorphism*

$$E_2^{n,s}(Y, X) \cong \text{Ext}_{E_*E}^{s, n+s}(E_*Y, E_*X),$$

where the left-hand side denotes the second page of the E -Adams spectral sequence.

(2) *The algebraic and topological E -Adams filtrations on $[Y, X]_*$ coincide.*

In particular, the above are true if $Y = \mathbf{S}$.

Proof. For the first part, see [Ada95, Chapter III.15]. The second part then follows from [PP23, Warning 3.21]. \blacksquare

Remark 2.105. The idea of the proof is that one can construct a modified version of the E -Adams spectral sequence whose second page is *always* given by Ext groups of comodules, and whose filtration is *always* the algebraic one. One might refer to this as the E_* -based Adams spectral sequence. There is a comparison map from E -based to the E_* -based version, and if Y has projective homology, then this is an isomorphism on second pages (and hence also on all later pages). For a further explanation of this point, we refer to the work of Patchkoria–Pstragowski [PP23]; see Section 3 of op. cit., particularly Warning 3.21 therein.

2.5.3 Convergence

Next, we deal with the issue of convergence. In this cosimplicial formulation, the object $\text{ASS}_E(X)$ is a complete filtration by Remark 2.73, so it might appear as if there are no convergence problems. The issue, however, is whether the underlying spectrum of $\text{ASS}_E(X)$ is isomorphic to X . Only in this case are we willing to speak of (conditional) convergence of the Adams spectral sequence.

We closely follow Bousfield's seminal discussion on convergence [Bou79, Sections 5 and 6]. Although those results are specific to the setting of spectra, some results have been generalised by Mantovani [Man21, Section 7].

We begin with some terminology.

Definition 2.106. Let X be a spectrum, and let E be spectrum with a map $\mathbf{S} \rightarrow E$. The E -nilpotent completion of X is the spectrum

$$X_E^\wedge := \text{Tot}(E^{[\bullet]} \otimes X).$$

The map $\mathbf{S} \rightarrow E^{[\bullet]}$ induces a natural map $X \rightarrow X_E^\wedge$. We say that X is E -nilpotent complete if this map is an isomorphism.

Bousfield shows that one may alternatively compute the E -nilpotent completion by a choice of what he calls an E -nilpotent resolution; see [Bou79, Proposition 5.8].

Essentially by definition, the E -Adams spectral sequence converges conditionally to $\text{map}(Y, X_E^\wedge)$. If Y is the sphere, then this is simply X_E^\wedge . In both cases then, the question of convergence is whether the map $X \rightarrow X_E^\wedge$ is an isomorphism.

One should think of the E -nilpotent completion as a more computable approximation to the E -localisation of X . As it is a limit of E -local objects, the E -nilpotent completion is E -local, so that the natural map $X \rightarrow X_E^\wedge$ factors through a map

$$L_E X \longrightarrow X_E^\wedge. \tag{2.107}$$

Unfortunately, this map can fail to be an isomorphism. This failure is connected to some unexpected behaviour of nilpotent completion: it can fail to be idempotent.

Proposition 2.108 (Bousfield). *Let E and X be as above. Then the map (2.107) is an isomorphism if and only if the natural map $X_E^\wedge \rightarrow (X_E^\wedge)_E^\wedge$ is an isomorphism.*

Proof. See the discussion following the proof of Proposition 5.5 in [Bou79]. ■

In the case where both E and X are bounded below, the E -nilpotent completion and E -localisation of X coincide and are of an arithmetic nature. For the following, recall that if R is an ordinary ring, then its *core* is its subring

$$\{x \in R \mid x \otimes 1 = 1 \otimes x \text{ holds in } R \otimes R\}.$$

Theorem 2.109 (Bousfield). *Let E be a bounded-below homotopy-associative ring spectrum.*

- (1) *Suppose that the core of the ring $\pi_0 E$ is $\mathbf{Z}[J^{-1}]$ for some set of primes J . Then for every bounded-below spectrum X , the natural map $L_E X \rightarrow X_E^\wedge$ is an isomorphism, and moreover*

$$L_E X \cong X[J^{-1}].$$

(2) Suppose that the core of the ring $\pi_0 E$ is \mathbf{Z}/n for some nonzero integer n . Then for every bounded-below spectrum X , the natural map $L_E X \rightarrow X_E^\wedge$ is an isomorphism, and moreover

$$L_E X \cong X_n^\wedge.$$

Proof. See [Bou79, Theorems 6.5 and 6.6]. Note that Bousfield uses the term *connective* to mean *bounded below*. ■

Remark 2.110. Bousfield computes all rings that can appear as the core of a ring. These fall into four types; the above two are the first two types. Bousfield shows [Bou79, Theorem 6.7] that if the core of $\pi_0 E$ is of one of the other two types, then the map $L_E \mathbf{Z} \rightarrow \mathbf{Z}_E^\wedge$ is not an isomorphism.

Example 2.111. The ring MU satisfies the first condition for $J = \emptyset$, so that for all spectra X that are bounded below, we have

$$L_{\text{MU}} X \cong X_{\text{MU}}^\wedge \cong X.$$

For a fixed prime p , the ring BP satisfies the first condition for J the set of primes different from p , so that for all bounded below spectra X , we have

$$L_{\text{BP}} X \cong X_{\text{BP}}^\wedge \cong X_{(p)}.$$

Lastly, the ring \mathbf{F}_p satisfies the second condition for $n = p$, so that for all bounded below spectra X , we have

$$L_{\mathbf{F}_p} X \cong X_{\mathbf{F}_p}^\wedge \cong X_p^\wedge.$$

Beware that for a general X not bounded below, none of these isomorphisms need to hold. ▲

Chapter 3

The τ -formalism

So far, we have introduced spectral sequences in the way that they are normally introduced. We will now rephrase them in terms of the τ -formalism. In Section 3.1, we introduce τ in filtered abelian groups. Like in the previous chapter, this is both as a warm-up and to describe the structure present on the homotopy groups of filtered spectra. In this context, τ is a helpful variable for keeping track of filtrations and hidden extensions; we showcase some examples in Section 3.1.1. Next, in Section 3.2 we lift this story to filtered spectra, and begin to rephrase spectral sequences in the language of τ . We compare various operations in the τ -formalism in spectra and abelian groups; for example, the issue of convergence can also be phrased as the difference between τ -completion in filtered spectra and in filtered abelian groups (Warning 3.32). This formalism becomes particularly convenient when discussing *total differentials*, which we introduce in Section 3.3, and where we discuss strengthened versions of the ordinary Leibniz rule.

After introducing this more foundational setup, our goal becomes to prove the *Omnibus Theorem*, showing that if X is a filtered spectrum, then the $\mathbf{Z}[\tau]$ -module $\pi_{*,*} X$ captures all of the structure of the spectral sequence underlying X . To help us prove this, we introduce the τ -Bockstein spectral sequence in Section 3.4. Although it is not strictly necessary to use this, it is a helpful organisational tool, and is useful to have available in general. With this in hand, we prove the Omnibus Theorem in Section 3.5.

Finally, we would like to export all of the aforementioned results and tools to other settings. We discuss the notion of a *deformation* in Section 3.6. For deformations arising in a special way, which we refer to as *monoidal deformations*, we show that indeed everything discussed above exports directly. Our main example of a deformation is that of *synthetic spectra* and is deferred to the next chapter, but we include a few additional examples at the end of this chapter.

Many of the results in this chapter are well known to experts. We draw from [Bar23], [BHS22, Appendices A–C], [Lur15], and [Pst25]. Our proof of the Omnibus Theorem is heavily inspired by [BHS23, Appendix A].

3.1 Filtered abelian groups

Recall that a module over the polynomial ring $\mathbf{Z}[x]$ is the same as an abelian group together 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 turn this into a graded ring by giving τ degree -1 . By forgetting the transition maps, a filtered abelian group 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. The forgetful functor thus factors through a functor $\text{FilAb} \rightarrow \text{Mod}_{\mathbf{Z}[\tau]}(\text{grAb})$.

Proposition 3.1. *The functor*

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

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 strict filtrations this is very helpful, particularly when recording filtration jumps and hidden relations. We give a few examples in Section 3.1.1.

Various properties of filtered abelian groups can be rephrased using τ . The ones we focus on are the following.

- (1) A filtered abelian group is *strict* if and only if the corresponding $\mathbf{Z}[\tau]$ -module is τ -power torsion free.
- (2) The *associated graded* of a filtered abelian group A is given by the graded abelian group A/τ .
- (3) The underlying object $A^{-\infty}$ of a filtered abelian group A can be identified with the τ -inversion of the corresponding $\mathbf{Z}[\tau]$ -module.
- (4) A filtered abelian group is *derived complete* (Definition 2.11) if and only if the corresponding $\mathbf{Z}[\tau]$ -module is τ -adically complete.

The first two claims are obvious; let us elaborate on the other two.

Definition 3.2. The **constant filtration functor** $\text{Const}: \text{Ab} \rightarrow \text{FilAb}$ is given by sending an abelian group A to the constant filtration on A , given by

$$\dots \longleftarrow A \longleftarrow A \longleftarrow A \longleftarrow \dots$$

It is easy to check that Const restricts to an equivalence from Ab onto the filtered abelian groups whose transition maps are all isomorphisms; let us say that such a filtered abelian

group is **constant**. In terms of τ , this says that τ acts invertibly on it. For this reason, we may also refer to such filtered abelian groups as being **τ -invertible**, and we write $\text{FilAb}[\tau^{-1}]$ for the full subcategory on these objects.

Proposition 3.3. *The inclusion $\text{FilAb}[\tau^{-1}] \subseteq \text{FilAb}$ admits both a left and a right adjoint. Under the equivalence $\text{Const}: \text{Ab} \simeq \text{FilAb}[\tau^{-1}]$, the adjunctions*

$$\begin{array}{ccc} \text{FilAb} & \begin{array}{c} \xrightarrow{\tau^{-1}} \\ \xleftarrow{\quad} \\ \xrightarrow{\quad} \end{array} & \text{FilAb}[\tau^{-1}] & \text{become} & \text{FilAb} & \begin{array}{c} \xrightarrow{A \mapsto A^{-\infty}} \\ \xleftarrow{\text{Const}} \\ \xrightarrow{A \mapsto A^{\infty}} \end{array} & \text{Ab}. \end{array}$$

Proof. By definition of the limit and colimit, the functor Const admits both a left and right adjoint, being given by the colimit and limit functor, respectively. The claims then follow immediately from the fact that Const restricts to an equivalence $\text{Ab} \simeq \text{FilAb}[\tau^{-1}]$. ■

To allow for a distinction to be made between an abelian group and its corresponding constant filtered abelian group, we introduce the following notation.

Notation 3.4. We write $\tau^{-1}: \text{FilAb} \rightarrow \text{FilAb}[\tau^{-1}]$ for the functor sending A to

$$A[\tau^{-1}] := \text{colim} \left(A \xrightarrow{\tau} A \xrightarrow{\tau} \cdots \right).$$

We write $(-)^{\tau=1}: \text{FilAb} \rightarrow \text{Ab}$ for the composite

$$\text{FilAb} \xrightarrow{\tau^{-1}} \text{FilAb}[\tau^{-1}] \xrightarrow{\simeq} \text{Ab}.$$

Both of these functors are naturally symmetric monoidal.

In terms of $\mathbf{Z}[\tau]$ -modules, these functors take the following form.

Variant 3.5. The subcategory of τ -invertible filtered abelian groups is (symmetric monoidally) equivalent to the category of graded $\mathbf{Z}[\tau^{\pm}]$ -modules. Rephrased like this, the functors τ^{-1} and $(-)^{\tau=1}$ are given, respectively, by sending a graded $\mathbf{Z}[\tau]$ -module M to

$$\mathbf{Z}[\tau^{\pm}] \otimes_{\mathbf{Z}[\tau]} M, \quad \text{respectively} \quad \mathbf{Z} \otimes_{\mathbf{Z}[\tau]} M,$$

where in the latter we let τ act on \mathbf{Z} by the identity.^[1] Note that these two functors are related by the (symmetric monoidal) equivalence

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

given by evaluation at degree zero.

This explains what we mean by the colimit of a filtered abelian group being the same as the τ -inversion of the corresponding $\mathbf{Z}[\tau]$ -module. The claim about completion follows from the following.

Definition 3.6. Let R be a (commutative) ring, let $x \in R$, and let M be an R -module. The x -adic filtration on M is the filtered R -module

$$\cdots \xrightarrow{x} M \xrightarrow{x} M \xrightarrow{x} M \longleftarrow \cdots,$$

which we index to be constant from degree 0 onwards. If M is a commutative R -algebra, then this is naturally a filtered commutative R -algebra.

Proposition 3.7. Let R be a ring, let $x \in R$, and let M be an R -module. Then the x -adic filtration on M is derived complete if and only if M is x -complete as an R -module, i.e., if the natural map

$$M \longrightarrow M_x^\wedge := \lim_k M/x^k$$

is an isomorphism.

Warning 3.8. The above should not be confused with the notion of *derived x -completion* (such as derived p -completion, a.k.a. L -completion; see [HS99, Appendix A]) in the sense of [GM92]. In fact, for filtered abelian groups, derived τ -completion is in general different from τ -completion; see Warning 3.32 below.

Proof of Proposition 3.7. Consider the diagram

$$\begin{array}{ccccc} \vdots & & \vdots & & \vdots \\ \downarrow & & \downarrow & & \downarrow \\ M & \xrightarrow{x^3} & M & \longrightarrow & M/x^3 \\ x \downarrow & & \parallel & & \downarrow \\ M & \xrightarrow{x^2} & M & \longrightarrow & M/x^2 \\ x \downarrow & & \parallel & & \downarrow \\ M & \xrightarrow{x} & M & \longrightarrow & M/x. \end{array}$$

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

$$0 \longrightarrow L \longrightarrow M \longrightarrow M_x^\wedge \longrightarrow K \longrightarrow 0$$

where

$$L = \lim(\cdots \xrightarrow{x} M \xrightarrow{x} M) \quad \text{and} \quad K = \lim^1(\cdots \xrightarrow{x} M \xrightarrow{x} M).$$

^[1]Note that in the second of these cases, we lose a grading, because letting τ act by the identity on \mathbf{Z} does not turn \mathbf{Z} into a *graded* $\mathbf{Z}[\tau]$ -module. Said differently, $(-)^{\tau=1}$ is given by taking the quotient by $\tau - 1$, which is not a homogeneous element, thus resulting in a loss of grading.

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

Corollary 3.9. *A filtered abelian group is derived complete (Definition 2.11) if and only if its corresponding graded $\mathbf{Z}[\tau]$ -module is τ -adically complete.*

Proof. Let A be a filtered abelian group. Applying the previous to the ring $R = \mathbf{Z}[\tau]$ and the element $x = \tau$, the limit and the first-derived limit of the τ -adic filtration on A are the constant filtered abelian groups on the limit and first-derived limit of A , respectively. ■

Variante 3.10. There is an obvious variant of all of the above for graded abelian groups, and all the analogous versions of the previous results hold true. A **filtered graded abelian group** is a functor $\mathbf{Z}^{\text{op}} \rightarrow \text{grAb}$. Note that this by definition means that the transition maps preserve degrees. We write $\text{FilgrAb} := \text{Fun}(\mathbf{Z}^{\text{op}}, \text{grAb})$ for the category of filtered graded abelian groups.

We give grAb the symmetric monoidal structure with the Koszul sign rule. We then regard FilgrAb as a symmetric monoidal category under Day convolution.

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

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

where 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. This indexing is designed to fit with the indexing conventions for spectral sequences from Section 2.3.

3.1.1 Examples

Consider the p -adic filtration (see Definition 3.6) on an abelian group A . The induced strict filtration on A is the **strict p -adic filtration**, given by

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

Note, however, that the maps in the p -adic filtration itself need not be injective, as A might contain p -torsion. The additional information in the p -adic filtration is that it remembers all possible choices of p -divisions of elements.

Example 3.11. 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

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

We think of \tilde{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. As with all $\mathbf{Z}[\tau]$ -modules, the filtration of an element corresponds to the τ -divisibility.

This filtration is not τ -complete however; for this, we would have to pass to \mathbf{Z}_p . ▲

In Remark 2.14, we remarked that in a strictly filtered ring, the filtration of elements is subadditive. The corresponding $\mathbf{Z}[\tau]$ -algebra records this very elegantly.

Example 3.12. Consider the ring

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

We give A a strict 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: the product $4 \cdot \nu$ is 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. Saying that 4ν lands in filtration 3 means that 4ν is the τ^2 -multiple of an element in filtration 3. Due to the lack of τ -torsion, in this case there is a unique such element, namely η^3 . The resulting graded $\mathbf{Z}[\tau]$ -algebra is given by

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

In an informal 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. Observe that if we put $\tau = 1$, then this indeed recovers the original ring A .

Note that, unlike in Example 3.11, we do not write $\tilde{\eta}$ or $\tilde{\nu}$, but instead use the symbols η and ν to directly record the filtration of the elements in the ring A . We do this because, once we fix a filtration on A , we think of the filtration of an element as an intrinsic property, not something to be witnessed by another element. We cannot use this type of notation in Example 3.11 however, because the symbol p is usually reserved for $1 + \cdots + 1$, and it is a bad idea to break this convention. (For instance, using the symbol 2 to denote anything other than $1 + 1$ is not advisable.) ▲

Elements that are in the kernel of a 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*.

3.2 Filtered spectra

By the Yoneda lemma, the natural transition map $\pi_{n,s} \rightarrow \pi_{n,s-1}$ is induced by a map $\mathbf{S}^{n,s-1} \rightarrow \mathbf{S}^{n,s}$.

Definition 3.13.

- (1) The map $\tau: \mathbf{S}^{0,-1} \rightarrow \mathbf{S}$ is the image of the morphism $-1 \rightarrow 0$ in \mathbf{Z} under the functor $i: \mathbf{Z} \rightarrow \text{FilSp}$ from Definition 2.17.
- (2) If X is a filtered spectrum, then tensoring $\tau: \mathbf{S}^{0,-1} \rightarrow \mathbf{S}$ with X results in a map $\Sigma^{0,-1}X \rightarrow X$. We will denote this map by τ_X , or simply by τ if there is little risk of confusion.

In a diagram, writing $\mathbf{S}^{0,-1}$ in the top row and \mathbf{S} in the bottom row, the map τ looks like

$$\begin{array}{ccccccc} \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \mathbf{S} \equiv \cdots \\ & & \downarrow & & \downarrow & & \parallel \\ \cdots & \longrightarrow & 0 & \longrightarrow & \mathbf{S} \equiv \mathbf{S} \equiv \cdots \end{array}$$

If X is a filtered spectrum, then the map τ_X looks like

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X^2 & \xrightarrow{f_1} & X^1 & \xrightarrow{f_0} & X^0 \longrightarrow \cdots \\ & & \downarrow f_1 & & \downarrow f_0 & & \downarrow f_{-1} \\ \cdots & \longrightarrow & X^1 & \xrightarrow{f_0} & X^0 & \xrightarrow{f_{-1}} & X^{-1} \longrightarrow \cdots \end{array}$$

In words: the components of the map τ_X are the transition maps of X .

Remark 3.14. Using the above notation, the functor $i: \mathbf{Z} \rightarrow \text{FilSp}$ from Definition 2.17 can be depicted as the diagram in FilSp given by

$$\cdots \xrightarrow{\tau} \mathbf{S}^{0,-1} \xrightarrow{\tau} \mathbf{S} \xrightarrow{\tau} \mathbf{S}^{0,1} \xrightarrow{\tau} \cdots$$

Previously, we considered the homotopy groups of a filtered spectrum as a filtered graded abelian group. Now, we rephrase this in terms of $\mathbf{Z}[\tau]$ -modules.

Definition 3.15. We define a functor

$$\pi_{*,*}: \text{FilSp} \longrightarrow \text{Mod}_{\mathbf{Z}[\tau]}(\text{bigrAb}), \quad X \longmapsto \pi_{*,*} X$$

where the $\mathbf{Z}[\tau]$ -module structure is given by the map τ of Definition 3.13. This is naturally a lax symmetric monoidal functor, where bigrAb is given the Koszul sign rule according to the first grading.

Our ultimate goal is to make precise that the bigraded $\mathbf{Z}[\tau]$ -module $\pi_{*,*} X$ captures the data of the spectral sequence underlying X . Before we can do this, we start by reformulating the basic building blocks of spectral sequences in terms of τ . We also compare these notions to the analogous notions for $\mathbf{Z}[\tau]$ -modules from the previous sections. Note, however, that these do not always coincide: modding out by τ and completing at τ are different in the stable than in the abelian setting.

3.2.1 Inverting τ

Definition 3.16. A filtered spectrum X is called **τ -invertible** if the map $\tau: \Sigma^{0,-1} X \rightarrow X$ is an isomorphism. We write $\text{FilSp}[\tau^{-1}]$ for the full subcategory of FilSp on the τ -invertible filtered spectra. Write $\tau^{-1}: \text{FilSp} \rightarrow \text{FilSp}$ for the functor sending X to the colimit

$$X[\tau^{-1}] := \text{colim}(X \xrightarrow{\tau} \Sigma^{0,1} X \xrightarrow{\tau} \dots).$$

It is easy to check that τ acts invertibly on $X[\tau^{-1}]$, so that τ -inversion lands in τ -invertible filtered spectra. Moreover, it participates in an adjunction

$$\text{FilSp} \xrightleftharpoons{\tau^{-1}} \text{FilSp}[\tau^{-1}].$$

Inverting τ is a particularly good kind of localisation: it is a *smashing localisation*. We refer to [GGN16, Section 3] for an introduction to such localisations. The practical upshot is that τ -invertible objects are closed under limits, colimits and tensor products, and that τ -local objects get an essentially unique structure of a $\mathbf{S}[\tau^{-1}]$ -module.

Proposition 3.17. *The functor of τ -localisation is a smashing localisation, i.e., it is given by tensoring with the idempotent object $\mathbf{S}[\tau^{-1}]$. In particular, the inclusion functor $\text{FilSp}[\tau^{-1}] \subseteq \text{FilSp}$ preserves colimits and has a further right adjoint.*

Proof. The tensor product of filtered spectra preserves colimits. It follows that

$$\begin{aligned} X[\tau^{-1}] &= \text{colim}(X \xrightarrow{\tau} \Sigma^{0,1} X \xrightarrow{\tau} \dots) \\ &\cong \text{colim}(\mathbf{S} \xrightarrow{\tau} \mathbf{S}^{0,1} \xrightarrow{\tau} \dots) \otimes X \\ &= \mathbf{S}[\tau^{-1}] \otimes X. \end{aligned} \quad \blacksquare$$

The notion of a τ -invertible filtered spectrum is not new: it is the same as a filtered spectrum whose transition maps are invertible. More precisely, we have the following identifications.

Proposition 3.18. *The symmetric monoidal functor $\text{Const}: \text{Sp} \rightarrow \text{FilSp}$ restricts to an equivalence onto the τ -invertible filtered spectra:*

$$\text{Const}: \text{Sp} \xrightarrow{\simeq} \text{FilSp}[\tau^{-1}].$$

Under this equivalence, the adjunctions

$$\begin{array}{ccc} \text{FilSp} & \begin{array}{c} \xrightarrow{\tau^{-1}} \\ \longleftarrow \\ \xrightarrow{\quad} \\ \longleftarrow \\ \xrightarrow{\quad} \end{array} & \text{FilSp}[\tau^{-1}] & \text{become} & \text{FilSp} & \begin{array}{c} \xrightarrow{X \mapsto X^{-\infty}} \\ \xrightarrow{\text{Const}} \\ \xrightarrow{\quad} \\ \xrightarrow{X \mapsto X^{\infty}} \end{array} & \text{Sp}. \end{array}$$

Paralleling Notation 3.4, we will sometimes use the following notation to distinguish between the two equivalent ∞ -categories Sp and $\text{FilSp}[\tau^{-1}]$. In practice however, we may refer to both functors as “ τ -inversion”.

Notation 3.19. We write $(-)^{\tau=1}$ for the composite

$$\text{FilSp} \xrightarrow{\tau^{-1}} \text{FilSp}[\tau^{-1}] \xrightarrow{\cong} \text{Sp}.$$

Remark 3.20. The operation of τ -inversion of filtered spectra is compatible with τ -inversion of filtered (graded) abelian groups from Section 3.1. More specifically, if X is a filtered spectrum, then the natural map provides an isomorphism

$$(\pi_{n,*} X)[\tau^{-1}] \xrightarrow{\cong} \pi_{n,*}(X[\tau^{-1}]),$$

due to the fact that bigraded homotopy groups preserve filtered colimits (as the filtered spheres are compact). As a result, this also induces an isomorphism

$$(\pi_{n,*} X)^{\tau=1} \cong \pi_n(X^{\tau=1}).$$

Remark 3.21 (Detection and τ -divisibility). Let X be a filtered spectrum, and let $\theta \in \pi_n X^{\tau=1}$ be nonzero. The statement that θ is detected (see Definition 2.36) in the spectral sequence underlying X in filtration s translates to the statement that θ lifts to a non- τ -divisible element α in $\pi_{n,s} X$. Indeed, saying that α is not τ -divisible is another way of saying that α does not lift to $\pi_{n,s+1} X$, which is equivalent to its image in $E_1^{n,s}$ not being zero, which is implied by the definition of detection from Definition 2.36.

Example 3.22. Recall the definition $\mathbf{S}^{n,s} = \Sigma^n i(s)$ from Definition 2.19. It is immediate from the definition of i that we have a natural isomorphism $i(-)[\tau^{-1}] \cong \text{Const } \mathbf{S}$. As τ -inversion and Const are exact functors, they preserve suspensions, so we find that

$$\mathbf{S}^{n,s}[\tau^{-1}] = \Sigma^n i(s)[\tau^{-1}] \cong \Sigma^n \text{Const } \mathbf{S} \cong \text{Const } \mathbf{S}^n.$$

In other words, $(\mathbf{S}^{n,s})^{\tau=1} \cong \mathbf{S}^n$. We can think of this as saying that inverting τ forgets the filtration. ▲

3.2.2 Modding out by τ

Notation 3.23. Let $k \geq 1$.

- ◆ We write $C\tau^k$ for the cofibre of the map $\tau^k: \mathbf{S}^{0, -k} \rightarrow \mathbf{S}$.
- ◆ If X is a filtered spectrum, then we write X/τ^k for $C\tau^k \otimes X$.

Concretely, $C\tau^k$ is the filtered spectrum

$$\cdots \longrightarrow 0 \longrightarrow \mathbf{S} \longleftarrow \cdots \longleftarrow \mathbf{S} \longrightarrow 0 \longrightarrow \cdots,$$

where the nonzero terms are in filtrations $0, -1, \dots, -k + 1$. If X is a filtered spectrum, then X/τ^k is in filtration s given by the cofibre of $X^{s+k} \rightarrow X^s$.

Unlike in the case of filtered abelian groups, in higher algebra, taking quotients in a monoidal way is a treacherous matter. In this case, it turns out we can do this in the best possible way.

Theorem 3.24 (Lurie). *For every $k \geq 0$, the filtered spectrum $C\tau^k$ admits (uniquely up to contractible choice) the structure of a filtered \mathbf{E}_∞ -ring such that its unit map $\mathbf{S} \rightarrow C\tau^k$ is an isomorphism in filtrations $0, -1, \dots, -k + 1$.*

Proof. See [Lur15, Proposition 3.2.5], bearing in mind that Lurie writes $\text{Rep}(\mathbf{Z})$ for FilSp , and writes \mathbb{A} for $C\tau$. ■

Using this ring structure on $C\tau$, we can make precise the way in which tensoring with $C\tau$ recovers the associated graded. We first define a functor

$$d: \text{grSp} \longrightarrow \text{FilSp}$$

given by left Kan extension along the functor $\mathbf{Z}^{\text{discr}} \rightarrow \mathbf{Z}^{\text{op}}$, with $\mathbf{Z}^{\text{discr}}$ denoting the discrete category with objects \mathbf{Z} . Informally, this functor is given by sending a graded spectrum $(X_n)_n$ to the filtered spectrum

$$\cdots \longrightarrow \bigoplus_{n \geq 1} X_n \longrightarrow \bigoplus_{n \geq 0} X_n \longrightarrow \bigoplus_{n \geq -1} X_n \longrightarrow \cdots,$$

with maps the natural inclusions. Being defined as the left Kan extension along a symmetric monoidal functor, this is naturally a symmetric monoidal functor.

Theorem 3.25 (Lurie). *The composite*

$$\text{grSp} \xrightarrow{d} \text{FilSp} \xrightarrow{C\tau \otimes -} \text{Mod}_{C\tau}(\text{FilSp})$$

is a symmetric monoidal equivalence. Moreover, this equivalence fits into a commutative diagram

$$\begin{array}{ccc} \text{FilSp} & \xrightarrow{\text{Gr}} & \text{grSp} \\ & \searrow^{C\tau \otimes -} & \downarrow \simeq \\ & & \text{Mod}_{C\tau}(\text{FilSp}). \end{array}$$

Proof. See [Lur15, Proposition 3.2.7], bearing in mind that Lurie writes $\text{Rep}(\mathbf{Z})$ for FilSp , writes $\text{Rep}(\mathbf{Z}^{\text{ds}})$ for grSp , writes \mathbb{A} for $C\tau$, and writes I for d . ■

Remark 3.26. The above in particular puts a symmetric monoidal structure on the associated graded functor $\text{Gr}: \text{FilSp} \rightarrow \text{grSp}$, because the functor $C\tau \otimes -: \text{FilSp} \rightarrow \text{Mod}_{C\tau}(\text{FilSp})$ is canonically symmetric monoidal. One could have also done this more directly: see [Hed20, Section II.1.3].

Warning 3.27. Being a module over $C\tau$ is not a property, but additional structure. One can see this by observing that $C\tau$ is not an idempotent, i.e., that $C\tau \otimes C\tau$ is not isomorphic to $C\tau$: we instead have

$$C\tau \otimes C\tau \cong C\tau \oplus \Sigma^{1,-1} C\tau.$$

Indeed, the map τ on $C\tau$ is nullhomotopic (because $C\tau$ is a ring), so the cofibre sequence defining $C\tau$ splits after tensoring with $C\tau$. Alternatively, note that the associated graded of $C\tau$ is concentrated in filtrations 0 and -1 , where it is \mathbf{S} and \mathbf{S}^1 , respectively. Via the equivalence of Theorem 3.25, this precisely gives us the above splitting.

Warning 3.28. The notion of modding out by τ in the spectral setting is decidedly different from modding out by τ in the setting of filtered abelian groups. Indeed, this is exactly the difference between a cofibre and a cokernel: the former also sees the kernel.

3.2.3 Completing at τ

Definition 3.29. The functor of **τ -adic completion** (or τ -completion for short) is $C\tau$ -localisation of FilSp , i.e., inverting those maps that become an isomorphism after tensoring with $C\tau$.

We write FilSp_τ^\wedge for the full subcategory of FilSp on the τ -adically complete filtered spectra. This results in an adjunction

$$\text{FilSp} \begin{array}{c} \xrightarrow{(-)_\tau^\wedge} \\ \longleftarrow \end{array} \text{FilSp}_\tau^\wedge.$$

This notion is, in fact, nothing but the notion of completeness previously introduced in Definition 2.28.

Proposition 3.30.

- (1) A filtered spectrum X is τ -adically complete if and only if its limit X^∞ vanishes.
- (2) A map of filtered spectra is a $C\tau$ -equivalence if and only if it induces an isomorphism on associated graded.
- (3) For every filtered spectrum X , the natural map $X \rightarrow \text{cofib}(X^\infty \rightarrow X)$ is the τ -completion of X .

Proof. By definition, a filtered spectrum X is τ -complete if and only if $\text{Map}(Y, X)$ is contractible for all τ -invertible Y . The first claim therefore follows immediately from Proposition 3.18. The second item is immediate from Theorem 3.25. For the final claim, we have to show that $\text{cofib}(X^\infty \rightarrow X)$ is τ -complete and that the map from X into it is a $C\tau$ -equivalence. The former is clear, and the latter follows from the tetrahedral axiom. ■

The operations of inverting and completing at τ are related in the following way.

Proposition 3.31. *For a filtered spectrum X , there is a natural pullback square of lax symmetric monoidal functors*

$$\begin{array}{ccc} X & \longrightarrow & X_\tau^\wedge \\ \downarrow & \lrcorner & \downarrow \\ X[\tau^{-1}] & \longrightarrow & (X_\tau^\wedge)[\tau^{-1}]. \end{array}$$

In particular, a map of filtered spectra $X \rightarrow Y$ is an isomorphism if and only if the maps

$$X[\tau^{-1}] \longrightarrow Y[\tau^{-1}] \quad \text{and} \quad C\tau \otimes X \longrightarrow C\tau \otimes Y$$

are both an isomorphism.

Proof. This is a standard result; see [Man21, Proposition 4.1.1]. ■

Concretely, this says that a map of filtered spectra is an isomorphism if and only if it is an isomorphism on the colimit and on the associated graded.

Finally, we compare τ -completeness of filtered spectra and filtered abelian groups.

Warning 3.32. If X is a filtered spectrum, then τ -completeness of X need not imply τ -completeness of the $\mathbf{Z}[\tau]$ -module $\pi_{*,*} X$. Indeed, this is part of the discussion of convergence. In detail: using Corollary 3.9, we see that part (a) of the definition of strong convergence in Definition 2.44 is asking $\pi_{*,*} X$ to be τ -complete. As explained in Warning 2.52, the τ -completeness of X (i.e., the vanishing of X^∞) need not imply this. The conditional convergence criteria of Boardman from Theorem 2.55 and Remark 2.57 give the further conditions needed to go from τ -completeness to τ -completeness of $\pi_{*,*} X$.

3.3 Total differentials

Our next topic concerns the differentials in the spectral sequence. Although the previous concepts involving τ are reformulations of ones from Chapter 2, the *total differentials* to be introduced in this section did not make an appearance there. However, this is not because one cannot phrase these without τ , but because we find that τ provides for an easier notational setup to introduce these concepts.

Many authors have used total differentials in both synthetic and motivic spectra; see [BHS23], [Chu22] and [Isa+24], for example. They naturally arise in the filtered setting

as well, and they are one of the big benefits of working at the filtered level: they lead to strengthened versions of the usual Leibniz rule. Further, knowledge of total differentials allows one to deduce hidden extensions from non-hidden extensions. We explain these applications at the end of this section.

Notation 3.33. Let X be a filtered spectrum, and let $n \geq 1$. Write ∂_n^∞ for the boundary map in the cofibre sequence

$$\Sigma^{0, -n} X \xrightarrow{\tau^n} X \longrightarrow X/\tau^n \xrightarrow{\partial_n^\infty} \Sigma^{1, -n} X.$$

For $N \geq n$, write ∂_n^N for the boundary map in

$$\Sigma^{0, -n} X/\tau^{N-n} \xrightarrow{\tau^n} X/\tau^N \longrightarrow X/\tau^n \xrightarrow{\partial_n^N} \Sigma^{1, -n} X/\tau^{N-n}.$$

We call ∂_1^∞ the **total differential** on X , and call ∂_1^N the **N -truncated total differential** on X .

The map ∂_n^N captures information about the d_n, \dots, d_{N-1} -differentials in X . Decreasing N results in a loss of information, but increasing the lower index n should be thought of as an increase of information: we will see that, roughly speaking, ∂_n^N is only defined on elements on which the differentials d_1, \dots, d_{n-1} vanish.

We will now make these ideas precise. Before we begin, let us point out what these maps are concretely. As colimits in FilSp are computed levelwise, if we evaluate the cofibre sequence defining ∂_1^∞ at filtration s , we obtain the cofibre sequence

$$X^{s+1} \longrightarrow X^s \longrightarrow \text{Gr}^s X \longrightarrow \Sigma X^{s+1},$$

where the first map is the transition map. Recall from Appendix A that the boundary map $\text{Gr}^s \rightarrow \Sigma X^{s+1}$ is precisely the map used to define all the differentials in the underlying spectral sequence. The other total differentials are variants on this map, and a similar diagram chase will make precise the intuition for ∂_n^N we gave above.

With this in mind as our intuition, we proceed to the formal proofs.

Proposition 3.34. *Let X be a filtered spectrum, let $n \geq k \geq 1$, and let $\infty \geq N \geq n$. Then we have commutative diagrams*

$$\begin{array}{ccc} X/\tau^n & \xrightarrow{\partial_n^\infty} & \Sigma^{1, -n} X \\ & \searrow \partial_n^N & \downarrow \\ & & \Sigma^{1, -n} X/\tau^{N-n} \end{array} \quad \text{and} \quad \begin{array}{ccc} X/\tau^n & \xrightarrow{\partial_n^\infty} & \Sigma^{1, -n} X \\ \downarrow & & \downarrow \tau^{n-k} \\ X/\tau^k & \xrightarrow{\partial_k^\infty} & \Sigma^{1, -k} X, \end{array}$$

where the unlabelled maps are the reduction maps. In words: ∂_n^N is the mod τ^{N-n} reduction of ∂_n^∞ , and $\tau^{n-k} \cdot \partial_n^\infty = \partial_k^\infty$.

Proof. For readability, we omit the bigraded suspensions in this proof. We start with the commutative diagram

$$\begin{array}{ccc} X & \xlongequal{\quad} & X \\ \tau^{N-n} \downarrow & & \downarrow \tau^N \\ X & \xrightarrow{\tau^n} & X. \end{array}$$

Taking pushouts in the vertical direction once, and repeatedly in the horizontal direction, we arrive at a commutative diagram

$$\begin{array}{ccccccc} X & \xrightarrow{\tau^n} & X & \longrightarrow & X/\tau^n & \xrightarrow{\partial_n^\infty} & X \\ \downarrow & & \downarrow & & \parallel & & \downarrow \\ X/\tau^{N-n} & \xrightarrow{\tau^n} & X/\tau^N & \longrightarrow & X/\tau^n & \xrightarrow{\partial_n^N} & X/\tau^{N-n}. \end{array}$$

The right-most square is the first claimed diagram.

The second diagram comes from the commutative diagram

$$\begin{array}{ccccccc} X & \xrightarrow{\tau^n} & X & \longrightarrow & X/\tau^n & \xrightarrow{\partial_n^\infty} & X \\ \tau^{n-k} \downarrow & & \parallel & & \downarrow & & \downarrow \tau^{n-k} \\ X & \xrightarrow{\tau^k} & X & \longrightarrow & X/\tau^k & \xrightarrow{\partial_k^\infty} & X \end{array}$$

obtained by taking horizontal pushouts of the left-most square. ■

Next, we relate the τ -divisibility of the (truncated) total differentials to the vanishing of differentials in the underlying spectral sequence.

For the case of truncated total differentials, we run into the subtlety that there are two kinds of τ -multiples in $\pi_{*,*} X/\tau^k$. Namely, we can either speak of the τ -multiple of an element from X/τ^k , or of the τ -multiple of an element in X/τ^{k-1} regarded as an element of X/τ^k . The latter of these is more general. This difference will come up a number of times, particularly when describing the structure of $\pi_{*,*} X/\tau^k$ later in Section 3.5.1. We use the following notation to distinguish between these.

Notation 3.35. Let X be a filtered spectrum, and let $k \geq m \geq 0$ be integers.

- ◆ For $\theta \in \pi_{*,*} X/\tau^k$, we write $\tau^m \cdot \theta$ for the τ^m -multiple of θ in the $\mathbf{Z}[\tau]$ -module $\pi_{*,*}(X/\tau^k)$.
- ◆ For $\theta \in \pi_{*,*} X/\tau^{m-k}$, we write $\tau^m(\theta)$ for the image of θ under the map

$$\tau^m : \Sigma^{0,-m} X/\tau^{k-m} \longrightarrow X/\tau^k.$$

Both versions have their advantages and disadvantages. The former of the two is, by definition, captured by the $\mathbf{Z}[\tau]$ -module $\pi_{*,*}(X/\tau^k)$, while the map τ^m in the latter of the

two participates in a cofibre sequence

$$\Sigma^{0,-m} X/\tau^{k-m} \longrightarrow X/\tau^k \longrightarrow X/\tau^m,$$

and as a result is closely tied to the truncated total differential ∂_m^k . They are, however, related in the following way.

Remark 3.36. There is a commutative diagram

$$\begin{array}{ccc} \Sigma^{0,-m} X/\tau^k & \xrightarrow{\tau^m} & X/\tau^k \\ \downarrow & \nearrow \tau^m & \\ \Sigma^{0,-m} X/\tau^{k-m} & & \end{array}$$

where the left vertical map is the reduction. It follows that for $\theta \in \pi_{*,*} X/\tau^k$, the elements $\tau^m \cdot \theta$ and $\tau^m(\theta)$ (where in the latter, we implicitly reduce $\theta \bmod \tau^{k-m}$) are the same. The difference between the notations of Notation 3.35, then, is that there could be more elements of the form $\tau^m(\alpha)$ (where $\alpha \in \pi_{*,*} X/\tau^{k-m}$) than elements of the form $\tau^m \cdot \theta$ (where $\theta \in \pi_{*,*} X/\tau^k$).

Proposition 3.37. *Let X be a filtered spectrum, and let $x \in E_1^{n,s}$. View x as an element of $\pi_{n,s} X/\tau$. Let $r \geq 0$.*

- (1a) *If $\partial_1^\infty(x)$ is τ^r -divisible, then $d_1(x), \dots, d_r(x)$ vanish.*
- (1b) *More generally, if $k > r$ and $\partial_1^k(x)$ is of the form $\tau^r(\alpha)$ for some $\alpha \in \pi_{*,*} X/\tau^{k-r-1}$, then $d_1(x), \dots, d_r(x)$ vanish.*
- (2a) *Suppose that $\alpha \in \pi_{n-1, s+r+1} X$ is an element such that $\tau^r \cdot \alpha = \partial_1^\infty(x)$. Then the mod τ reduction of α is a representative for $d_{r+1}(x)$ in $E_1^{n-1, s+r+1}$.*
- (2b) *More generally, if $k > r + 1$ and $\alpha \in \pi_{*,*} X/\tau^{k-r-1}$ is such that $\tau^r(\alpha) = \partial_1^k(x)$, then the mod τ reduction of α is a representative for $d_{r+1}(x)$.*

Proof. We begin with the claims regarding ∂_1^∞ . These are a rephrasing of the definition of the differentials in the spectral sequence associated to X . Indeed, evaluating the cofibre sequence of filtered spectra

$$\Sigma^{0,-r} X \xrightarrow{\tau^r} X \longrightarrow X/\tau^r \xrightarrow{\partial_r^\infty} \Sigma^{1,-r} X$$

at filtration s is exactly the cofibre sequence of spectra

$$X^{s+r} \longrightarrow X^s \longrightarrow \text{Gr}^s X \longrightarrow \Sigma X^{s+r}.$$

Finally, the analogous claims for ∂_1^k follow by combining the ones for ∂_1^∞ with Proposition 3.34. ■

Later, when we have more of an understanding of how $\pi_{*,*} X/\tau^k$ relates to the spectral sequence, we will be able to rephrase the above result in a convenient way; see Construction 3.68 and Corollary 3.69.

Warning 3.38. The converse of either item (1a) or (1b) is not true in general. The reason is that the r -th differential is only well defined up to the images of shorter differentials. As a result, one cannot in general use $d_r(x) = 0$ to deduce that $\partial_1^\infty(x)$ is τ^r -divisible if $r > 1$. This can be done, of course, if previous differentials vanish in the appropriate range: more specifically, if bidegree $(n-1, s+r)$ receives no differentials of length shorter than r .

3.3.1 Applications

The reason for going to the trouble of using total differentials is that they allow for more sophisticated differential-deduction techniques. Particularly, it allows one to deduce longer differentials from shorter ones.

A simple example of this is to use linearity of the total differential.

Proposition 3.39 (Linearity of the total differential). *Let $n \geq 1$, and let $n \leq N \leq \infty$. Let X be a (left) homotopy-module over a homotopy-associative ring R in FilSp . Then the map $\pi_{*,*} \partial_n^N$ on X is $\pi_{*,*} R$ -linear. If $N < \infty$, then the map $\pi_{*,*} \partial_n^N$ is also $\pi_{*,*} R/\tau^N$ -linear.*

Proof. This is immediate. ■

A more powerful version of this is the following result, due to Burklund [Bur22, Chapter 3]. It is also referred to as the *synthetic Leibniz rule*, but as we will see, the synthetic version follows directly from the filtered version.

Theorem 3.40 (Filtered Leibniz rule, Burklund). *Let R be a homotopy-associative ring in FilSp . For any $n \geq 1$, the map*

$$\partial_n^{2n} : \pi_{*,*}(R/\tau^n) \longrightarrow \pi_{*-1, *+n}(R/\tau^n)$$

is a derivation. In particular, for any two classes $x, y \in \pi_{,*}(R/\tau^n)$, we have the relation*

$$\partial_n^{2n}(xy) = \partial_n^{2n}(x) \cdot y + (-1)^{|x|} x \cdot \partial_n^{2n}(y).$$

Proof. See [CDvN24, Theorem 2.34]. ■

Because mod τ reduction is a ring map, it follows from this that ∂_n^{n+1} is also a derivation. This recovers the ordinary Leibniz rule for d_n , or more precisely, for a first-page representative of d_n .

Finally, we discuss how knowledge of total differentials is, to some extent, the same as hidden extensions. This was used in, e.g., [BHS23, Proof of Proposition A.20 (14)] and [Mar24, Remark 4.1.3], and is explained in detail in [Isa+24, Method 2.17, Example 2.18, Proposition 4.5].

Remark 3.41 (Total differentials and hidden extensions). Fix a filtered homotopy-ring spectrum X . Let $x, y \in E_1^{*,*}$ be elements, and write $\alpha = \partial_1^\infty(x)$ and $\beta = \partial_1^\infty(y)$. In other words, we have a differential on x hitting (the mod τ reduction of) α , and likewise from y to (the mod τ reduction of) β , possibly of different lengths. We describe a technique for deducing (possibly hidden) extensions between α and β from a (non-hidden) extension between x and y on the E_1 -page.

Let $\theta \in \pi_{*,*} X$ be another element, write $t \in E_1^{*,*}$ for its mod τ reduction, and suppose that we have a multiplicative relation in $E_1^{*,*}$ given by

$$t \cdot x = y.$$

Using linearity of ∂_1^∞ over $\pi_{*,*} X$, we find that

$$\beta = \partial_1^\infty(y) = \partial_1^\infty(t \cdot x) = \partial_1^\infty(\theta \cdot x) = \theta \cdot \partial_1^\infty(x) = \theta \cdot \alpha.$$

In words: the extension $t \cdot x = y$ allows us to deduce the extension $\theta \cdot \alpha = \beta$, provided that we know that α and β are the total differentials on x and y , respectively. In fact, reading the argument backwards, we see that knowing the relation $\theta \cdot \alpha = \beta$ and the total differential $\partial_1^\infty(x) = \alpha$ allows us to deduce the total differential on y .

The resulting extension $\theta \cdot \alpha = \beta$ corresponds to a hidden extension when the differential on y is longer than the differential on x . Indeed, set up properly, the length of the differential corresponds to the τ -divisibility of the total differential. If $\alpha = \tau^r \alpha'$ and $\beta = \tau^s \beta'$, where α' and β' are not τ -divisible, and $0 < r < s$, then we learn that

$$\theta \cdot \alpha' = \tau^{s-r} \beta' \quad \text{up to } \tau^r\text{-divisible elements.}$$

Modulo τ , this relation is hidden: the right-hand side reduces to zero modulo τ . Yet, the relation $t \cdot x = y$ is *not* a hidden extension, so that the knowledge of total differentials allows us to turn non-hidden extensions into hidden extensions between the targets of the differentials.

The only downside of this approach is that it requires α and β to be images of a total differential, or equivalently, that they have to be τ -power torsion (in other words, they come from classes that are hit by differentials). Such classes define the zero element in $\pi_* X^{-\infty}$, making it seem like this method has limited use for deducing extensions between elements in $\pi_* X^{-\infty}$. In practice, one can get around this using the following trick. If we start with non- τ -power torsion elements α and β we would like to find a relation between, we may be able to find another element γ such that $\alpha \cdot \gamma$ and $\beta \cdot \gamma$ are τ -power torsion. By exactness this must mean they are in the image of a total differential. If the sources of these total differentials are related by a multiplicative extension on the E_1 -page, then (up to dividing by γ) we learn about a (possibly hidden) extension between α and β .

This method can be modified with other total differentials in the place of ∂_1^∞ . If we work with ∂_1^N , then we can learn about hidden extensions of length at most $N - 1$. These

truncated total differentials may be easier to obtain, giving it a practical upshot at the cost of a less powerful outcome. On the other hand, we can also work with ∂_n^∞ (or even ∂_n^N). The key difference is that there, we can take multiplicative relations X/τ^n as input, which may themselves correspond to hidden extensions of length smaller than n . In this way, this technique becomes a method to turn hidden extensions into hidden extensions of a possibly greater length.

Example 3.42. This is a slightly simplified example of a hidden extension in Tmf that is shown to hold in [CDvN24]. In Tmf , there is a 2-extension in the 110-stem:

$$2 \cdot \kappa_4 = \eta_1^2 \bar{\kappa}^3. \quad (3.43)$$

In the DSS for Tmf , this is a hidden extension, with filtration jump 12: the element κ_4 has filtration 2, the element η_1 has filtration 1, and the element $\bar{\kappa}$ has filtration 4. Working in the filtered spectrum giving rise to this DSS, we have canonical lifts of these elements to their respective filtrations, which we will denote by the same name. Although neither κ_4 nor $\eta_1^2 \bar{\kappa}^3$ is τ -power torsion, they become so after multiplying with $\bar{\kappa}^3$: we even have the total differentials

$$\partial_4^{28}(2\nu\Delta^7) = \tau^8 \kappa_4 \bar{\kappa}^3 \quad \text{and} \quad \partial_4^{28}(\tau^2 \eta^3 \Delta^7) = \tau^{20} \eta_1^2 \bar{\kappa}^6.$$

The hidden extension $4\nu = \tau^2 \eta^3$, which is of length 2, now stretches to a length 12 extension:

$$2 \cdot \tau^8 \kappa_4 \bar{\kappa}^3 = \partial_4^{28}(4\nu\Delta^7) = \partial_4^{28}(\tau^2 \eta^3 \Delta^7) = \tau^{20} \eta_1^2 \bar{\kappa}^6.$$

From this, the desired extension (3.43) follows. ▲

3.4 Digression: the τ -Bockstein spectral sequence

To prove the Omnibus Theorem, we will use the τ -Bockstein spectral sequence of a filtered spectrum. This is a spectral sequence that compute the bigraded homotopy of a filtered spectrum X . Its usefulness is due to the fact that the τ -Bockstein differentials can be identified with the differentials in the spectral sequence underlying X .

This is one of the more technical sections in these notes. In this text, we only need this spectral sequence for the proof of the Omnibus Theorem in the next section, so the reader who is willing to take the statements of the Omnibus Theorem on faith does not need to read this section. Although one could have instead phrased that proof with a more hands-on argument, the τ -Bockstein spectral sequence is a useful device in and of itself, so we thought it worthwhile to give an account of this spectral sequence.

Remark 3.44. The setup of the τ -Bockstein spectral sequence is, in some sense, the most general setup of a (stable) Bockstein spectral sequence. Through the use of deformations of Section 3.6 (specifically Proposition 3.82), one can recover other Bockstein spectral sequences; see Example 3.93 for the case of the p -Bockstein spectral sequence of spectra, for instance.

Definition 3.45. Let X be a filtered spectrum. The τ -Bockstein filtration on X , denoted $\text{BF}_\tau X$, is the bifiltered spectrum $\mathbf{Z}^{\text{op}} \rightarrow \text{FilSp}$ given by

$$\dots \xrightarrow{\tau} \Sigma^{0,-2} X \xrightarrow{\tau} \Sigma^{0,-1} X \xrightarrow{\tau} X \longleftarrow \dots$$

indexed to be constant from filtration 0 onwards.

Remark 3.46. It is by design that the filtration becomes constant from filtration 0 onwards. This way, its underlying object is X , and consequently this filtration helps us understand $\pi_{*,*} X$. If we had continued the pattern of placing τ 's going all the way to the right, then the underlying object would be $X[\tau^{-1}]$, and in fact, we would learn no more than if we had used the underlying spectral sequence of X directly.

This leads to a spectral sequence in the same way as for singly filtered spectra, except that the indexing is more involved. To help us index it, we use the philosophy from Remark 2.30. First of all, this means that the filtration variable s records the location in the diagram as depicted above. Second, homotopy groups in FilSp are naturally bigraded, so we would like to understand $\pi_{n,w}$ of the colimit. Accordingly, we apply $\pi_{n,w}$ to the above diagram, and study the resulting behaviour using the long exact sequences involving the associated graded.

To avoid potential confusion regarding indexing, we discuss this spectral sequence in the case of a general bifiltered spectrum. In this text, we will only apply this in the case of (variants of) the τ -Bockstein filtration.

Construction 3.47. Let $X: \mathbf{Z}^{\text{op}} \rightarrow \text{FilSp}$ be a bifiltered spectrum. We define a trigraded exact couple

$$A^{n,w,s} = \pi_{n,w}(X^s) \quad \text{and} \quad E^{n,w,s} = \pi_{n,w}(\text{Gr}^s X),$$

fitting into the following diagram, where each map is annotated by its tridegree.

$$\begin{array}{ccc} A^{n,w,s} & \xrightarrow{(0,0,-1)} & A^{n,w,s} \\ & \swarrow (-1,0,1) & \searrow (0,0,0) \\ & E^{n,w,s} & \end{array}$$

In this indexing, the differential d_r has tridegree $(-1, 0, r)$: indeed, a d_r -differential is given by applying a boundary map of degree $(-1, 0, 1)$ once, lifting against a map of degree $(0, 0, -1)$ a total of $r - 1$ times, and then projecting down by a map of degree zero. The resulting spectral sequence is of signature

$$E_1^{n,w,s} = \pi_{n,w}(\text{Gr}^s X) \implies \pi_{n,w} X^{-\infty}.$$

The induced strict filtration on $\pi_{n,w} X^{-\infty}$ is

$$F^s \pi_{n,w} X^{-\infty} = \text{im}(\pi_{n,w} X^s \longrightarrow \pi_{n,w} X^{-\infty}).$$

Example 3.48. Let X be a filtered spectrum. We refer to the trigraded spectral sequence underlying $\mathrm{BF}_\tau X$ as the **τ -Bockstein spectral sequence** (τ -BSS) of X . Plugging in the definitions, we see that

$$\pi_{n,w}(\mathrm{BF}_\tau^s X) = \pi_{n,w}(\Sigma^{0,-s} X) = \pi_{n,w+s} X \quad \text{and} \quad \mathrm{Gr}^s(\mathrm{BF}_\tau X) = \Sigma^{0,-s} X/\tau.$$

In particular, the first page is of the form

$$E_1^{n,w,s} = \begin{cases} \pi_{n,w+s}(X/\tau) & \text{if } s \geq 0, \\ 0 & \text{else.} \end{cases} \quad (3.49)$$

The induced strict filtration on $\pi_{*,*} X$ is

$$F^s \pi_{n,w} X = \mathrm{im}(\tau^s : \pi_{n,w+s} X \longrightarrow \pi_{n,w} X). \quad \blacktriangle$$

It will be useful to think of an element in filtration s in the τ -BSS as a formal τ^s -multiple.

Notation 3.50. We define a formal variable $\bar{\tau}$ to have tridegree $(0, -1, 1)$. Let X be a filtered spectrum, and $\{E_r^{*,*,*}\}$ denote its τ -BSS. We put a $\mathbf{Z}[\bar{\tau}]$ -module structure on $E_1^{*,*,*}$ in such a way that the isomorphism (3.49) becomes an isomorphism trigraded $\mathbf{Z}[\bar{\tau}]$ -modules

$$E_1^{*,*,*} \cong \mathbf{Z}[\bar{\tau}] \otimes \pi_{*,*}(X/\tau),$$

where $\pi_{n,w}(X/\tau)$ is placed in tridegree $(n, w, 0)$.

Remark 3.51. Even though the τ -BSS is a trigraded spectral sequence, one can to a certain extent depict it as a bigraded one, as follows. Since the d_r^τ -differential has tridegree $(-1, 0, s)$, by fixing a constant value for w , we get a bigraded spectral sequence trying to converge to $\pi_{*,w} X$. An element in filtration s will be a formal τ^s -multiple, the only catch being that the class it is a formal multiple of lives in a spectral sequence for a different w -value (namely $w + s$).

The behaviour of the τ -Bockstein spectral sequence is the same as any Bockstein spectral sequence: differentials capture τ -power torsion in $\pi_{*,*} X$, where a differential of length r corresponds to τ^r -torsion. What is special to the τ -BSS of X is that the differentials are exactly the differentials in the spectral sequence underlying X . More specifically, a differential $d_r(x) = y$ in the spectral sequence underlying X corresponds to a τ -Bockstein differential

$$d_r^\tau(x) = \bar{\tau}^r \cdot y.$$

The insertion of this $\bar{\tau}$ -power means that, instead of killing elements directly, the differentials in the spectral sequence underlying X introduce τ -power torsion in $\pi_{*,*} X$. Together with $\bar{\tau}$ -linearity of the d_r^τ -differentials, this determines all differentials in the τ -BSS.

Stating this in a precise way requires some care. For the sake of completeness, we prove this here in detail. We base the statement of this result on [Pal05] and [BHS23,

Appendix A]. The full statement about τ -power torsion will be given by the Omnibus Theorem of the next section.

In the statement of the following theorem, all claims about differentials should be interpreted as up to boundaries of lower differentials. That is, we allow ourselves to speak of a differential $d_r(x) = y$ where x and y are elements of E_1 that are $d_{\leq r-1}$ -cycles, implicitly considering these elements as defining classes in E_r .

Theorem 3.52. *Let X be a filtered spectrum. Let $\{E_r^{*,*}, d_r\}$ denote the spectral sequence underlying X , and let $\{E_r^{*,*}, d_r^\tau\}$ denote the τ -Bockstein spectral sequence of X .*

- (1) *There is a natural isomorphism of trigraded $\mathbf{Z}[\bar{\tau}]$ -modules*

$$E_1^{*,*} \cong \mathbf{Z}[\bar{\tau}] \otimes \pi_{*,*}(X/\tau) = \mathbf{Z}[\bar{\tau}] \otimes E_1^{*,*},$$

where $\pi_{n,w}(X/\tau)$ is placed in tridegree $(n, w, 0)$ and $\bar{\tau}$ has tridegree $(0, -1, 1)$.

- (2) *The differentials are $\bar{\tau}$ -linear and $\bar{\tau}$ -divisible: for $x \in E_1^{n,w,s}$ and $y \in E_1^{n-1,w,s+r}$, there is a differential*

$$d_r^\tau(x) = y,$$

if and only if for any (hence all) $m \geq 0$, there is a differential

$$d_r^\tau(\bar{\tau}^m x) = \bar{\tau}^m y.$$

In particular, the $\mathbf{Z}[\bar{\tau}]$ -module structure on the first page induces a $\mathbf{Z}[\bar{\tau}]$ -module structure on later pages.

- (3) *The target of a d_r^τ -differential is a $\bar{\tau}^r$ -multiple: for every $x \in E_r^{n,w,s}$, there is an element $y \in E_r^{n-1,w+r,s}$ such that*

$$d_r^\tau(x) = \bar{\tau}^r \cdot y,$$

where the multiplication denotes the $\mathbf{Z}[\bar{\tau}]$ -module structure on $E_r^{*,*}$ from (2).

- (4) *If $x \in E_1^{n,w}$ and $y \in E_r^{n-1,w+r}$, then there is a differential*

$$d_r(x) = y$$

if and only if there is a differential

$$d_r^\tau(x) = \bar{\tau}^r \cdot y.$$

- (5) *Suppose that $x \in E_1^{n,w,s}$ detects an element $\theta \in \pi_{n,w} X$. Then θ is τ^s -divisible. Moreover, for every $m \geq 0$, the class $\bar{\tau}^m \cdot x$ detects $\tau^m \cdot \theta$.*

- (6) *The τ -Bockstein spectral sequence for X converges conditionally to $\pi_{*,*} X$ if and only if the spectral sequence underlying X converges conditionally to $\pi_* X^{-\infty}$. If this is the case, then the τ -Bockstein spectral sequence converges strongly if and only if $\text{RE}_\infty^{n,s}$ (the derived ∞ -term for the spectral sequence underlying X) vanishes for all n, s .*

Here, items (2) and (4) should be read inductively, in the following way.

- ◆ Item (1) provides a $\mathbf{Z}[\bar{\tau}]$ -module structure on the first page, so that the statement of $\bar{\tau}$ -linearity of the d_1^τ -differential is well defined. As the second page is the homology of the first page, this endows the second page with a natural $\mathbf{Z}[\bar{\tau}]$ -module structure, so that we can talk about $\bar{\tau}$ -linearity of d_2^τ -differentials, and so forth.
- ◆ Using item (1), the comparison of d_1 with d_1^τ makes sense. Once this is established, it follows that the isomorphism from item (1) induces an isomorphism

$$E_2^{n,w,s} \cong E_2^{n,w+s}$$

for all $s \geq 1$. Since the d_2^τ -differential only hits filtrations $s \geq 2$, we can use these isomorphisms to make sense of the comparison between d_2 and d_2^τ , and so forth.

Proof. Item (1) is true by definition of $\bar{\tau}$ in Notation 3.50.

Item (2) is clear after unwrapping definitions. Let us do this in the case $m = 1$, which is sufficient to prove the entire statement. If x is an element of tridegree (n, w, s) , this means that x is an element of

$$\pi_{n,w}(\mathrm{Gr}^s \mathrm{BF}_\tau X) = \pi_{n,w+s} X / \tau.$$

The meaning of a differential $d_r^\tau(x) = y$ is as follows: there is an element α in $\pi_{n,w} \mathrm{BF}_\tau^{s+r} X = \pi_{n,w+s+r} X$ such that

$$\tau^{r-1} \cdot \alpha = \partial_1^\infty(x) \quad \text{in } \pi_{n,w} \mathrm{BF}_\tau^{s+1} X = \pi_{n,w+s+1} X,$$

and such that the mod τ reduction of α is equal to y (up to boundaries of shorter differentials). The element $\bar{\tau} \cdot x$ is given by considering x as an element in

$$\pi_{n,w-1}(\mathrm{Gr}^{s+1} \mathrm{BF}_\tau X) \cong \pi_{n,w-1+s+1} X / \tau = \pi_{n,w} X / \tau,$$

and the differential on it is calculated in exactly the same way. This proves the ‘only if’ statement. To also deduce the ‘if’ statement, it suffices to observe that by induction on r , the shorter differentials entering tridegree $(n, w, s + r)$ are $\bar{\tau}$ -power multiples of differentials originating from filtration 0. As a result, both the differential $d_r^\tau(x) = y$ and $d_r^\tau(\bar{\tau}x) = \bar{\tau}y$ have the same boundaries as their indeterminacy, so the claim follows.

Item (3) will follow from item (2), combined with the following claim: for every r, n, w , and s , multiplication by $\bar{\tau}$ induces a surjection

$$\bar{\tau}: E_r^{n,w,s} \twoheadrightarrow E_r^{n,w-1,s+1}.$$

This claim, in turn, we prove by induction on r . For $r = 1$ it is clear from item (1). Assume that for some $r \geq 1$, multiplication by $\bar{\tau}$ induces a surjection $E_r^{n,w,s} \twoheadrightarrow E_r^{n,w-1,s+1}$. Write $(\ker d_r)^{n,w,s}$ for the kernel of the d_r -differential out of $E_r^{n,w,s}$, and write $(\mathrm{im} d_r)^{n,w,s}$ for the image of d_r into $E_r^{n,w,s}$. Then we have

$$E_{r+1}^{n,w,s} = \frac{(\ker d_r)^{n,w,s}}{(\mathrm{im} d_r)^{n,w,s}} \quad \text{and} \quad E_{r+1}^{n,w-1,s+1} = \frac{(\ker d_r)^{n,w-1,s+1}}{(\mathrm{im} d_r)^{n,w-1,s+1}}.$$

The $\bar{\tau}$ -divisibility of the differentials from (2) implies that multiplication by $\bar{\tau}$ restricts to a surjection $(\ker d_r)^{n,w,s} \rightarrow (\ker d_r)^{n,w-1,s+1}$, which implies the desired statement.

To prove item (4), it suffices to prove this for one fixed value of $w = w_0$ at a time. The exact couple computing the differentials going out of filtration w is also computed by the exact couple associated with the truncated version of X :

$$\dots \longrightarrow X^{w_0+2} \longrightarrow X^{w_0+1} \longrightarrow X^{w_0} \longleftarrow X^{w_0} \longleftarrow \dots,$$

where X^{w_0} is placed in filtration w_0 . Let us denote this filtered spectrum by Y . (In a picture, the spectral sequence underlying Y is obtained from the one associated to X by removing the elements in filtration strictly below w_0 .) Next, we observe that by levelwise evaluating $\mathrm{BF}_\tau X: \mathbf{Z}^{\mathrm{op}} \rightarrow \mathrm{FilSp}$ at degree w_0 , we obtain $\Sigma^{0,-w_0} Y$. As a result, the functor of levelwise evaluating at degree w_0 induces an isomorphism of exact couples

$$A^{n,w_0,s}(\mathrm{BF}_\tau X) \xrightarrow{\cong} A^{n,w_0+s}(Y) \quad E^{n,w_0,s}(\mathrm{BF}_\tau X) \xrightarrow{\cong} E^{n,w_0+s}(Y)$$

where $s \geq 0$. This identifies the d_r^τ -differentials with the d_r -differentials going out of in filtration w_0 .

For item (5), recall from Definition 2.36 that x detecting θ means that there is a lift $\alpha \in \pi_{n,w} \mathrm{BF}_\tau^s X = \pi_{n,w+s} X$ of x that maps to θ under $\pi_{n,w} \mathrm{BF}_\tau^s X \rightarrow \pi_{n,w} \mathrm{BF}_\tau^0 X$. This transition map is given by multiplication by τ^s , so this is saying that $\tau^s \cdot \alpha = \theta$. This proves the first clause. We prove the second clause for $m = 1$, from which the general case follows by iterating. Using the E_∞ -structure on $C\tau$ from Theorem 3.24, it follows that multiplication by τ is zero on $\pi_{*,*}(X/\tau)$, so that $\tau \cdot \theta$ is detected in filtration at least $s + 1$. Next, the element α , considered as an element of $\pi_{n,w-1} \mathrm{BF}_\tau^{s+1} X = \pi_{n,w+s} X$, is a lift of $\bar{\tau} \cdot x \in E_1^{n,w-1,s+1}$. Clearly, the element α in $\pi_{n,w-1} \mathrm{BF}_\tau^{s+1} X$ maps to $\tau^{s+1} \cdot \alpha = \tau \cdot \theta$ under $\tau^{s+1}: \pi_{n,w+s} X \rightarrow \pi_{n,w-1} X$, so indeed $\bar{\tau} \cdot x$ detects $\tau \cdot \theta$.

Finally, we address the claims regarding convergence. Conditional convergence of the τ -Bockstein spectral sequence is the claim that the limit of $\mathrm{BF}_\tau X$ vanishes. By Proposition 3.18, this limit is isomorphic to the constant spectrum on X^∞ , proving the claim about conditional convergence. Suppose now that both spectral sequences converge conditionally. By the trigraded analogue^[2] of Theorem 2.55, the τ -Bockstein spectral sequence converges strongly if and only if $\mathrm{RE}_\infty^{n,w,s}$ vanishes for all n, w, s . Using the isomorphism from (4), we find that (for $s \geq 0$)

$$\mathrm{RE}_\infty^{n,w,s} = \lim_r^1 Z_r^{n,w,s} \cong \lim_r^1 Z_r^{n,w+s} = \mathrm{RE}_\infty^{n,w+s}. \quad \blacksquare$$

^[2]As explained in [Hed20, Remark 1.2.4], one needs certain requirements on an abelian category \mathcal{A} to ensure that Boardman's arguments apply to spectral sequences valued in \mathcal{A} . Here, we are working with spectral sequences valued in bigraded abelian groups, where these conditions are certainly met, and Boardman's arguments go through without any change (merely tagging on an additional grading).

Remark 3.53. A different way of stating (part of) items (3) and (4) above is to say that the isomorphism of item (1) induces isomorphisms (for $s \geq 0$)

$$Z_r^{n,w,s} \cong Z_r^{n,w+s} \quad B_r^{n,w,s} \cong B_{\min(r,s)}^{n,w+s} \quad E_{r+1}^{n,w,s} \cong Z_r^{n,w+s} / B_{\min(r,s)}^{n,w+s}.$$

In particular, we have isomorphisms

$$Z_\infty^{n,w,s} \cong Z_\infty^{n,w+s} \quad B_\infty^{n,w,s} \cong B_s^{n,w+s} \quad E_\infty^{n,w,s} \cong Z_\infty^{n,w+s} / B_s^{n,w+s}.$$

The appearance of the s -boundaries instead of the ∞ -boundaries in the last expression is the reason that a d_r -differential in the spectral sequence underlying X leads to τ^r -power torsion in $\pi_{*,*} X$.

Remark 3.54 (Convergence of the τ -BSS). Notably, in part (6), we do not need to assume strong convergence of the spectral sequence underlying X to ensure strong convergence of the τ -BSS for X . This distinction is relevant in the case where X is neither left nor right concentrated, in which case additionally Boardman's whole-plane obstruction W from Remark 2.57 needs to vanish to ensure strong convergence of the spectral sequence underlying X . The reason this does not appear for the τ -BSS is explained by Warning 2.52. Indeed, the group W is the obstruction to a limit commuting with a colimit, but the abutment of the τ -BSS is $\pi_{n,w} X$, where we have not yet taken the colimit over w (which would result in $\pi_n X^{-\infty}$). Of course, to be able to study $\pi_* X^{-\infty}$ using $\pi_{*,*} X$, one would then need W to vanish.

Remark 3.55 (Second-page indexing). As per usual up till this point, we have used first-page indexing for the τ -Bockstein spectral sequence. Even in situations where one indexes filtered spectra using second-page indexing, there is something to be said for using first-page indexing for the τ -BSS: only in this indexing does filtration correspond to τ -power divisibility. In that case, a $d_r^{\bar{\tau}}$ -differential would correspond to a d_{r+1} -differential. (See Variant 4.39 for an example.) If desired however, second-page indexing for the τ -BSS can be achieved via

$$\tilde{E}_{r+1}^{n,w,s} := E_r^{n,w,s+w}.$$

In this indexing, $\bar{\tau}$ has tridegree $(0, -1, 0)$, and the non- $\bar{\tau}$ -divisible groups are located in tridegrees of the form (n, w, w) .

Warning 3.56. The $\bar{\tau}$ -divisions appearing in items (2) and (3) are not necessarily unique, due to the $\bar{\tau}$ -torsion caused by shorter differentials.

3.4.1 The truncated τ -Bockstein spectral sequence

There is a truncated version of the τ -Bockstein spectral sequence, which instead computes the bigraded homotopy groups of X/τ^k . One could, of course, apply the τ -BSS directly to X/τ^k , but it is more efficient to use the following modification of the τ -BSS.

Definition 3.57. Let X be a filtered spectrum and let $k \geq 1$. The k -truncated τ -Bockstein filtration on X is the bifiltered spectrum $\mathrm{tr}_k \mathrm{BF}_\tau X$ given in nonnegative filtrations by

$$\cdots \longrightarrow 0 \longrightarrow \Sigma^{0, -k+1} X/\tau \xrightarrow{\tau} \Sigma^{0, -k+2} X/\tau^2 \xrightarrow{\tau} \cdots \xrightarrow{\tau} X/\tau^k$$

and indexed to be constant from filtration 0 onwards. The resulting spectral sequence we call the k -truncated τ -Bockstein spectral sequence for X .

Remark 3.58. Upon evaluation at a fixed filtration, we obtain the filtered spectrum described by [Ant24, Construction 3.17].

Rather than having to re-prove the analogous version of Theorem 3.52 from the ground up, we can instead deduce this from Theorem 3.52 by means of the following map.

Construction 3.59. Let X be a filtered spectrum. Then for every $k \geq 1$, we have a morphism of bifiltered spectra $\mathrm{BF}_\tau X \rightarrow \mathrm{tr}_k \mathrm{BF}_\tau X$ of the form

$$\begin{array}{ccccccccccc} \cdots & \xrightarrow{\tau} & \Sigma^{0, -k} X & \xrightarrow{\tau} & \Sigma^{0, -k+1} X & \xrightarrow{\tau} & \Sigma^{0, -k+2} X & \xrightarrow{\tau} & \cdots & \xrightarrow{\tau} & X \\ & & \downarrow & & \downarrow & & \downarrow & & & & \downarrow \\ \cdots & \longrightarrow & 0 & \longrightarrow & \Sigma^{0, -k+1} X/\tau & \xrightarrow{\tau} & \Sigma^{0, -k+2} X/\tau^2 & \xrightarrow{\tau} & \cdots & \xrightarrow{\tau} & X/\tau^k \end{array}$$

constructed by letting each nontrivial square be a pushout. Formally, it is left Kan extended from the subdiagram

$$\begin{array}{ccccccc} \cdots & \xrightarrow{\tau} & \Sigma^{0, -k-1} X & \xrightarrow{\tau} & \Sigma^{0, -k} X & \xrightarrow{\tau} & \cdots \xrightarrow{\tau} X. \\ & & \downarrow & & \downarrow & & \\ \cdots & \longrightarrow & 0 & \longrightarrow & 0 & & \end{array}$$

Likewise, for $m \geq k \geq 1$, we have a morphism $\mathrm{tr}_m \mathrm{BF}_\tau X \rightarrow \mathrm{tr}_k \mathrm{BF}_\tau X$, and these fit into a tower

$$\mathrm{BF}_\tau X \longrightarrow \cdots \longrightarrow \mathrm{tr}_k \mathrm{BF}_\tau X \longrightarrow \mathrm{tr}_{k-1} \mathrm{BF}_\tau X \longrightarrow \cdots \longrightarrow \mathrm{tr}_1 \mathrm{BF}_\tau X.$$

On colimits (equivalently, on filtration 0), this tower is the τ -adic tower of X :

$$X \longrightarrow \cdots \longrightarrow X/\tau^k \longrightarrow X/\tau^{k-1} \longrightarrow \cdots \longrightarrow X/\tau.$$

The only subtlety with the truncated version is the distinction between the two kinds of τ -multiples, as discussed in Notation 3.35.

Theorem 3.60. Let X be a filtered spectrum and let $k \geq 1$. Let $\{E_r^{*,*}, d_r\}$ denote the spectral sequence underlying X , and let $\{E_r^{*,*}, d_r^\tau\}$ denote the k -truncated τ -Bockstein spectral sequence of X .

- (1) The map on spectral sequences induced by $\mathrm{BF}_\tau X \rightarrow \mathrm{tr}_k \mathrm{BF}_\tau X$ from Construction 3.59 is, on first pages, given by the quotient

$$\mathbf{Z}[\bar{\tau}] \otimes \pi_{*,*}(X/\tau) \longrightarrow \mathbf{Z}[\bar{\tau}]/\bar{\tau}^k \otimes \pi_{*,*}(X/\tau).$$

For $m \leq k$, the map on spectral sequences induced by $\mathrm{tr}_k \mathrm{BF}_\tau X \rightarrow \mathrm{tr}_m \mathrm{BF}_\tau X$ from Construction 3.59 is, on first pages, given by the quotient

$$\mathbf{Z}[\bar{\tau}]/\tau^k \otimes \pi_{*,*}(X/\tau) \longrightarrow \mathbf{Z}[\bar{\tau}]/\bar{\tau}^m \otimes \pi_{*,*}(X/\tau).$$

- (2) The differentials are $\bar{\tau}$ -linear: for $x \in E_1^{n,w,s}$ and $y \in E_1^{n-1,w,s+r}$, if there is a differential

$$d_r^\tau(x) = y,$$

then for all $m \geq 0$, there is a differential

$$d_r^\tau(\bar{\tau}^m x) = \bar{\tau}^m y.$$

In particular, the $\mathbf{Z}[\bar{\tau}]$ -module structure on the first page induces a $\mathbf{Z}[\bar{\tau}]$ -module structure on later pages.

- (3) The target of a d_r^τ -differential is a $\bar{\tau}^r$ -multiple: for every $x \in E_r^{n,w,s}$, there is an element $y \in E_r^{n-1,w+r,s}$ such that

$$d_r^\tau(x) = \bar{\tau}^r \cdot y,$$

where the multiplication denotes the $\mathbf{Z}[\bar{\tau}]$ -module structure on $E_r^{*,*,*}$ from (2).

- (4) Let $r \leq k - 1$. If $x \in E_1^{n,w}$ and $y \in E_r^{n-1,w+r}$, then there is a differential

$$d_r(x) = y$$

if and only if there is a differential

$$d_r^\tau(x) = \bar{\tau}^r \cdot y.$$

- (5) Suppose that $x \in E_1^{n,w,s}$ detects an element $\theta \in \pi_{n,w} X/\tau^k$. Then θ is in the image of the map

$$\tau^s: \Sigma^{0,-s} X/\tau^{k-s} \longrightarrow X/\tau^k,$$

i.e., it is of the form $\tau^s(\alpha)$ for some $\alpha \in \pi_{n,w+s} X/\tau^{k-s}$.

Moreover, for every m such that $s + m \leq k - 1$, the class $\bar{\tau}^m \cdot x$ detects $\tau^m \cdot \theta$. For every m such that $s + m \geq k$, we instead have $\tau^m \cdot \theta = 0$.

- (6) The k -truncated τ -Bockstein spectral sequence converges strongly to $\pi_{*,*} X/\tau^k$.

Take particular note that in item (2), the truncated differentials are no longer $\bar{\tau}$ -divisible in general, and that item (4) only applies to differentials of length $r \leq k - 1$ (longer d_r^τ -differentials vanish for degree reasons).

Proof. The definition of the map $\mathrm{BF}_\tau X \rightarrow \mathrm{tr}_k \mathrm{BF}_\tau X$ from Construction 3.59 as a left Kan extension shows that in filtrations $0 \leq s \leq k-1$, the map induces an isomorphism on associated graded. This proves the first part of item (1), and the second is analogous.

Items (2) to (4) follow immediately from items (2) to (4) of Theorem 3.52, using the map from item (1). For item (6), it is enough to note that $\mathrm{tr}_k \mathrm{BF}_\tau X$ vanishes in filtrations $s \geq k$, so that (the trigraded analogue of) Proposition 2.56 implies the strong convergence.

It remains to prove item (5). The statement that x detects θ means that there is a lift $\alpha \in \pi_{n, w+s} X / \tau^{k-s}$ of x that maps to θ under the transition map

$$\mathrm{tr}_k \mathrm{BF}_\tau^s X = \Sigma^{0, -s} X / \tau^{k-s} \longrightarrow X / \tau^k = \mathrm{tr}_k \mathrm{BF}_\tau^0 X,$$

in other words, that $\tau^s(\alpha) = \theta$. If $s + m \leq k - 1$, then

$$\pi_{n, w-m}(\mathrm{Gr}^{s+m} \mathrm{tr}_k \mathrm{BF}_\tau X) = \pi_{n, w} X / \tau,$$

and $\bar{\tau}^m \cdot x$ is given by the element x considered as an element of this group via this identification. Write β for the mod τ^{k-s-m} reduction of α . Then β is a lift of $\bar{\tau}^m \cdot x$ to $\pi_{n, w-m}(\mathrm{tr}_k \mathrm{BF}_\tau^{s+m} X) = \pi_{n, w+s} X / \tau^{k-s-m}$. Evidently, the image of this element in $\mathrm{tr}_k \mathrm{BF}_\tau^0 X$ is given by

$$\tau^{s+m}(\beta) = \tau^s(\tau^m(\beta)) = \tau^s(\tau^m \cdot \alpha) = \tau^m \cdot \tau^s(\alpha) = \tau^m \cdot \theta.$$

Lastly, if instead $s + m \geq k$, then this means $m \geq k - s$. Since $\tau^m \cdot \tau^s(\alpha) = \tau^s(\tau^m \cdot \alpha)$ and α lives in $\pi_{*,*} X / \tau^{k-s}$, it follows that $\tau^m \cdot \alpha = 0$ (using the ring structure on $C\tau^{k-s}$), proving the final claim. \blacksquare

Warning 3.61. Unlike in the integral case, if an element in the truncated τ -BSS is detected in filtration s , this does not imply that it is a τ^s -multiple in $\pi_{*,*} X / \tau^k$. Instead, in general this only implies that it is of the form $\tau^s(\theta)$ for some θ in $\pi_{*,*} X / \tau^{k-s}$. (See Notation 3.35 for the distinction between these two.)

3.5 The Omnibus Theorem

We can summarise Section 3.2 as follows: if X is a filtered spectrum, then its underlying spectral sequence is of the form

$$E_1^{n,s} = \pi_{n,s}(C\tau \otimes X) \implies \pi_n(X^{\tau=1}).$$

In Appendix A, we argued that the bigraded homotopy groups $\pi_{*,*} X$ as a $\mathbf{Z}[\tau]$ -module should capture the differentials in this spectral sequence. The *Omnibus Theorem* makes this precise, by describing the structure of $\pi_{*,*} X$ in terms of the underlying spectral sequence of X . We prove it here in the context of filtered spectra, which should be regarded as the most general (stable) setting for it: later in Section 3.6, we will show how it extends to any *monoidal deformation*. In Chapter 4, we will see how this recovers and extends the synthetic Omnibus Theorem of Burklund–Hahn–Senger, and compare our proof to theirs; see Section 4.5 in particular.

Theorem 3.62 (Omnibus). *Let X be a complete filtered spectrum, and assume that in the spectral sequence underlying X , we have $\mathrm{RE}_\infty^{*,*} = 0$ (for instance, this happens if the spectral sequence converges strongly). Let $x \in E_1^{n,s} = \pi_{n,s}(X/\tau)$ be a nonzero class. Then the following are equivalent.*

- (1a) *The element x is a permanent cycle.*
- (1b) *The element $x \in \pi_{n,s}(X/\tau)$ lifts to an element of $\pi_{n,s} X$.*

For any such lift α to $\pi_{n,s} X$, the following are true.

- (2a) *If x survives to page r , then $\tau^{r-1} \cdot \alpha$ is nonzero.*
- (2b) *If x survives to page ∞ , then α maps to a nonzero element in $\pi_n X^{\tau=1} = \pi_n X^{-\infty}$, and this element is detected by x .*

Moreover, if x lifts to X , then there exists a lift α with either one of the following additional properties.

- (3a) *If x is the target of a d_r -differential, then $\tau^r \cdot \alpha = 0$.*
- (3b) *If $\theta \in \pi_n X^{\tau=1}$ is detected by x , then α is sent to θ under $\pi_{n,s} X \rightarrow \pi_n X^{\tau=1}$.*

Finally, we have the following generation statement.

- (4) *Let $\{\alpha_i\}$ be a collection of elements of $\pi_{n,*} X$ such that their mod τ reductions generate the permanent cycles in stem n . Then the τ -adic completion of the $\mathbf{Z}[\tau]$ -submodule of $\pi_{n,*} X$ generated by the $\{\alpha_i\}$ is equal to $\pi_{n,*} X$.*

Proof. By Theorem 3.52 (6), the hypotheses on X are equivalent to strong convergence of the τ -Bockstein spectral sequence for X .

We begin with (1). The map $\mathrm{BF}_\tau X \rightarrow \mathrm{tr}_1 \mathrm{BF}_\tau X$ from Construction 3.59 is, on underlying objects, equal to the reduction map $X \rightarrow X/\tau$. Meanwhile the induced map on the first page of the resulting spectral sequences is given by quotienting by $\bar{\tau}$:

$$\mathbf{Z}[\bar{\tau}] \otimes \pi_{*,*}(X/\tau) \twoheadrightarrow \pi_{*,*}(X/\tau),$$

so in particular it is an isomorphism in filtration 0. By strong convergence, every permanent cycle in the τ -BSS lifts to $\pi_{*,*} X$. Accordingly, it follows that $x \in \pi_{n,s}(X/\tau)$ lifts to $\pi_{n,s} X$ if and only if the element x considered in tridegree $(n, s, 0)$ of the τ -BSS is a permanent cycle. The equivalence of (1a) and (1b) therefore follows from the identification of the differentials in the τ -BSS from Theorem 3.52 (4).

Next, we check (2a). Suppose that $\alpha \in \pi_{n,s} X$ is a lift of x . Then since $\mathrm{BF}_\tau X$ is constant from filtration 0 and onwards, this means that the class x detects α . By Theorem 3.52 (5), the element $\bar{\tau}^{r-1} \cdot x$ detects $\tau^{r-1} \cdot \alpha$. To show that $\tau^{r-1} \cdot \alpha$ is nonzero, by (the trigraded analogue of) Remark 2.46 we need to show that $\bar{\tau}^{r-1} \cdot x$ survives to page ∞ . Again using Theorem 3.52 (4), we see that $\bar{\tau}^{r-1} \cdot x$ survives to page r if and only if x in the spectral

sequence underlying X survives to page r . If this happens, then for degree reasons $\bar{\tau}^{r-1} \cdot x$ also survives to page ∞ .

Properties (2b) and (3b) are a restatement of the definition of detection from Definition 2.36, with (2b) specifically following from Remark 2.46.

We prove (3a) by induction on r . Suppose that x is the target of a d_r -differential, say $d_r(y) = x$ for $y \in \pi_{n+1, s-r}(X/\tau)$. This implies that $d_r^\tau(y) = \bar{\tau}^r \cdot x$. Unrolling the definition of the d_r^τ -differential, this means the following. There exists an element $\alpha \in \pi_{n, s} X$ such that

$$\tau^{r-1} \cdot \alpha = \partial_1^\infty(y) \quad \text{in } \pi_{n-1, s-r+1} X,$$

and such that the mod τ reduction of α agrees with x up to the images of differentials of length shorter than r .

$$\begin{array}{ccccc} \Sigma^{0, -r} X & \xrightarrow{\tau^{r-1}} & \Sigma^{0, -1} X & \xrightarrow{\tau} & X \\ \downarrow & & \swarrow \partial_1^\infty & & \downarrow \\ \Sigma^{0, -r} X/\tau & & & & X/\tau \end{array}$$

By induction, we may assume that the images of shorter differentials have τ^{r-1} -torsion lifts, so that without loss of generality, we may assume α is a lift of x . By exactness, we have $\tau \cdot \partial_1^\infty(y) = 0$. But then $\tau^r \cdot \alpha = \tau \cdot \partial_1^\infty(y) = 0$, so that α is a lift satisfying (3a).

For the final claim, let M denote the $\mathbf{Z}[\tau]$ -submodule of $\pi_{n,*} X$ generated by the α_i . By strong convergence of the τ -BSS, the $\mathbf{Z}[\tau]$ -module $\pi_{n,*} X$ is τ -complete; see Warning 3.32. It follows that M_τ^\wedge is naturally a submodule of $\pi_{n,*} X$. To show that this inclusion is an equality, it suffices to show that it becomes surjective after quotienting by τ , as both modules are τ -complete. After quotienting by τ , the inclusion becomes

$$M/\tau \longrightarrow (\pi_{n,*} X)/\tau = F_\tau^0 \pi_{n,*} X,$$

where the last identification is Theorem 3.52 (5). By the identifications of the differentials in the τ -BSS of Theorem 3.52 (4), the assumption on the α_i translates to this map being a surjection. This finishes the proof. \blacksquare

Warning 3.63. Suppose that $x \in E_1^{n,s}$ is the target of a d_r -differential, and suppose that $\alpha \in \pi_{n,s} X$ is a lift of x . Then item (3a) does *not* necessarily imply that α is τ^r -torsion (in fact, in general α need not even be τ -power torsion): the theorem only guarantees that there exists *some* lift of x that is τ^r -torsion.

We can draw a number of simpler conclusions from this result.

Corollary 3.64. *Let X be a filtered spectrum satisfying the conditions of Theorem 3.62, and let n be an integer. Then the $\mathbf{Z}[\tau]$ -module $\pi_{n,*} X$ is τ -power torsion free if and only if the n -stem in the spectral sequence underlying X does not receive any nonzero differentials.*

Proof. The lack of incoming differentials implies that every permanent cycle survives to page ∞ , so one direction follows from item (2a). Conversely, if the n -stem does receive a nonzero differential, then by item (3a) there exists a τ -power torsion element in $\pi_{n,*} X$. ■

Corollary 3.65. *Let X be a filtered spectrum satisfying the conditions of Theorem 3.62, and let n, s be integers. If $\pi_{n,s+d} X/\tau$ vanishes for all $d \geq 0$, then $\pi_{n,s} X$ vanishes also.*

Proof. This follows directly from item (4). ■

Remark 3.66. By inspecting the proof of Theorem 3.62, we see that the convergence conditions on X are only needed for items (1) and (4). Without these assumptions, item (2) still holds for any lift, but such a lift is no longer guaranteed to exist. In item (3) meanwhile, the assumptions on x in both (3a) and (3b) imply that a lift exists (see Definition 2.36).

3.5.1 The truncated Omnibus Theorem

There is also a version of the Omnibus Theorem that describes the structure of $\pi_{*,*}(X/\tau^k)$ in terms of the spectral sequence underlying X . In this case, we no longer need any convergence conditions, but the generation statement is more involved to state. Let us begin therefore with the other parts of the Omnibus Theorem.

Theorem 3.67 (Truncated Omnibus, part 1). *Let X be a filtered spectrum, and let $k \geq 1$. Let $x \in E_1^{n,s} = \pi_{n,s}(X/\tau)$ be a class. Then the following are equivalent.*

- (1a) *The differentials $d_1(x), \dots, d_{k-1}(x)$ vanish.*
- (1b) *The element $x \in \pi_{n,s}(X/\tau)$ lifts to an element of $\pi_{n,s}(X/\tau^k)$.*

For any such lift α to $\pi_{n,s}(X/\tau^k)$, the following are true.

- (2a) *If x survives to page r for $r \leq k$, then $\tau^{r-1} \cdot \alpha$ is nonzero.*
- (2b) *The image of α under $\partial_k^{k+1}: \pi_{n,s}(X/\tau^k) \rightarrow \pi_{n-1,s+r}(X/\tau)$ is a representative for $d_k(x)$.*

Moreover, if x lifts to X/τ^k , then there exists a lift α with the following additional property.

- (3) *If x is the target of a d_r -differential for $r < k$, then $\tau^r \cdot \alpha = 0$.*

Proof. The k -truncated τ -Bockstein spectral sequence for X converges strongly to $\pi_{*,*} X/\tau^k$ by Theorem 3.60 (6).

The morphism $\mathrm{tr}_k \mathrm{BF}_\tau X \rightarrow \mathrm{tr}_1 \mathrm{BF}_\tau X$ from Construction 3.59 is, on underlying objects, the reduction map $X/\tau^k \rightarrow X/\tau$. On first pages of the underlying spectral sequences on the other hand, the map takes the form of quotienting $\bar{\tau}$:

$$\mathbf{Z}[\bar{\tau}]/\bar{\tau}^k \otimes \pi_{*,*} X/\tau \longrightarrow \pi_{*,*} X/\tau.$$

In the same way as in the proof of Theorem 3.62, the equivalence between (1a) and (1b) follows by strong convergence and from the identification of the differentials in the truncated τ -BSS arising from Theorem 3.60 (1).

Next, we prove item (2a). Let $\alpha \in \pi_{n,s} X/\tau^k$ be a lift of x . By Theorem 3.60 (5), the element $\tau^{r-1} \cdot \alpha$ is detected by $\bar{\tau}^{r-1} \cdot x$. The assumption that x survives to page r implies that $\bar{\tau}^{r-1} \cdot x$ is a permanent cycle that survives to page r , and for degree reasons it then also survives to page ∞ . Using Remark 2.46, it follows that $\tau^{r-1} \cdot \alpha$ is nonzero.

Next, we prove item (2b). Let $\alpha \in \pi_{n,s} X/\tau^k$ be a lift of x . As a consequence of Theorem 3.60 (1), it is enough to show that $\partial_k^{k+1}(\alpha)$ is a representative for $d_k^r(x)$ in the non-truncated τ -BSS, where we regard x in tridegree $(n, s, 0)$. Recall how the d_k -differential on x in the non-truncated τ -BSS is computed: we apply the boundary map $\partial_1^\infty(x)$, choose a τ^{k-1} -division of this element, and reduce this mod τ .

$$\begin{array}{ccccc} \Sigma^{0,-k} X & \xrightarrow{\tau^{k-1}} & \Sigma^{0,-1} X & \xrightarrow{\tau} & X \\ \downarrow & & \swarrow \partial_1^\infty & & \downarrow \\ \Sigma^{0,-k} X/\tau & & & & X/\tau \end{array}$$

We have $\tau^{k-1} \cdot \partial_k^\infty(\alpha) = \partial_1^\infty(x)$. In other words, $\partial_k^\infty(\alpha)$ is a valid choice of τ^{k-1} -division of $\partial_1^\infty(x)$, so that its projection to $\pi_{n,s+k} X/\tau$ is a representative for $d_k^r(x)$. But the mod τ reduction of ∂_k^∞ is ∂_k^{k+1} , proving item (2b).

Next, we prove item (3). Suppose that x is the target of a d_r -differential for $r < k$. We consider the element $\bar{\tau}^{r-1} \cdot x$ in tridegree $(n, s-r+1, r-1)$ of the k -truncated τ -Bockstein spectral sequence. Then the d_r -differential hitting x translates to a d_r^r -differential

$$d_r^r(y) = \bar{\tau}^r x.$$

Unrolling what this means, we learn the following. There exists an element $\beta \in \pi_{n,s}(X/\tau^{k-r})$ such that $\tau^{r-1}(\beta) = \partial_1^k(y)$, and such that β reduces to x modulo τ up to the images of shorter differentials.

$$\begin{array}{ccccc} \Sigma^{0,-r} X/\tau^{k-r} & \xrightarrow{\tau^{r-1}} & \Sigma^{0,-1} X/\tau^{k-1} & \xrightarrow{\tau} & X/\tau^k \\ \downarrow & & \swarrow \partial_1^k & & \downarrow \\ \Sigma^{0,-r} X/\tau & & & & X/\tau \end{array}$$

Like in the proof of Theorem 3.62, by induction we may assume without loss of generality that β reduces to x . By exactness, the element $\partial_1^k(y)$ satisfies $\tau(\partial_1^k(y)) = 0$, so any choice of β satisfies $\tau^r(\beta) = 0$. As a result, it suffices to show that there is a choice of $\beta \in \pi_{n,s} X/\tau^{k-r}$ that lifts to an element in $\pi_{n,s} X/\tau^k$. Indeed, if such a lift α exists, then by Remark 3.36, we have $\tau^r \cdot \alpha = \tau^r(\beta) = 0$, which would mean that α is the lift proving item (3).

To produce such an α , we first show that $y \in \pi_{n+1,s-r} X/\tau$ lifts to X/τ^r . This follows from (1) because y is a $d_{\leq r-1}$ -cycle. Choose a lift \tilde{y} . It then follows that $\partial_r^k(\tilde{y})$ is a valid

choice for β as above. To show that this β lifts to X/τ^k , we need to show that $\partial_{k-r}^k(\beta) = 0$. Note that $\partial_{k-r}^k \circ \partial_r^k$ can be written as (omitting shifts for readability)

$$X/\tau^{k-1} \xrightarrow{\partial_k^\infty} X \longrightarrow X/\tau^{k-r} \xrightarrow{\partial_{k-r}^\infty} X \longrightarrow X/\tau^r.$$

The middle two maps are part of a cofibre sequence, so in particular their composition is zero. This means that indeed $\partial_{k-r}^k(\beta) = 0$, showing that β lifts to the desired α , thus proving item (3). ■

Construction 3.68. Let X be a filtered spectrum, and let n, s be integers.

- (1) By Theorem 3.62 (1), the reduction map $\pi_{n,s} X \rightarrow \pi_{n,s} X/\tau = E_1^{n,s}$ has image given by the subgroup $Z_\infty^{n,s}$ of permanent cycles. Postcomposing with the quotients, we in particular obtain, for every $1 \leq r \leq \infty$, a map

$$\pi_{n,s} X \longrightarrow E_r^{n,s}$$

which is surjective in the case $r = \infty$.

- (2) Let $r \geq 1$. By Theorem 3.67 (1), the reduction map $\pi_{n,s} X/\tau^r \rightarrow \pi_{n,s} X/\tau = E_1^{n,s}$ has image given by the subgroup $Z_{r-1}^{n,s}$ of $(r-1)$ -cycles. Postcomposing with the quotient, we in particular obtain a surjective map

$$\pi_{n,s} X/\tau^r \twoheadrightarrow E_r^{n,s}.$$

These maps are compatible with each other in the obvious way.

Using these maps, we can now reformulate Proposition 3.37.

Corollary 3.69. *Let X be a filtered spectrum, let $r \geq 1$, and let n and s be integers. Then we have a commutative diagram*

$$\begin{array}{ccc} \pi_{n,s} X/\tau^r & \xrightarrow{\partial_r^\infty} & \pi_{n-1,s+r} X \\ \downarrow & & \downarrow \\ E_r^{n,s} & \xrightarrow{d_r} & E_r^{n-1,s+r}. \end{array}$$

More generally, for $R \geq r$, we have a commutative diagram

$$\begin{array}{ccc} \pi_{n,s} X/\tau^r & \xrightarrow{\partial_r^R} & \pi_{n-1,s+r} X/\tau^{R-r} \\ \downarrow & & \downarrow \\ E_r^{n,s} & \xrightarrow{d_r} & Z_m^{n-1,s+r}/B_{r-1}^{n-1,s+r} \end{array}$$

where m is the minimum of $r-1$ and $R-r-1$.

Proof. This follows directly from Proposition 3.37. ■

Next, we turn to the question of finding generators for $\pi_{*,*} X/\tau^k$. In the non-truncated version (Theorem 3.62 (4)), we could start with lifts of permanent cycles, and take the (τ -complete) $\mathbf{Z}[\tau]$ -module that they generate to reconstruct all of $\pi_{*,*} X$. In the truncated case, taking τ -multiples is a more subtle notion. In order to generate $\pi_{*,*} X/\tau^k$, we need to take the τ -multiples of elements from all lower truncations X/τ^i for $i \leq k$ into account. Unfortunately, stating this precisely makes the indexing get a little out of hand.

For applications, we also need a relative version describing the kernel of $X/\tau^k \rightarrow X/\tau^m$. To compensate for the more intricate phrasing of this result, we give a simplified, more coarse description in Corollary 3.72 below.

Theorem 3.70 (Truncated Omnibus, part 2). *Let X be a filtered spectrum, let $n, s \in \mathbf{Z}$, and let $r \geq 1$.*

(1) *Let $k \geq 1$ be fixed. Suppose that for every $1 \leq i \leq k$, we have a collection of elements*

$$\{\beta_j^i\}_j \subseteq \pi_{n, s+k-i} X/\tau^i$$

whose mod τ reductions generate the abelian group

$$\mathbf{Z}_{i-1}^{n, s+k-i} / \mathbf{B}_{k-i}^{n, s+k-i}. \quad (3.71)$$

(Note that by Theorem 3.67 (1), such a collection exists for every i .) Write

$$\alpha_j^i := \tau^{k-i}(\beta_j^i) \in \pi_{n, s} X/\tau^k.$$

Then $\{\alpha_j^i\}_{i,j}$ is a set of generators for the abelian group $\pi_{n, s} X/\tau^k$.

(2) *Let $1 \leq m \leq k$ be fixed. Suppose that for every $1 \leq i \leq k - m$, we have a collection of elements*

$$\{\beta_j^i\}_j \subseteq \pi_{n, s+k-i} X/\tau^i$$

whose mod τ reductions generate the abelian group (3.71). Write $\alpha_j^i := \tau^{k-i}(\beta_j^i)$. Then $\{\alpha_j^i\}_{i,j}$ is a set of generators for the abelian group

$$\ker(\pi_{n, s} X/\tau^k \longrightarrow \pi_{n, s} X/\tau^m).$$

Proof. The k -truncated Bockstein spectral sequence converges strongly to $\pi_{*,*} X/\tau^k$ by Theorem 3.60 (6). The first result therefore follows from Theorem 3.60 (4); see also Remark 3.53. The second follows analogously by considering the natural map from the k -truncated τ -BSS to the m -truncated τ -BSS for X induced by the map from Construction 3.59. ■

Sometimes, the following simplified criterion is sufficient.

Corollary 3.72. *Let X be a filtered spectrum, let $n, s \in \mathbf{Z}$, and let $k \geq m \geq 1$.*

- (1) *If $\pi_{n, s+d} X/\tau$ vanishes for $0 \leq d \leq k-1$, then $\pi_{n, s} X/\tau^k$ vanishes also.*
- (2) *If $\pi_{n, s+d} X/\tau$ vanishes for $m \leq d \leq k-1$, then the reduction map $\pi_{n, s} X/\tau^k \rightarrow \pi_{n, s} X/\tau^m$ is injective.*

Proof. In the notation of Theorem 3.70, we have $E_1^{n, s} = \pi_{n, s} X/\tau$, and $Z_r^{n, s}$ is a subgroup of this. It follows that the relevant groups in (3.71) vanish, so the claims follow. ■

3.6 Deformations

So far, we have seen that the ∞ -category FilSp is the natural home for (stable) spectral sequences. For specific purposes however, this category might be a bit unwieldy, and it might be helpful to find a modification of FilSp that is more suited to the problem at hand. The main example of such a modification in these notes is that of *synthetic spectra*. However, much of the setup of synthetic spectra holds much more generally, and leads one to a broad theory of ‘modifications’ of FilSp . These have become known as *deformations*. Readers only interested in synthetic spectra may move on to the next chapter, referring back to this section as needed.

This section is concerned with the general properties of deformations. Of particular interest is the case where this deformation structure arises from a (symmetric) monoidal left adjoint out of FilSp ; we call these (*symmetric*) *monoidal deformations*, which are the subject of Section 3.6.1. For monoidal deformations, we can prove much more: we prove all that we need in order to import all results about filtered spectra into a monoidal deformation; see Theorem 3.88. In particular, the Omnibus Theorem holds in any monoidal deformation, where the underlying spectral sequence is replaced by what we call the *signature spectral sequence*. We will make particular use of this in the case of synthetic spectra in the next chapter.

Later, in Chapter 5, we will continue our study of deformations and discuss *cellularity*, *filtered models*, and *evenness*. For now, our goals are more modest, and our main aim is to show how to deduce results in a deformation from results in filtered spectra.

Much of the material in this section is based on the treatment of deformations given by Barkan [Bar23, Section 2] and Burklund–Hahn–Senger [BHS22, Appendices A–C]. While these sources also discuss constructing new deformations out of old ones, we will focus on studying phenomena within a fixed deformation.

Definition 3.73 ([Bar23], Definition 2.2). A **(1-parameter, stable) deformation** is a left module over FilSp in $\text{Pr}_{\text{st}}^{\mathbf{L}}$.

A module over FilSp in $\text{Pr}_{\text{st}}^{\mathbf{L}}$ is also called an *FilSp-linear ∞ -category*. We refer to [NPR24, Section 3] for further background on \mathcal{D} -linear ∞ -categories where \mathcal{D} is a presentably symmetric monoidal ∞ -category.

A deformation \mathcal{C} is in particular left tensored over FilSp . As a result, we will use the same notation of bigraded shifts $\Sigma^{n,s}$ on \mathcal{C} to mean tensoring with the filtered spectrum $\mathbf{S}^{n,s}$. Likewise, if $X \in \mathcal{C}$, we will generally denote by $\tau: \Sigma^{0,-1}X \rightarrow X$ the map given by tensoring the map $\tau: \mathbf{S}^{0,-1} \rightarrow \mathbf{S}$ of filtered spectra with X . Moreover, if A is a filtered ring spectrum, then we can speak of modules over A in \mathcal{C} . Two cases of A deserve a special name.

Notation 3.74. Let \mathcal{C} be a deformation. The **generic fibre** of \mathcal{C} is defined by

$$\mathcal{C}[\tau^{-1}] := \text{Mod}_{\mathbf{S}[\tau^{-1}]}(\mathcal{C}),$$

and we refer to the functor

$$\mathcal{C} \longrightarrow \mathcal{C}[\tau^{-1}], \quad X \longmapsto X[\tau^{-1}] := \mathbf{S}[\tau^{-1}] \otimes X$$

as the **τ -inversion functor**. Further, the **special fibre** of \mathcal{C} is defined by

$$\text{Mod}_{C\tau}(\mathcal{C}),$$

and for $X \in \mathcal{C}$, we write

$$X/\tau := C\tau \otimes X.$$

Finally, we write \mathcal{C}_τ^\wedge for the full subcategory of \mathcal{C} on the $C\tau$ -local objects, which we call **τ -complete**. This results in a localisation

$$\mathcal{C} \xrightarrow[\leftarrow]{(-)_\tau^\wedge} \mathcal{C}_\tau^\wedge.$$

Remark 3.75. Using that the tensoring over FilSp preserves colimits in each variable, we find that

$$X[\tau^{-1}] = \text{colim}(X \xrightarrow{\tau} \Sigma^{0,1}X \xrightarrow{\tau} \dots)$$

and

$$X/\tau = \text{cofib}(\tau: \Sigma^{0,-1}X \longrightarrow X).$$

Moreover, since $\mathbf{S}[\tau^{-1}]$ is an idempotent in FilSp , it follows that $\mathcal{C} \rightarrow \mathcal{C}[\tau^{-1}]$ is a smashing localisation, and the forgetful functor $\mathcal{C}[\tau^{-1}] \rightarrow \mathcal{C}$ is fully faithful with essential image consisting of those X on which τ is an isomorphism. In most examples, the forgetful functor $\text{Mod}_{C\tau}(\mathcal{C}) \rightarrow \mathcal{C}$ is not fully faithful.

In a moment, we will see a way to obtain examples of deformations. For now, let us mention the following.

Example 3.76. The universal example of a deformation is that of FilSp itself. In this case, Proposition 3.18 identifies the generic fibre with Sp , and Theorem 3.25 identifies the special fibre with grSp . ▲

Example 3.77 ([BHS22], Appendix C.1). Let R be a filtered \mathbf{E}_1 -ring. Then the ∞ -category $\mathrm{Mod}_R(\mathrm{FilSp})$ is naturally a deformation. Its generic and special fibres are equivalent to, respectively,

$$\mathrm{Mod}_{R^{\tau=1}}(\mathrm{Sp}) \quad \text{and} \quad \mathrm{Mod}_{R/\tau}(\mathrm{grSp}).$$

If R is a filtered \mathbf{E}_∞ -ring, then these equivalences are naturally symmetric monoidal. In fact, in this case the deformation $\mathrm{Mod}_R(\mathrm{FilSp})$ is a *symmetric monoidal deformation*, a concept to be introduced in Section 3.6.1 below. \blacktriangle

Remark 3.78. The use of the terms *deformation*, *special fibre* and *generic fibre* is inspired by algebraic geometry. We think of τ as the deformation parameter, and the ‘deformation’ is one of the special fibre to the generic fibre. Geometrically, FilSp plays the role of $\mathbf{A}^1/\mathbf{G}_m$. One can in fact make this comparison precise; see [Mou21].

The analogous statement to Proposition 3.31 holds in any deformation. In this sense, the parameter τ governs the large-scale structure of a deformation.

Proposition 3.79. *Let \mathcal{C} be a deformation. For $X \in \mathcal{C}$, there is a natural pullback square*

$$\begin{array}{ccc} X & \longrightarrow & X_\tau^\wedge \\ \downarrow & \lrcorner & \downarrow \\ X[\tau^{-1}] & \longrightarrow & (X_\tau^\wedge)[\tau^{-1}]. \end{array}$$

In particular, a map $X \rightarrow Y$ in \mathcal{C} is an isomorphism if and only if the maps

$$X[\tau^{-1}] \rightarrow Y[\tau^{-1}] \quad \text{and} \quad C\tau \otimes X \rightarrow C\tau \otimes Y$$

are both an isomorphism.

Proof. This follows by tensoring the pullback square from Proposition 3.31 for the unit in FilSp with the object $X \in \mathcal{C}$. As tensoring preserves colimits in each variable, and we are in the stable setting, the resulting square is again a pullback square. \blacksquare

Remark 3.80. Because the inclusion $\mathcal{C}[\tau^{-1}] \subseteq \mathcal{C}$ admits both a left and a right adjoint, it follows from [HA, Proposition A.8.20] that this equips \mathcal{C} with the structure of a stable recollement

$$\begin{array}{ccccc} & \xleftarrow{\tau^{-1}} & & \xleftarrow{(-)_\tau^\wedge} & \\ & \curvearrowright & & \curvearrowright & \\ \mathcal{C}[\tau^{-1}] & \xrightarrow{\quad} & \mathcal{C} & \xrightarrow{\quad} & \mathcal{C}_\tau^\wedge \\ & \curvearrowleft & & \curvearrowleft & \end{array}$$

and the pullback square of Proposition 3.79 is the one corresponding to this recollement.

A deformation is naturally enriched in filtered spectra; more loosely speaking, it is enriched in spectral sequences.

Construction 3.81. Let \mathcal{C} be a deformation. Let $X, Y \in \mathcal{C}$. Recall the functor $i: \mathbf{Z} \rightarrow \text{FilSp}$ from Definition 2.17. Define the **filtered mapping spectrum** $\text{filmap}_{\mathcal{C}}(X, Y)$ from X to Y as the filtered spectrum

$$\text{map}_{\mathcal{C}}(i(-) \otimes X, Y): \mathbf{Z}^{\text{op}} \longrightarrow \text{Sp},$$

where $\text{map}_{\mathcal{C}}(-, -)$ denotes the mapping spectrum of \mathcal{C} . This is naturally functorial in X and Y , leading to a functor $\text{filmap}: \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \text{FilSp}$. Concretely, the $\text{filmap}_{\mathcal{C}}(X, Y)$ is given by

$$\cdots \longrightarrow \text{map}_{\mathcal{C}}(\Sigma^{0,1} X, Y) \longrightarrow \text{map}_{\mathcal{C}}(X, Y) \longrightarrow \text{map}_{\mathcal{C}}(\Sigma^{0,-1} X, Y) \longrightarrow \cdots$$

with transition maps induced by τ .

3.6.1 Monoidal deformations

One way to obtain the structure of a deformation on \mathcal{C} is to give it the additional structure of an *algebra* over FilSp . The universal property of FilSp provides a way to construct this. In the following, we are careful with the distinction between monoidal and symmetric monoidal functors, because of the existence of important examples where these functors are not symmetric. Nevertheless, in the case of synthetic spectra, these issues do not arise, so the distinction will not appear much in the later text.

Recall the symmetric monoidal functor $i: \mathbf{Z} \rightarrow \text{FilSp}$ from Definition 2.17 which, by Remark 3.14, is of the form

$$\cdots \xrightarrow{\tau} \mathbf{S}^{0,-1} \xrightarrow{\tau} \mathbf{S} \xrightarrow{\tau} \mathbf{S}^{0,1} \xrightarrow{\tau} \cdots .$$

The universal property of FilSp says that it is the universal presentable stable ∞ -category on a diagram of this form. In more loose terms, this says it is the universal category on the endomorphism τ of the unit.

Proposition 3.82 (Universal property of filtered spectra).

- (1) *Let \mathcal{C} be a presentable stable ∞ -category. Then there is an equivalence*

$$\text{Fun}(\mathbf{Z}, \mathcal{C}) \simeq \text{LFun}(\text{FilSp}, \mathcal{C})$$

such that, if $f: \mathbf{Z} \rightarrow \mathcal{C}$ corresponds to $F: \text{FilSp} \rightarrow \mathcal{C}$, then we have a natural isomorphism

$$F \circ i \cong f.$$

- (2) *Let \mathcal{C} be a presentably (symmetric) monoidal stable ∞ -category. Then there is an equivalence*

$$\text{Fun}^{\otimes}(\mathbf{Z}, \mathcal{C}) \simeq \text{LFun}^{\otimes}(\text{FilSp}, \mathcal{C}),$$

where Fun^{\otimes} denotes the ∞ -category of (symmetric) monoidal functors. Moreover, when forgetting the (symmetric) monoidal structure on the functors, this equivalence coincides with the equivalence from item (1), and the natural isomorphism $F \circ i \cong f$ is naturally (symmetric) monoidal.

Proof. Note that we have a symmetric monoidal equivalence $\text{FilSp} \simeq \text{Sp}(\text{PSh}(\mathbf{Z}))$, where $\text{PSh}(\mathbf{Z})$ is also equipped with the Day convolution symmetric monoidal structure. The first universal property is therefore the combination of the universal property of presheaves from [Ker, Tag 03W9] and the presentable universal property of stabilisation from [HA, Corollary 1.4.4.5]. The universal property of Day convolution in the monoidal (respectively, symmetric monoidal) case from [HA, Example 2.2.6.10] (respectively, Example 2.2.6.9 of op. cit.) then upgrades this to the second claimed equivalence. ■

Notation 3.83. Let \mathcal{C} be a presentably monoidal ∞ -category, and let $f: \mathbf{Z} \rightarrow \mathcal{C}$ be a monoidal functor.

- ◆ We typically reserve the letter ρ for for the monoidal functor $\text{FilSp} \rightarrow \mathcal{C}$ induced by f . This functor in particular turns \mathcal{C} into a (left) FilSp -module, which informally is given by (where $A \in \text{FilSp}$ and $X \in \mathcal{C}$)

$$A \otimes X := \rho(A) \otimes X.$$

We summarise this by saying that ρ gives \mathcal{C} the structure of a **monoidal deformation**.^[3]

- ◆ We typically reserve the letter σ for the (lax monoidal) right adjoint $\mathcal{C} \rightarrow \text{FilSp}$ to ρ . We call σ the **signature functor**; if $X \in \mathcal{C}$, then we refer to σX as the **signature** of X .
- ◆ If \mathcal{C} is symmetric monoidal and f is a symmetric monoidal functor, then ρ is naturally symmetric monoidal, and σ is naturally lax symmetric monoidal. In this case, we say ρ gives \mathcal{C} the structure of a **symmetric monoidal deformation**.

Roughly speaking, the functor ρ is characterised by preserving colimits and by sending τ in FilSp to the map $f(-1 \rightarrow 0)$. Thus, we can think of f as a ‘prescription’ for what the map τ in \mathcal{C} ought to be.

Example 3.84. Let R be a filtered \mathbf{E}_2 -ring. Then $\text{Mod}_R(\text{FilSp})$ is naturally an \mathbf{E}_1 -monoidal ∞ -category, and the functor $R \otimes -: \text{FilSp} \rightarrow \text{Mod}_R(\text{FilSp})$ is a monoidal functor, turning $\text{Mod}_R(\text{FilSp})$ into a monoidal deformation. If R is a filtered \mathbf{E}_∞ -ring, then this turns it into a symmetric monoidal deformation.

In a precise sense, this example includes many deformations: all symmetric monoidal deformations with compact unit that are generated by the spheres are of this form; see Section 5.2. ▲

^[3]This is a bad choice of terminology. If \mathcal{O} is an ∞ -operad, one should define an \mathcal{O} -monoidal deformation as an \mathcal{O} -algebra in $\text{Pr}_{\text{st}}^{\text{L}}$. A monoidal functor $f: \mathbf{Z} \rightarrow \mathcal{C}$ gives rise to an \mathbf{E}_1 -monoidal functor $\text{FilSp} \rightarrow \mathcal{C}$, but this does not equip \mathcal{C} with the structure of an \mathbf{E}_1 -algebra in $\text{Pr}_{\text{st}}^{\text{L}}$. (This is the usual difference between an \mathbf{E}_1 -algebra in $\text{Mod}_A(\mathcal{C})$ and an \mathbf{E}_1 -map $A \rightarrow R$ in \mathcal{C} .) However, in these notes, we are mostly concerned with the symmetric monoidal (i.e., \mathbf{E}_∞) case where this distinction goes away, so this abuse of terminology does not have many ramifications for the rest of these notes.

Using that ρ is monoidal, we can import structure from FilSp into \mathcal{C} . For example, the \mathbf{E}_∞ -structure on $C\tau$ induces an \mathbf{E}_1 -structure on $\rho(C\tau)$. If ρ is symmetric monoidal, then $\rho(C\tau)$ also acquires an \mathbf{E}_∞ -structure.

Remark 3.85. Using that σ is right adjoint to ρ , it follows that for $X \in \mathcal{C}$, the filtered spectrum σX is given by

$$\cdots \longrightarrow \text{map}_{\mathcal{C}}(f(1), X) \longrightarrow \text{map}_{\mathcal{C}}(f(0), X) \longrightarrow \text{map}_{\mathcal{C}}(f(-1), X) \longrightarrow \cdots,$$

where $\text{map}_{\mathcal{C}}(-, -)$ denotes the mapping spectrum of \mathcal{C} , and where the transition maps are induced by f . We see that, after forgetting the lax monoidal structure on σ , it is naturally isomorphic to $\text{filmap}(\mathbf{1}_{\mathcal{C}}, -)$ from Construction 3.81.

Remark 3.86. It is difficult in general to obtain symmetric monoidal functors out of \mathbf{Z} , as it is not free as a symmetric monoidal ∞ -category. We learned from Shaul Barkan that t -structures are a source of such functors, using the Whitehead filtration; see [BvN25, Section 2]. We will use this in Section 4.3 to give synthetic spectra the structure of a symmetric monoidal deformation.

So far, we have used a deformation structure on an ∞ -category \mathcal{C} to define similar-looking operations in \mathcal{C} as we have in FilSp , such as inverting and modding out by τ . In the monoidal case, the functor σ shows that these operations translate back to the respective operations in FilSp . This both gives a more concrete interpretation of these operations, and also ties the story of deformations into the study of spectral sequences.

To prove the desired properties of σ , we use the following lemma. We refer to, e.g., [NPR24, Definition 3.5] for a definition of the projection map, and Section 3 of op. cit. for an introduction to these ideas.

Lemma 3.87. *Let \mathcal{C} be a monoidal deformation. Let $A \in \text{FilSp}$ and $X \in \mathcal{C}$. Consider the natural projection map*

$$A \otimes \sigma(X) \longrightarrow \sigma(\rho(A) \otimes X).$$

- (1) *If A is dualisable, then the projection map is an isomorphism for all $X \in \mathcal{C}$.*
- (2) *If σ preserves colimits, then the projection map is an isomorphism for all $A \in \text{FilSp}$ and all $X \in \mathcal{C}$.*

Proof. The first part follows from [NPR24, Lemma 3.8 (b)]. The second part follows from the fact that, in this case, both sides preserve colimits in A and that FilSp is generated under colimits by dualisable objects. ■

In particular, the functor σ commutes with bigraded suspensions.

Theorem 3.88. *Let \mathcal{C} be a (symmetric) monoidal deformation which arises from a (symmetric) monoidal functor $f: \mathbf{Z} \rightarrow \mathcal{C}$.*

(1) For every $X \in \mathcal{C}$, we have a natural isomorphism

$$\tau_{\sigma X} \cong \sigma(\tau_X).$$

In particular, σ sends τ -invertible objects in \mathcal{C} to τ -invertible (a.k.a. constant) filtered spectra.

(2) There are commutative diagrams of lax (symmetric) monoidal functors

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{\sigma} & \text{FilSp} \\ (-)/\tau \downarrow & & \downarrow (-)/\tau \\ \text{Mod}_{C\tau}(\mathcal{C}) & \xrightarrow{\sigma} & \text{Mod}_{C\tau}(\text{FilSp}). \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{C} & \xrightarrow{\sigma} & \text{FilSp} \\ (-)_{\tau}^{\wedge} \downarrow & & \downarrow (-)_{\tau}^{\wedge} \\ \mathcal{C}_{\tau}^{\wedge} & \xrightarrow{\sigma} & \text{FilSp}_{\tau}^{\wedge}. \end{array}$$

In particular, for $X \in \mathcal{C}$, we have natural isomorphisms

$$\sigma(C\tau \otimes X) \cong C\tau \otimes \sigma(X) \quad \text{and} \quad \sigma(X_{\tau}^{\wedge}) \cong \sigma(X)_{\tau}^{\wedge}.$$

(3) The functor σ preserves small colimits if and only if the monoidal unit of \mathcal{C} is compact.

If this happens, then there is a commutative diagram of lax (symmetric) monoidal functors

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{\sigma} & \text{FilSp} \\ \tau^{-1} \downarrow & & \downarrow \tau^{-1} \\ \mathcal{C}[\tau^{-1}] & \xrightarrow{\sigma} & \text{FilSp}[\tau^{-1}]. \end{array}$$

In particular, if σ preserves colimits, then for $X \in \mathcal{C}$, we have a natural isomorphism

$$\sigma(X[\tau^{-1}]) \cong \sigma(X)[\tau^{-1}].$$

(4) The functor σ is conservative if and only if the image of f generates \mathcal{C} as a stable ∞ -category under colimits.

Proof. Item (1) and the $C\tau$ -part of item (2) follow from Lemma 3.87 (1). To prove that σ also preserves τ -completion, we need to check that σ preserves $C\tau$ -equivalences and preserves $C\tau$ -local objects. The first again follows from the projection formula, and the second follows because ρ preserves $C\tau$ -acyclics (being an FilSp -linear functor).

For item (3), note that since σ is exact, it preserves colimits if and only if it preserves filtered colimits. The latter is equivalent to its left adjoint ρ preserving compact objects. For this, it is equivalent to check that ρ sends a collection of compact generators of FilSp to compact objects of \mathcal{C} . As the filtered spheres $\mathbf{S}^{0,s}$ for $s \in \mathbf{Z}$ form stable generators for FilSp , and $\rho(\mathbf{S}^{0,s}) \cong f(s)$, we see that σ preserves colimits if and only if f lands in compact objects of \mathcal{C} . Since f is monoidal and hence sends the unit to the unit, the latter

condition implies that the unit of \mathcal{C} is compact. Conversely, if the unit of \mathcal{C} is compact, then all dualisable objects of \mathcal{C} are compact. Because f is monoidal and all objects of \mathbf{Z} are dualisable, it then follows that all values of f are compact. We conclude that σ preserves colimits if and only if the unit of \mathcal{C} is compact.

If σ preserves colimits, then Lemma 3.87 (2) implies that σ preserves τ -inversion.

Finally, for item (4), we use again that the filtered spheres are generators, so that σX is zero if and only if for all $s \in \mathbf{Z}$, the mapping spectrum

$$\mathrm{map}_{\mathrm{FilSp}}(\mathbf{S}^{0,s}, \sigma X) \cong \mathrm{map}_{\mathcal{C}}(\rho(\mathbf{S}^{0,s}), X) \cong \mathrm{map}_{\mathcal{C}}(f(s), X)$$

vanishes. This shows the final claim. ■

Remark 3.89. As explained in [NPR24, Definition 3.9, Remark 3.10], it follows from Theorem 3.88 (3) that ρ is an *internal left adjoint* in FilSp -linear ∞ -categories (i.e., its right adjoint σ is itself an FilSp -linear functor) if and only if the unit of \mathcal{C} is compact.

We now introduce the notion of the signature spectral sequence. We will be more brief, as we will develop this in detail in the next chapter in the case of synthetic spectra (which is our main case of interest).

Definition 3.90. Let \mathcal{C} be a monoidal deformation. If the unit of \mathcal{C} is compact, then the functor $\pi_{n,s}: \mathcal{C} \rightarrow \mathrm{Ab}$ defined by

$$\pi_{n,s}(-) := [\Sigma^{n,s} \mathbf{1}_{\mathcal{C}}, -]$$

preserves filtered colimits. It follows that for $X \in \mathcal{C}$, the spectral sequence underlying σX is of the form

$$E_1^{n,s} = \pi_{n,s}(X/\tau) \implies [\Sigma^n \mathbf{1}_{\mathcal{C}}, X[\tau^{-1}]] \cong \pi_{n,*}(X)[\tau^{-1}].$$

Accordingly, we refer to this as the **signature spectral sequence** of X .

Theorem 3.88 tells us that we can understand this spectral sequence σX through computations in \mathcal{C} . For instance, item (2) tells us that σX is conditionally convergent whenever X is τ -complete, and that the associated graded is given by applying σ to X/τ . This is particularly useful in cases where $\mathrm{Mod}_{\mathcal{C}\tau}(\mathcal{C})$ is simpler than the original case of filtered spectra, where $\mathrm{Mod}_{\mathcal{C}\tau}(\mathrm{FilSp}) \simeq \mathrm{grSp}$ is topological in nature.

The above also tells us that the Omnibus Theorems from Section 3.5 carry over to the deformation \mathcal{C} , with all occurrences of the underlying spectral sequence replaced by the signature spectral sequence. We have to be a little careful if the unit of \mathcal{C} is not compact, as then σ need not preserve τ -inversion. Even if the unit of \mathcal{C} is not compact, the truncated Omnibus Theorem of Theorems 3.67 and 3.70 apply in \mathcal{C} without change. In the non-truncated case of Theorem 3.62, all results that do not compare $\pi_{*,*} X$ with $\pi_* X^{\tau=1}$ apply in \mathcal{C} as well. If the unit of \mathcal{C} is compact, then all of the Omnibus Theorems apply in their

entirety. In summary then, we learn that the $\mathbf{Z}[\tau]$ -module $\pi_{*,*} X$ captures the signature spectral sequence of X .

Not just the Omnibus Theorem carries over to a deformation, but also the τ -Bockstein spectral sequence we used to prove it.

Variant 3.91. Let \mathcal{C} be a monoidal deformation, and $X \in \mathcal{C}$. The τ -adic filtration on X is the filtered object $\mathbf{Z}^{\text{op}} \rightarrow \mathcal{C}$ given by

$$\dots \xrightarrow{\tau} \Sigma^{0,-2} X \xrightarrow{\tau} \Sigma^{0,-1} X \xrightarrow{\tau} X = \dots$$

Analogously to Construction 3.47, this leads to a trigraded spectral sequence which we call the τ -Bockstein spectral sequence of X , which is of the form

$$E_1^{n,w,s} \cong \pi_{n,w+s}(X/\tau) \implies \pi_{n,w} X.$$

From Theorem 3.88, it follows that when we apply σ to the τ -adic filtration on X , we obtain the τ -adic filtration on σX . It follows that the τ -BSS of X is the τ -BSS of σX . We may therefore freely use the results of Section 3.4 for this spectral sequence; in particular, it captures the signature spectral sequence of X .

Two structural properties of deformations, namely *cellularity* and *evenness*, will be discussed later in Sections 5.1 and 5.3. For now, we finish this chapter with a few examples of deformations, which are of a different flavour than the one we will meet in the next chapter.

We learned the following example from Christian Carrick and Lennart Meier. See also [BHS22, Examples A.8 and A.9].

Example 3.92. Let Sp_{C_2} denote the ∞ -category of genuine C_2 -spectra. The Euler class a_σ gives Sp_{C_2} the structure of a deformation. (To avoid notational confusion, we will avoid the names ρ and σ for the deformation functors.) More specifically, let σ denote the sign representation of C_2 . Then the inclusion of fixed points results in a map $S^0 \rightarrow S^\sigma$, which stably results in a map called the *Euler class*

$$a_\sigma: \mathbf{S}^{-\sigma} \longrightarrow \mathbf{S}.$$

We expect, but do not check in detail, that this assembles to a monoidal functor

$$\dots \xrightarrow{a_\sigma} \mathbf{S}^{-\sigma} \xrightarrow{a_\sigma} \mathbf{S} \xrightarrow{a_\sigma} \mathbf{S}^\sigma \xrightarrow{a_\sigma} \dots$$

(Note, however, that this cannot be made symmetric monoidal, due to a nontrivial switch map for $\mathbf{S}^\sigma \otimes \mathbf{S}^\sigma$.) This monoidal functor induces a monoidal left adjoint $\text{FilSp} \rightarrow \text{Sp}_{C_2}$, resulting in the promised monoidal deformation structure. The generic fibre is equivalent to Sp . Interestingly, its special fibre is also spectra: it is given by modules over $\mathbf{S}/a_\sigma \cong \Sigma_+^\infty C_2$, which by [BDS15, Theorem 1.1] is equivalent to Sp (via a lift of the restriction-coinduction adjunction).

Under these identifications, τ -inversion is identified with geometric fixed points, and τ -completion is identified with Borel completion (i.e., inverting those maps that induce an isomorphism on homotopy fixed points). As a result, the pullback square of Proposition 3.79 becomes, after taking (genuine) fixed points, the *Tate square*

$$\begin{array}{ccc} X^{C_2} & \longrightarrow & X^{hC_2} \\ \downarrow & \lrcorner & \downarrow \\ X^{\Phi C_2} & \longrightarrow & X^{tC_2}. \end{array}$$

The resulting τ -BSS is identified with the a_σ -BSS, which is the homotopy fixed-point spectral sequence.

One could mimic these constructions for Sp_G for a general finite group G in the place of C_2 . However, the resulting deformation would be rather contrived: for general G , the structure of Sp_G is better captured by taking all subgroups of G into account. Only if $G = C_2$ is the resulting structure exactly that of a deformation. \blacktriangle

Finally, we end with an example that is not strictly speaking a deformation, but which is close enough in that it also allows theorems from filtered spectra to be imported over.

Example 3.93 (Recovering the p -Bockstein spectral sequence). Fix a prime p , and consider the functor $f: \mathbf{Z} \rightarrow \mathrm{Sp}$ given by

$$\cdots \xrightarrow{p} \mathbf{S} \xrightarrow{p} \mathbf{S} \xrightarrow{p} \mathbf{S} \xrightarrow{p} \cdots.$$

This functor induces an adjunction

$$\mathrm{FilSp} \xrightleftharpoons[\sigma]{\rho} \mathrm{Sp}.$$

However, neither f nor ρ can be made monoidal. Indeed, the functor ρ sends τ to p , and hence sends $C\tau$ to \mathbf{S}/p . The latter does not admit an \mathbf{E}_1 -structure for any p , preventing ρ (and hence f) from being monoidal.

While strictly speaking not a deformation, we can still use this adjunction to import information from the τ -Bockstein spectral sequence. Namely, we can check by hand that σ sends τ_X to $\tau_{\sigma X}$. In Sp , the map τ is given by p , which σ sends to p because it is additive. Further, by Remark 4.32, the functor σ sends a spectrum X to

$$\cdots \xrightarrow{p} X \xrightarrow{p} X \xrightarrow{p} X \xrightarrow{p} \cdots$$

so indeed σ sends τ_X to $\tau_{\sigma X}$. It follows that σ sends the p -Bockstein filtration on X to the τ -Bockstein filtration on σX . The resulting trigraded spectral sequence we compute to be the p -BSS for X with an additional filtration tagged on. More specifically, at every level of this new filtration, it is the p -BSS for X , and the transition maps for this new filtration

are all given by multiplication by p . Using this, we can deduce the analogous version of Theorem 3.52 for the p -BSS, at least those parts that do not depend on any monoidality properties of τ (such as Theorem 3.52 (5)).

For a monoidal version of this example, consider instead the functor $\mathbf{Z} \rightarrow \mathcal{D}(\text{Ab})$ given by

$$\cdots \xrightarrow{p} \mathbf{Z} \xrightarrow{p} \mathbf{Z} \xrightarrow{p} \mathbf{Z} \xrightarrow{p} \cdots .$$

This functor is symmetric monoidal: it lands in the heart of $\mathcal{D}(\text{Ab})$, so equivalently is given by a functor $\mathbf{Z} \rightarrow \text{Ab}$. It is easy to check that this functor describes a strict, multiplicative filtration on the commutative ring $\mathbf{Z}[\frac{1}{p}]$, and as a result is canonically a symmetric monoidal functor. It follows that this gives $\mathcal{D}(\text{Ab})$ the structure of a symmetric monoidal deformation. We can therefore directly deduce the analogue of Theorem 3.52 by using Theorem 3.88, thereby recovering, e.g., [Pal05, Theorem 3.8].

Informally, we may summarise both situations by saying that we “put τ equal to p ” and thereby recover the p -Bockstein spectral sequence. However, this slogan should be taken with a grain of salt, as it ignores the monoidality issues raised above, which depend on the specific category one is working with. In particular, only in the monoidal case will we be able to import multiplicative properties of the τ -BSS. ▲

Chapter 4

Synthetic spectra

Previously in Section 2.5, we discussed the classical definition of the Adams spectral sequence. The more modern way to interact with Adams spectral sequences is to use *synthetic spectra*. This chapter is intended both as a first introduction to and as a manual for working with synthetic spectra. Compared to most of the existing literature, our distinctive focus is to understand these through the lens of filtered spectra.

We review the main categorical features of synthetic spectra in Sections 4.1 and 4.2. There, among other things, we encounter the synthetic map τ . In Section 4.3, we show that this gives the ∞ -category of synthetic spectra the structure of a (symmetric monoidal) deformation in the sense of Section 3.6.1, and we study the structure of this deformation. In particular, this results in a functor

$$\sigma: \text{Syn}_E \longrightarrow \text{FilSp}$$

that preserves limits and colimits and sends τ to the filtered map τ . We refer to σX as the *signature* of X . Consequently, we obtain a synthetic Omnibus Theorem, describing the homotopy groups of a synthetic spectrum X in terms of the spectral sequence underlying σX .

The question of which spectral sequence this is requires a computation. In Section 4.4, our goal is to compute this for so-called *synthetic analogues*; this is both an important foundational result, and also showcases how to work with synthetic spectra. There is a lax symmetric monoidal functor

$$v: \text{Sp} \longrightarrow \text{Syn}_E$$

called the *synthetic analogue functor*. We show that the signature of the synthetic analogue of a spectrum is its E -Adams spectral sequence; see Section 4.4, particularly Theorem 4.71. Finally, in Section 4.5 we collect some implications of this result, and briefly discuss some notational conventions.

Remark 4.1. In addition to providing a new interface for interacting with Adams spectral sequences, synthetic spectra have also been used in setting up obstruction theories; see

[Bar23; HL17; PV22]. We do not (yet) discuss these obstruction theories in this version of the notes.

For the most part, this chapter consists of an overview of results from [Pst22]. We learned much of the relationship with filtered spectra from [Bar23], [BHS22, Appendix C], [CD24, Section 1], and [Pst25]. Some results in this chapter are modifications of results appearing in [CDvN25, Section 1].

4.1 Categorical properties

Before we can do more serious work, we need to know the basic properties and structure of the category we are dealing with. We do not review the constructions given by Pstragowski, but content ourselves with summarising the main properties, preferring to work with synthetic spectra in a ‘model-independent way’.^[1]

Construction 4.2. Let E be a homotopy associative ring spectrum of Adams type. In [Pst22], Pstragowski constructs a symmetric monoidal ∞ -category Syn_E of **E -based synthetic spectra**, together with a unital lax symmetric monoidal functor $\nu: \mathrm{Sp} \rightarrow \mathrm{Syn}_E$. We call ν the **synthetic analogue functor**.

We may refer to E -based synthetic spectra as *E -synthetic spectra*, or even simply by *synthetic spectra* if E is clear from the context. On the opposite end, when we want to vary the variable E , we would write ν_E for ν , emphasising it as the *E -synthetic analogue*. It would be more principled to write $\mathrm{Syn}_E(\mathrm{Sp})$ instead of Syn_E , indicating that it is a modification of the ∞ -category Sp . For simplicity, we will stick to the shorter name in this chapter.

Notation 4.3. An **E -synthetic \mathbf{E}_∞ -ring** is an \mathbf{E}_∞ -algebra object in Syn_E . If E is clear from the context, then we may also refer to such an object as a *synthetic \mathbf{E}_∞ -ring*.

Remark 4.4. Just as with the Adams spectral sequence, the role of the ∞ -category of spectra here is not of fundamental importance: there should be a similar modification of any nice enough stable ∞ -category with a type of Adams spectral sequence. We stick to the case of spectra here to more conveniently cite Pstragowski’s construction. A more general theory can be found in [PP23], but this construction differs from the one in [Pst22] in various ways; see [PP23, Section 6.5].

Remark 4.5. Although the notation seems to suggest otherwise, the symmetric monoidal ∞ -category Syn_E depends on much less data than the ring spectrum E . This is because the Adams spectral sequence depends on less data than the ring spectrum E ; see Remark 2.89. In particular, Syn_E is not sensitive to a potential coherent multiplicative structure on E , nor does it require it for its construction as a symmetric monoidal ∞ -category.

^[1]We mean this in a loose way: we specifically choose the construction from [Pst22] over the other available ones. However, phrasing our arguments and computations in this way should make them more robust and more easily adapted to other synthetic contexts.

Warning 4.6. The ∞ -category Syn_E is functorial in E , and in fact the construction does not require the assumption that E is of Adams type. However, it only gives the correct answer if E is of Adams type. For instance, as observed by Pstrągowski and explained by Schäppi in [Sch20, Theorem 2.3.7], taking $E = \mathrm{MU}$ or $E = \mathbf{Z}$ results in the *same* symmetric monoidal ∞ -category, even though the \mathbf{Z} -Adams spectral sequence is wildly different from the MU -Adams spectral sequence (see Example 2.95).

Because of this warning, throughout this chapter we stick to the following assumption.

Notation 4.7. For the remainder of this chapter, E denotes a fixed choice of a homotopy-associative ring spectrum of Adams type.

We begin by studying some categorical properties of synthetic spectra. Recall the notion of a *finite E -projective spectrum* from Definition 2.91.

Proposition 4.8.

- (1) *The ∞ -category Syn_E is stable.*
- (2) *The ∞ -category Syn_E is presentable, and the symmetric monoidal structure preserves colimits in each variable separately; that is to say, Syn_E is presentably symmetric monoidal.*
- (3) *If P is a finite E -projective spectrum, then vP is a compact and dualisable object in Syn_E , with dual $v(P^\vee)$. In particular, the monoidal unit is compact.*
- (4) *As a stable ∞ -category, Syn_E is compactly generated under colimits by the synthetic analogues of finite E -projectives. That is, the collection of $\Sigma^k vP$, for P finite E -projective and $k \in \mathbf{Z}$, forms a set of compact dualisable generators. In particular, Syn_E is compactly generated by dualisables.*
- (5) *The monoidal ∞ -category is rigid in the sense that an object is compact if and only if it is dualisable.*

Proof. Items (1) and (2) are [Pst22, Proposition 4.2], and items (3) and (4) are [Pst22, Remark 4.14]. Item (5) then follows formally from the fact that the unit is compact and that it has a set of compact dualisable generators; see, e.g., (the footnote to) [NPR24, Terminology 4.8]. ■

Next, we turn to properties of the functor v .

Proposition 4.9.

- (1) *The functor $v: \mathrm{Sp} \rightarrow \mathrm{Syn}_E$ is fully faithful, additive, and preserves filtered colimits. In particular, v preserves arbitrary coproducts.*
- (2) *Consider a cofibre sequence of spectra*

$$X \xrightarrow{f} Y \xrightarrow{g} Z.$$

Then the induced sequence

$$\nu X \xrightarrow{\nu f} \nu Y \xrightarrow{\nu g} \nu Z$$

is a cofibre sequence of synthetic spectra if and only if

$$0 \longrightarrow E_* X \xrightarrow{f_*} E_* Y \xrightarrow{g_*} E_* Z \longrightarrow 0$$

is short exact, or in other words, if the boundary map $Z \rightarrow \Sigma X$ is zero on E_* -homology.

- (3) The comparison map $\nu X \otimes \nu Y \rightarrow \nu(X \otimes Y)$ coming from the lax monoidal structure on ν is an isomorphism whenever X or Y is a filtered colimit of finite E -projective spectra.

More generally, if the E_* -homology of X or Y is flat as an E_* -module, then the map $\nu X \otimes \nu Y \rightarrow \nu(X \otimes Y)$ is a νE -equivalence.

Proof. Item (1) follows from [Pst22, Lemma 4.4 and Corollary 4.38], item (2) is [Pst22, Lemma 4.23], and item (3) is [Pst22, Lemma 4.24]. ■

Both conditions of Proposition 4.9 (3) are a type of flatness condition. This is obvious for the second one. For the first, compare this with the algebraic result that a module over a ring is flat if and only if it can be written as a filtered colimit of finite free modules; see [Stacks, Tag 058G].

Example 4.10. The definition of Adams type directly implies that $\nu E \otimes X \rightarrow \nu(E \otimes X)$ is an isomorphism for all spectra X . ▲

Example 4.11. Suppose $E = \mathbf{F}_p$, or more generally that E is a homotopy-associative ring spectrum such that $\pi_* E$ is a graded field. Then every finite spectrum is finite E -projective. The smallest subcategory of Sp that contains all finite spectra and is closed under filtered colimits is equal to all of Sp . We therefore learn from Proposition 4.9 (3) that ν is actually a strong symmetric monoidal functor if $\pi_* E$ is a graded field. ▲

Remark 4.12. The reader may gain intuition for the above properties by thinking of ν as being similar to the Whitehead filtration functor $\mathrm{Wh}: \mathrm{Sp} \rightarrow \mathrm{FilSp}$. In fact, this is more than a formal analogy: in the case $E = \mathbf{S}$, the ∞ -category $\mathrm{Syn}_{\mathbf{S}}$ is equivalent to $\mathrm{Mod}_{\mathrm{Wh}_{\mathbf{S}}}(\mathrm{FilSp})$, and this equivalence identifies ν with the Whitehead filtration functor; see Corollary 5.13.

Take particular note that ν is *not* an exact functor, even though it is a functor between stable ∞ -categories. For example, Proposition 4.9 (2) implies that $\Sigma(\nu X) \cong \nu(\Sigma X)$ if and only if $E_* X = 0$. In terms of the E -Adams spectral sequences, having vanishing E -homology is a very degenerate case, so the functor ν practically never preserves suspensions.

The difference between suspending in spectra and in synthetic spectra has a conceptual meaning as well: the former has the effect of shifting its Adams spectral sequence one to the right, while suspending its synthetic analogue also shifts it down by one filtration.

This is made precise by the following definition of the *synthetic bigraded spheres*. The indexing convention we use here turns out to be the most practical; we defer a more detailed explanation to Remark 4.37.

Definition 4.13 (Synthetic bigraded spheres). Let n and s be integers.

- (1) The **synthetic (n, s) -sphere** is

$$\mathbf{S}^{n,s} := \Sigma^{-s} \nu(\mathbf{S}^{n+s}).$$

We refer to n as the **stem**, and to s as the **filtration**.

- (2) We write $\Sigma^{n,s}: \text{Syn}_E \rightarrow \text{Syn}_E$ for the functor given by tensoring with $\mathbf{S}^{n,s}$ on the left.
 (3) We write $\pi_{n,s}: \text{Syn}_E \rightarrow \text{Ab}$ for the functor

$$\pi_{n,s}(-) := [\mathbf{S}^{n,s}, -].$$

- (4) The map $\tau: \mathbf{S}^{0,-1} \rightarrow \mathbf{S}^{0,0}$ is the colimit-comparison map

$$\tau: \mathbf{S}^{0,-1} = \Sigma(\nu\mathbf{S}^{-1}) \longrightarrow \nu\mathbf{S} = \mathbf{S}^{0,0}.$$

If X is a synthetic spectrum, then tensoring it with the map $\tau: \mathbf{S}^{0,-1} \rightarrow \mathbf{S}$ results in a map $\Sigma^{0,-1}X \rightarrow X$, which we denote by τ_X , or by τ when there is no risk of confusion.

Note that $\pi_{*,*}$ naturally lifts to a functor $\text{Syn}_E \rightarrow \text{Mod}_{\mathbf{Z}[\tau]}(\text{bigrAb})$.

Remark 4.14. If X is a spectrum, then we also have the natural colimit-comparison map $\Sigma\nu(\Sigma^{-1}X) \rightarrow \nu X$. This coincides with the map $\tau \otimes \nu X$ by [Pst22, Proposition 4.28].

Remark 4.15 (Cellularity). It is not necessarily true that bigraded homotopy groups detect isomorphisms of synthetic spectra. If this is the case, we say that Syn_E is *cellular*. For many E , this is the case. We discuss this issue more in Section 5.1. For applications to spectral sequences, one can equally well work with the cellularisation of Syn_E , so we view this as a technicality.

Remembering the precise definition of $\mathbf{S}^{n,s}$ is not the most important; it is enough to remember the following key facts.

Example 4.16.

- (1) For every n , the synthetic spectrum $\nu(\mathbf{S}^n)$ is the bigraded sphere $\mathbf{S}^{n,0}$. This is the first instance where we see that ν places everything in *Adams filtration zero*; see Section 4.5 below for a further discussion. Later we will also see that these are the only synthetic spheres that are in the essential image of ν : see Example 4.66.

We will abuse notation and abbreviate $\mathbf{S}^{0,0}$ simply by \mathbf{S} , and refer to it as the **synthetic sphere**. Most of the time, the context will allow one to infer whether the sphere spectrum or the synthetic sphere is meant by this notation.

(2) More generally, if X is a spectrum, then we have a natural isomorphism

$$\Sigma^{n,0} \nu X \cong \nu(\Sigma^n X).$$

Indeed, this follows since ν is strong symmetric monoidal when one factor is a sphere; see Proposition 4.9 (3).

(3) Categorical suspension is given by the bigraded suspension $\Sigma^{1,-1}$. ▲

Remark 4.17 (Koszul sign rule). It is possible to set up matters so that if A is a homotopy-commutative algebra in Syn_E , then $\pi_{*,*} A$ becomes a bigraded ring with a Koszul sign rule according to the first variable (i.e., the stem). Doing this involves choices, as explained by Dugger [Dug14], see also [Dug+24]. The choice described by Pstrągowski in [Pst22, Remark 4.10], and explained in detail by Chua in [Chu22, Section 6], results in this sign on homotopy groups.

As with any symmetric monoidal stable ∞ -category, synthetic spectra have an internal notion of homology. We regard it as a bigraded object.

Notation 4.18. Let A and X be synthetic spectra, and let n and s be integers. We write $A_{n,s}(X)$ for $\pi_{n,s}(A \otimes X)$.

Finally, we make a few comments about notation and indexing compared to the literature.

Remark 4.19. In the specific case of \mathbf{F}_p -synthetic spectra, it is becoming more and more common to use the letter λ to denote the map otherwise denoted by τ ; see, e.g., [BIX25] (particularly Section 1.1 therein). This is done to allow for computations that involve both BP-synthetic and \mathbf{F}_p -synthetic arguments at the same time. Because these notes are not aimed at these computations, we will still use the letter τ even in the \mathbf{F}_p -synthetic case.

Although the grading convention of Definition 4.13 has become more standard, it is not the only one in the literature. We refer to the indexing of Definition 4.13 as **Adams grading** of synthetic spectra. This is not the convention used in [Pst22], which instead follows the *motivic grading*. Unless explicitly said otherwise, we will not use motivic grading in these notes.

Remark 4.20 (Motivic grading). The **motivic grading** on synthetic spectra is to define

$$\mathbf{S}^{t,w} := \Sigma^{t-w} \nu(\mathbf{S}^w).$$

Conversion from Adams to motivic grading, and vice versa, is given respectively by

$$(n, s) \longmapsto (n, n + s) \quad \text{and} \quad (t, w) \longmapsto (t, w - t).$$

The only cases in which motivic grading agrees with Adams grading are those where the stem is 0. (In particular, τ has bidegree $(0, -1)$ in both conventions.) In [Pst22], the degree w is called the **weight**, and the difference $t - w$ is called the **Chow degree**. In Adams grading, the weight of $\mathbf{S}^{n,s}$ is given by $n + s$, while the Chow degree is given by $-s$. The

motivic grading is designed to match with the indexing conventions of motivic homotopy theory; see Section 5.4 for more information.

4.2 The homological t-structure

One of the most important features of Syn_E that distinguishes it from the category of FilSp is the existence of a particular t-structure. Following Burklund–Hahn–Senger [BHS23], we refer to this t-structure as the *homological t-structure*; it is referred to as the *natural t-structure* in [Pst22]. Although this is indeed the default t-structure on synthetic spectra for our purposes, we prefer this more descriptive name.

As the name suggests, the defining feature of this t-structure is that it looks at vE -homology of E -synthetic spectra to measure (co)connectivity, *not* at the bigraded homotopy groups. This has both its upsides and downsides. On the one hand, this homology tends to be a lot simpler than the homotopy, due to the special role that E plays for E -synthetic spectra. On the other hand, this means that taking truncations or connective covers can have very unpredictable effects on bigraded homotopy groups.

Remark 4.21. Carrick and Davies [CD24, Section 2] define t-structures on Syn_E based on the bigraded homotopy groups.

Definition 4.22. The **homological t-structure** on Syn_E is the t-structure where a synthetic spectrum X is connective if and only if

$$vE_{n,s}(X) = 0 \quad \text{whenever } s > 0.$$

We will write $\tau_{\geq n}$ and $\tau_{\leq n}$ for the n -connective cover and n -truncation functors with respect to this t-structure, respectively.

Theorem 4.23. *The homological t-structure is an accessible t-structure on Syn_E that satisfies the following.*

(a) *A synthetic spectrum X is connective if and only if*

$$vE_{n,s}(X) = 0 \quad \text{whenever } s > 0.$$

(b) *A synthetic spectrum X is 0-truncated if and only if X is vE -local and*

$$vE_{n,s}(X) = 0 \quad \text{whenever } s < 0.$$

(c) *Let X be a synthetic spectrum. The connective cover $\tau_{\geq 0}X \rightarrow X$ induces an isomorphism*

$$vE_{n,s}(\tau_{\geq 0}X) \xrightarrow{\cong} vE_{n,s}(X) \quad \text{whenever } s \leq 0.$$

Likewise, the 0-truncation $X \rightarrow \tau_{\leq 0}X$ induces an isomorphism

$$vE_{n,s}(X) \xrightarrow{\cong} vE_{n,s}(\tau_{\leq 0}X) \quad \text{whenever } s \geq 0.$$

- (d) For every spectrum X , the synthetic spectrum νX is connective.
(e) There exists a monoidal equivalence of categories

$$\mathrm{Syn}_E^\heartsuit \simeq \mathrm{grComod}_{E_*E}$$

under which $\Sigma^{n,0}$ becomes the functor $[n]$, and which fits into a commutative diagram of lax monoidal functors

$$\begin{array}{ccc} \mathrm{Sp} & \xrightarrow{\tau_{\leq 0\nu}} & \mathrm{Syn}_E^\heartsuit \\ & \searrow E_*(-) & \downarrow \simeq \\ & & \mathrm{grComod}_{E_*E}. \end{array}$$

Moreover, if E is homotopy commutative, then this equivalence and the above diagram are naturally symmetric monoidal.

- (f) The t-structure is right complete (but for general E , not even left separated).
(g) The t-structure is compatible with filtered colimits.
(h) The t-structure is compatible with the monoidal structure.

Proof. Property (a) alone determines this t-structure uniquely. By Theorem 4.18 of [Pst22], this definition of the t-structure agrees with the definition from Proposition 2.16 of op. cit., which has the remaining desired properties by Propositions 4.16, 4.18, and 4.21 of op. cit. ■

In fact, as we will re-prove in Example 4.66 below, the νE -homology of νX has a very simple form: there is an isomorphism of bigraded $\mathbf{Z}[\tau]$ -modules

$$\nu E_{**}(\nu X) \cong E_*X[\tau]$$

where $E_n X$ is placed in bidegree $(n, 0)$. This property has useful consequences for working with ν ; see Remark 4.46 for instance. It is also one of the first instances where we see that Syn_E is more suited for working with Adams spectral sequences than FilSp : as we will see in Warning 4.35, the filtered spectrum underlying νX is rarely connective in the diagonal t-structure on FilSp , making it harder to work with that filtered spectrum directly.

Remark 4.24. It follows from item (g) that $(\mathrm{Syn}_E)_{\geq 0}$ is a Grothendieck prestable ∞ -category; see [SAG, Proposition C.1.4.1].

Remark 4.25. There exist examples for which the monoidal equivalence $\mathrm{Syn}_E^\heartsuit \simeq \mathrm{grComod}_{E_*E}$ cannot be made symmetric monoidal. One can think of this as saying that the ‘correct’ braiding on $\mathrm{grComod}_{E_*E}$ is not the usual algebraic one, but rather a more exotic ‘topological’ one. For an example of this phenomenon, see [HL17, Section 6] and [BP23, Section 4], where $E = K(n)$ and where Sp is replaced by $(K(n)$ -local) modules over Morava E -theory.

Finally, let us make a few comments regarding notation and indexing.

Remark 4.26. The description of the connective objects is somewhat confusing, in that an object is connective when certain groups in a *positive* degree vanish. This clash is because the filtration in Adams spectral sequences is indexed cohomologically, while (at least in homotopy theory) we usually index t-structures homologically. Arguably, it would be less confusing to index this t-structure cohomologically instead (as is more common in algebraic geometry), writing $\tau^{\leq n}$ for what we normally denote by $\tau_{\geq -n}$, and $\tau^{\geq n}$ for $\tau_{\leq -n}$. However, to prevent confusion with the standard convention in homotopy theory, we will refrain from doing this.

Warning 4.27. Often with t-structures, one writes π_n^\heartsuit for the functor $\Sigma^{-n} \tau_{\leq n} \tau_{\geq n}$ considered as landing in the heart of the t-structure. Because this t-structure is measured by homology instead of homotopy, this notation can get confusing: the functor π_n^\heartsuit is *not* given by bigraded homotopy groups. Instead, for $X \in \text{Syn}_E$, by (the comment following) [Pst22, Theorem 4.18], we have an isomorphism of graded E_*E -comodules

$$\pi_n^\heartsuit(X) \cong vE_{*+n, -n}(X) = vE_{*, -n}(X)[-n].$$

Note also the minus sign in the filtration on the right-hand side; this is again due to the difference between homological and cohomological grading (cf. Remark 4.26). To avoid the potential confusion with the bigraded homotopy groups, we will generally avoid the notation π_n^\heartsuit .

4.3 Synthetic spectra as a deformation

With the foundational properties and structure in hand, we can relate synthetic spectra to spectral sequences. We begin by defining a deformation structure.

Lemma 4.28. *There is a natural symmetric monoidal structure on the functor $\mathbf{Z} \rightarrow \text{Syn}_E$ given by the multiplication-by- τ tower on the unit:*

$$\dots \xrightarrow{\tau} \mathbf{S}^{0,-1} \xrightarrow{\tau} \mathbf{S} \xrightarrow{\tau} \mathbf{S}^{0,1} \xrightarrow{\tau} \dots$$

Proof. We follow the argument given in the proof of [Law24b, Corollary 6.1]. In Pstragowski's model [Pst22], the synthetic sphere $\mathbf{S}^{0,s}$ is defined as the sheafification of the presheaf $\tau_{\geq -s} \text{map}(-, \mathbf{S})$, with τ induced by the suspension-comparison map. The Whitehead filtration functor is lax symmetric monoidal (Remark 2.25), as is sheafification, so the \mathbf{E}_∞ -structure on the sphere spectrum induces a symmetric monoidal structure on the multiplication-by- τ tower. ■

We will use the notation and terminology introduced in Notation 3.83; let us repeat it here for convenience.

Notation 4.29. By the universal property of FilSp from Proposition 3.82, the symmetric

monoidal functor $\mathbf{Z} \rightarrow \mathrm{Syn}_E$ from Lemma 4.28 induces an adjunction

$$\mathrm{FilSp} \begin{array}{c} \xrightarrow{\rho} \\ \xleftarrow{\sigma} \end{array} \mathrm{Syn}_E$$

where the left adjoint ρ is a symmetric monoidal functor. As a result, the functor σ is naturally lax symmetric monoidal. If X is a synthetic spectrum, then we refer to the filtered spectrum σX as its **signature**.

Remark 4.30 (History). This functor has appeared before in [BHS22, Appendix C] under the name i_* . The name *signature* was introduced in [CD24], with the letter σ starting to be used in [CDvN25; CDvN24] (and in later revisions of [CD24]).

The functor ρ lets us import important structure from FilSp . For instance, for every $k \geq 1$, the synthetic spectrum $C\tau^k$ inherits an \mathbf{E}_∞ -structure from the filtered spectrum $C\tau^k$.^[2]

As explained in Section 3.6, the functor σ can be thought of as an ‘underlying spectral sequence’ functor. More precisely, as a consequence of Theorem 3.88, it sends the synthetic map τ_X to the transition map of σX , and preserves modding out by τ . We now check that it also preserves colimits, and in particular preserves τ -inversion.

Proposition 4.31.

- (1) *The functor σ preserves colimits, and is even FilSp -linear. In particular, σ preserves τ -inversion, and ρ is an internal left adjoint in FilSp -linear ∞ -categories.*
- (2) *The functor σ is conservative if and only if Syn_E is cellular.*

Proof. This follows from Theorem 3.88 (3) and (4), and Remark 3.89, using that the synthetic sphere is compact, and that $\mathbf{S}^{0,s}$ for $s \in \mathbf{Z}$ form stable generators if and only if Syn_E is cellular. ■

Remark 4.32. As a special case of Remark 3.85, the functor σ can be described as follows. Write $\mathrm{map}(-, -)$ for the mapping spectrum functor of the stable ∞ -category Syn_E . Then σ is given by levelwise applying $\mathrm{map}(\mathbf{S}, -)$ to the multiplication-by- τ tower functor. In diagrams: for $X \in \mathrm{Syn}_E$, the filtered spectrum σX is given by

$$\dots \xrightarrow{\tau} \mathrm{map}(\mathbf{S}, \Sigma^{0,-1}X) \xrightarrow{\tau} \mathrm{map}(\mathbf{S}, X) \xrightarrow{\tau} \mathrm{map}(\mathbf{S}, \Sigma^{0,1}X) \xrightarrow{\tau} \dots$$

Remark 4.33. The adjunction $\rho \dashv \sigma$ is very close to a monadic adjunction. More precisely, it is a monadic adjunction if and only if Syn_E is cellular. What requires more assumptions is to then identify the monad on filtered spectra without making reference to the synthetic category. We discuss these things more in Section 5.2.

^[2]Alternatively, one can use Proposition 4.47 and the monoidality of the homological t-structure to give $C\tau$ an \mathbf{E}_∞ -structure; this is how it is done in [Pst22, Corollary 4.30].

4.3.1 The signature spectral sequence

The deformation picture tells us how to understand the bigraded homotopy groups of a synthetic spectrum. Namely, the synthetic bigraded spheres are in the image of the left adjoint ρ ; as a result, understanding synthetic homotopy groups is equivalent to understanding the filtered homotopy groups of σ applied to the synthetic spectrum. The latter, as explained by the Omnibus Theorem, captures a spectral sequence.

The only subtlety in this story is that there is a reindexing taking place when passing between filtered and synthetic spectra. To avoid confusion, we will for the moment distinguish the filtered and synthetic settings by writing

$$\mathbf{S}_{\text{fil}}^{n,s} \quad \text{and} \quad \mathbf{S}_{\text{syn}}^{n,s}$$

for the filtered and synthetic spheres, respectively, and similarly $\pi_{*,*}^{\text{fil}}$ and $\pi_{*,*}^{\text{syn}}$ for the homotopy groups.

Proposition 4.34.

(1) For all n and s , we have an isomorphism

$$\rho(\mathbf{S}_{\text{fil}}^{n,s}) \cong \mathbf{S}_{\text{syn}}^{n,s-n}.$$

(2) For all n , we have a natural isomorphism of graded $\mathbf{Z}[\tau]$ -modules (where $X \in \text{Syn}_E$)

$$\pi_{n,*}^{\text{syn}}(X) \cong \pi_{n,*+n}^{\text{fil}}(\sigma X).$$

(3) The functor ρ is right t -exact (with respect to the diagonal t -structure on FilSp and the homological t -structure on Syn_E); equivalently, the functor σ is left t -exact.

Proof. The functor ρ is characterised by preserving colimits and sending $\mathbf{S}_{\text{fil}}^{0,s}$ to $\mathbf{S}_{\text{syn}}^{0,s}$ for all s . In particular, ρ is exact, so it preserves arbitrary suspensions. Using the identifications

$$\mathbf{S}_{\text{fil}}^{n,s} \cong \Sigma^n \mathbf{S}_{\text{fil}}^{0,s} \quad \text{and} \quad \Sigma_{\text{syn}}(-) \cong \mathbf{S}_{\text{syn}}^{1,-1} \otimes -,$$

the first isomorphism follows. Using that ρ is left adjoint to σ , this implies that for every n and s , we have a natural isomorphism of abelian groups (where $X \in \text{Syn}_E$)

$$\pi_{n,s}^{\text{syn}}(X) \cong \pi_{n,s+n}^{\text{fil}}(\sigma X).$$

By Theorem 3.88, the functor σ sends τ_X to $\tau_{\sigma X}$, so this assembles to the claimed isomorphism of graded $\mathbf{Z}[\tau]$ -modules.

For the final claim, recall that $(\text{FilSp})_{\geq 0}$ is the smallest subcategory generated under colimits by the objects $\mathbf{S}_{\text{fil}}^{n,s}$ for $n - s \geq 0$. Since synthetic analogues are connective, it follows that $\mathbf{S}_{\text{syn}}^{k,u} = \Sigma^{-u} \nu(\mathbf{S}^{k+u})$ is $(-u)$ -connective in the homological t -structure. It follows that ρ sends $\mathbf{S}_{\text{fil}}^{n,s}$ for $n - s \geq 0$ to a connective synthetic spectrum. Because ρ preserves colimits, it follows that it restricts to a functor $(\text{FilSp})_{\geq 0} \rightarrow (\text{Syn}_E)_{\geq 0}$, proving the claim. \blacksquare

Warning 4.35. Even though $\rho(\mathbf{S}_{\text{fil}}^{n,s})$ is a synthetic sphere, the filtered spectrum $\sigma(\mathbf{S}_{\text{syn}}^{n,s})$ is *very* different from a filtered sphere. Indeed, the spectral sequence associated to a filtered sphere is uninteresting (see Example 2.32), while the spectral sequence underlying $\sigma(\mathbf{S}_{\text{syn}}^{n,s})$ is (a shift of) the E -Adams spectral sequence for the sphere spectrum (see Theorem 4.71 below), which is very interesting and highly nontrivial for many E . In particular, ρ is very far from preserving bigraded homotopy groups. This also implies σ is not right t-exact.

Proposition 4.34 tells us that the synthetic homotopy groups capture a spectral sequence; we give it a special name.

Definition 4.36. Let X be a synthetic spectrum. The **signature spectral sequence** of X is the spectral sequence underlying the filtered spectrum σX .

We will use second-page indexing for this spectral sequence, so that it is of the form

$$E_2^{n,s} = \pi_{n,s}(C\tau \otimes X) \implies \pi_n(X[\tau^{-1}]).$$

The Omnibus Theorem makes precise the way in which the synthetic homotopy groups capture this spectral sequence. One has to be slightly careful in that we reindexed the above to start on the second page, which subtracts one from powers of τ in the Omnibus Theorem. For example, d_r -differentials in the signature spectral sequence of X introduce τ^{r-1} -torsion in $\pi_{*,*} X$. We spell this out in the non-truncated case in Section 4.5 below.

Remark 4.37. We use second-page indexing because, for X a synthetic analogue, the filtered spectrum σX is the *décalage* of an Adams spectral sequence. As a result, it makes most sense to index this spectral sequence to agree with the usual indexing for Adams spectral sequences. The definition of synthetic spheres from Definition 4.13 was chosen exactly to fit with second-page indexing. Because we use first-page indexing on filtered spectra, this has the unfortunate side effect of causing the reindexing as in Proposition 4.34 (1). The reindexing $(n, s) \mapsto (n, s - n)$ of Proposition 4.34 is precisely the reindexing of Remark 2.40.

Remark 4.38. The distinction between the similar, but different, terms *signature* and *signature spectral sequence* is intentional. The former is a filtered spectrum, and as a result is able to capture more intricate structures (e.g., \mathbf{E}_n -structures), while the latter is only an algebraic object. However, we will not need to make this distinction very often.

The bare formalism only takes us so far: it does not tell us which spectral sequences arise in this way, nor does it tell us what the structure of τ -inverted synthetic spectra or $C\tau$ -modules in synthetic spectra are. We will investigate the second question first, leaving the computation of signatures for Section 4.4. As advertised in the introduction, we will show that $C\tau$ -modules are of an algebraic nature, forming a type of derived ∞ -category of an abelian category. This gives signature spectral sequences a structural advantage over the one coming from a bare filtered spectrum: the starting page is, in a sense, entirely algebraic.

Before we begin, let us briefly record the definition (and reindexing) of the τ -BSS in the synthetic setting.

Variante 4.39. If X is a synthetic spectrum, then its **τ -adic filtration** is the filtered synthetic spectrum $\mathbf{Z}^{\text{op}} \rightarrow \text{Syn}_E$ given by

$$\dots \xrightarrow{\tau} \Sigma^{0,-2} X \xrightarrow{\tau} \Sigma^{0,-1} X \xrightarrow{\tau} X = \dots$$

Analogously to Construction 3.47, this leads to a trigraded spectral sequence that we call the **τ -Bockstein spectral sequence** of X , which is of the form

$$E_1^{n,w,s} \cong \begin{cases} \pi_{n,w+s}(X/\tau) & \text{if } s \geq 0 \\ 0 & \text{else} \end{cases} \implies \pi_{n,w} X.$$

Its differential d_r^τ is of tridegree $(-1, 1, r)$ for $r \geq 1$. Note that when we apply σ to the τ -adic filtration on X , we obtain the τ -adic filtration on σX . Using Proposition 4.34, it follows that the τ -BSS of X is merely a reindexing of the τ -BSS of σX ; the reindexing is given by

$$E_1^{n,w,s}(X) \cong E_1^{n,w+n,s}(\sigma X).$$

We may therefore freely use the results of Section 3.4 for this spectral sequence; in particular, it captures the signature spectral sequence of X . Note that this is a situation where we use second-page indexing for ordinary spectral sequences (coming to us from the conventions for synthetic spectra), but nevertheless use first-page indexing for the corresponding τ -BSS; see Remark 3.55 for a further discussion of this. If we use second page indexing for both (see the previously cited remark), then we obtain the indexing of the τ -BSS used in, e.g., [BHS23, Theorem A.8].

4.3.2 Inverting τ

Recall that the functor $\nu: \text{Sp} \rightarrow \text{Syn}_E$ is fully faithful. However, since it is not an exact functor, we should not think too strongly of the image of ν as an embedding of spectra into synthetic spectra. If we make ν exact in a universal way, then this does result in a good embedding of spectra into synthetic spectra, and these happen to be exactly the τ -invertible synthetic spectra.

As in Notation 3.74, we write $\text{Syn}_E[\tau^{-1}]$ for the generic fibre of the deformation Syn_E . By Remark 3.75, this is the full subcategory of Syn_E on the τ -invertible synthetic spectra, and is moreover a smashing localisation of Syn_E .

Definition 4.40. Write $\mathfrak{y}: \text{Sp} \rightarrow \text{Syn}_E$ for the functor $\nu(-)[\tau^{-1}]$.

The functor \mathfrak{y} is also referred to as the *spectral Yoneda embedding*. By definition, \mathfrak{y} lands in τ -invertible synthetic spectra.

Theorem 4.41 ([Pst22], Theorem 4.37). *The functor \mathfrak{J} is fully faithful, exact, and symmetric monoidal, and restricts to a symmetric monoidal equivalence*

$$\mathfrak{J} : \mathbf{Sp} \xrightarrow{\simeq} \mathbf{Syn}_E[\tau^{-1}].$$

Notation 4.42. We write $(-)^{\tau=1}$ for the composite

$$\mathbf{Syn}_E \xrightarrow{\tau^{-1}} \mathbf{Syn}_E[\tau^{-1}] \simeq \mathbf{Sp}.$$

The following example is the analogous one to Example 3.22.

Example 4.43. Recall the definition $\mathbf{S}^{n,s} = \Sigma^{-s} \nu(\mathbf{S}^{n+s})$ from Definition 4.13. As τ -inversion is an exact functor on synthetic spectra, it preserves suspensions, so we find that

$$\mathbf{S}^{n,s}[\tau^{-1}] = \Sigma^{-s} \nu(\mathbf{S}^{n+s})[\tau^{-1}] = \Sigma^{-s} \mathfrak{J}(\mathbf{S}^{n+s}) \cong \mathfrak{J}(\mathbf{S}^n).$$

In other words, $(\mathbf{S}^{n,s})^{\tau=1} \cong \mathbf{S}^n$. We can think of this as saying that inverting τ forgets the Adams filtration. \blacktriangle

This identification is compatible with our earlier identification of $\mathbf{FilSp}[\tau^{-1}]$ from Section 3.2.1.

Proposition 4.44. *The adjunction $\rho \dashv \sigma$ restricts to an adjoint equivalence between τ -invertible objects. Moreover, this equivalence fits into a commutative diagram of symmetric monoidal equivalences*

$$\begin{array}{ccc} & \mathbf{Sp} & \\ \text{Const} \swarrow & & \searrow \mathfrak{J} \\ \mathbf{FilSp}[\tau^{-1}] & \xrightleftharpoons[\sigma]{\rho} & \mathbf{Syn}_E[\tau^{-1}]. \end{array}$$

In particular, for every spectrum X , the colimit of the filtered spectrum $\sigma(\nu X)$ is naturally isomorphic to X .

Proof. The functors ρ Const and \mathfrak{J} are symmetric monoidal colimit-preserving functors $\mathbf{Sp} \rightarrow \mathbf{Syn}_E$. By the universal property of \mathbf{Sp} , it follows that they are naturally isomorphic as symmetric monoidal functors; see [HA, Corollary 4.8.2.19]. By two-out-of-three, the functor ρ restricts to an equivalence between τ -invertible objects. Because σ is right adjoint to ρ , it follows that σ restricts to an inverse for it.

The final claim follows from the isomorphism $\sigma(\nu(X)[\tau^{-1}]) \cong \text{Const } X$ and the fact that σ preserves τ -inversion by Proposition 4.31. \blacksquare

So far, we have focussed on ν and defined \mathfrak{J} in terms of it. It is also possible to go in the other direction and characterise ν in terms of \mathfrak{J} , using the homological t-structure.

Proposition 4.45. *Let X be a spectrum.*

(1) The τ -inversion map

$$\nu X \longrightarrow \mathfrak{J}(X)$$

is a connective cover with respect to the homological t -structure.

(2) There is a natural isomorphism of functors $\mathbf{Z}^{\text{op}} \rightarrow \text{Syn}_E$ between the Whitehead filtration of $\mathfrak{J}(X)$,

$$\cdots \longrightarrow \tau_{\geq 1} \mathfrak{J}(X) \longrightarrow \tau_{\geq 0} \mathfrak{J}(X) \longrightarrow \tau_{\geq -1} \mathfrak{J}(X) \longrightarrow \cdots$$

and the multiplication-by- τ tower on νX ,

$$\cdots \xrightarrow{\tau} \Sigma^{0,-1} \nu X \xrightarrow{\tau} \nu X \xrightarrow{\tau} \Sigma^{0,1} \nu X \xrightarrow{\tau} \cdots$$

Proof. The first is [Pst22, Proposition 4.36], and the second follows from the same argument as in Lemma 4.28. \blacksquare

Remark 4.46. The functor $\nu: \text{Sp} \rightarrow \text{Syn}_E$ is neither a left nor right adjoint, as it is not even an exact functor. When considered as landing in connective synthetic spectra however, its categorical properties improve: it is then right adjoint to inverting τ . This follows from the isomorphism $\nu \cong \tau_{\geq 0} \circ \mathfrak{J}$ and by pasting adjunctions: the horizontal composites in

$$\text{Sp} \xleftarrow[\mathfrak{J}]{(-)^{\tau=1}} \text{Syn}_E[\tau^{-1}] \xleftarrow{\tau^{-1}} \text{Syn}_E \xleftarrow[\tau_{\geq 0}]{} (\text{Syn}_E)_{\geq 0}$$

form the adjunction

$$\text{Sp} \xleftarrow[\nu]{(-)^{\tau=1}} (\text{Syn}_E)_{\geq 0}.$$

4.3.3 Modding out by τ

Although τ -invertible synthetic spectra are equivalent to τ -invertible filtered spectra, modules over the cofibre of τ are very different in synthetic spectra compared to filtered spectra. This distinction is controlled by the homological t -structure.

Proposition 4.47. *For every spectrum X , the natural map $\nu X \rightarrow C\tau \otimes \nu X$ exhibits the target as the 0-truncation of the source. In particular, we have a natural isomorphism*

$$C\tau \otimes \nu X \cong E_*(X)$$

where we regard the right-hand side as an element of $\text{grComod}_{E_*E} \simeq \text{Syn}_E^\heartsuit$.

Proof. The first statement is [Pst22, Lemma 4.29]. The final claim follows by combining this with Theorem 4.23 (e). \blacksquare

Morally, $C\tau$ -modules in Syn_E are equivalent to the derived ∞ -category of graded E_*E -comodules. This is not entirely true, as for most E , the unit of $\mathcal{D}(\text{grComod}_{E_*E})$ is not a compact object.^[3] Accordingly, to make such a statement true, we need a modification of this derived ∞ -category to have a compact unit. It turns out that this modification is actually an enlargement.

The idea behind this construction is that instead of inverting homology-isomorphisms of chain complexes of comodules, we should invert homotopy-isomorphisms. Hovey [Hov04] constructs a model category doing this; see also the introduction to op. cit. for a further motivation for this construction. Barthel–Heard–Valenzuela [Bar+21] give the following description of the underlying ∞ -category of Hovey’s model category; see also [Pst22, Section 3.2] for a summary.

Recollection 4.48. Let (A, Γ) be a graded Hopf algebroid. Write Perf_Γ for the thick subcategory of $\mathcal{D}(\text{grComod}_{(A, \Gamma)})$ generated by the dualisable (ordinary, non-derived) comodules over (A, Γ) (considered as objects in the heart the derived ∞ -category). Define the **stable comodule ∞ -category** as

$$\text{Stable}_{(A, \Gamma)} = \text{Ind}(\text{Perf}_\Gamma).$$

The inclusion functor $\text{Perf}_\Gamma \rightarrow \mathcal{D}(\text{grComod}_{(A, \Gamma)})$ induces a functor $\text{Stable}_{(A, \Gamma)} \rightarrow \mathcal{D}(\text{grComod}_{(A, \Gamma)})$, and this turns out to have a fully faithful right adjoint:

$$\text{Stable}_{(A, \Gamma)} \xleftarrow{\text{fully faithful}} \mathcal{D}(\text{grComod}_{(A, \Gamma)}). \quad (4.49)$$

As Perf_Γ is closed under tensor products, the ∞ -category $\text{Stable}_{(A, \Gamma)}$ is naturally a symmetric monoidal functor, and the localisation functor to the derived is symmetric monoidal. By definition, the unit of $\text{Stable}_{(A, \Gamma)}$ is compact. Moreover, the localisation (4.49) is precisely given by Γ -localisation, i.e., inverting those maps that become isomorphisms after tensoring with Γ .

If E is a homotopy-associative ring spectrum, then we also write Stable_{E_*E} for the stable comodule ∞ -category of the Hopf algebroid (E_*, E_*E) .

Theorem 4.50. *There is a right t -exact fully faithful left adjoint of monoidal ∞ -categories*

$$\text{Mod}_{C\tau}(\text{Syn}_E) \hookrightarrow \text{Stable}_{E_*E}$$

with the following properties.

^[3]Note that this is a phenomenon that does not occur in $\mathcal{D}(\text{Mod}_A)$ or $\mathcal{D}(\text{grMod}_A)$ for a (graded) commutative ring A , and is specific to the comodule setting.

(1) *This functor sits in a commutative diagram of lax monoidal functors*

$$\begin{array}{ccc} \mathrm{Sp} & \xrightarrow{\nu(-)/\tau} & \mathrm{Mod}_{C\tau}(\mathrm{Syn}_E) \\ E_*(-) \downarrow & & \downarrow \\ \mathrm{grComod}_{E_*E} & \xrightarrow{\quad} & \mathrm{Stable}_{E_*E}. \end{array}$$

(2) *If E is Landweber exact or is the sphere spectrum, then this functor is an equivalence.*

(3) *If E is homotopy commutative, then this functor is naturally symmetric monoidal, and the diagram of (1) is naturally one of lax symmetric monoidal functors.*

Proof. The main result and item (3) are [Pst22, Theorem 4.46]. Item (1) follows by combining this with Theorem 4.23 (e), and item (2) is [Pst22, Proposition 4.53]. \blacksquare

In particular, the special fibre of synthetic spectra is entirely algebraic (except for possibly the braiding if E is not homotopy commutative). The failure of the above functor to be essentially surjective is the problem of Stable_{E_*E} not being generated by the objects E_*P where P ranges over the finite E -projective spectra. We regard this as a minor technical issue.

Remark 4.51. We now have two different generalisations of E -homology for synthetic spectra, namely νE -homology and $C\tau$ -homology. They both differ from E -homology for ordinary spectra, but in different ways.

- ◆ For νE -homology, we obtain a second grading, and even a $\mathbf{Z}[\tau]$ -module structure. Although the νE -homology of a synthetic analogue is very simple (having no τ -torsion for instance, see Example 4.66 below), the νE -homology of a general synthetic spectrum can be a highly nontrivial $\mathbf{Z}[\tau]$ -module.
- ◆ For $C\tau$ -homology, this takes values in (a modification of) the derived ∞ -category of E_*E -comodules. For a synthetic analogue, this lands in the heart, and by Proposition 4.47 is identified with E -homology in the ordinary sense. For a general synthetic spectrum however, the resulting object will rarely be an honest E_*E -comodule, but will generally be a derived or stable one.

Remark 4.52. Strengthening Proposition 4.47, the essential image of ν in fact consists precisely of those synthetic spectra X that are connective and for which $C\tau \otimes X$ is discrete in the homological t-structure; see [PV22, Proposition 2.16]. Using Theorem 4.50, we can equivalently state this condition as asking $C\tau \otimes \nu X$ to be an honest comodule, rather than a derived one.

Recall that the derived ∞ -category is obtained from the stable comodule ∞ -category by localising at E_*E -equivalences. Translated into synthetic terms, E_*E corresponds to $\nu E/\tau$, leading to the following. Moreover, note that the generation issues go away after νE -localisation, and we obtain an actual equivalence.

Theorem 4.53 ([Pst22], Theorem 4.54). *The functor from Theorem 4.50 restricts to a monoidal equivalence*

$$L_{\nu E} \text{Mod}_{C\tau}(\text{Syn}_E) \simeq \mathcal{D}(\text{grComod}_{E_*E})$$

which is naturally symmetric monoidal if E is homotopy commutative.

Notation 4.54. Following [Pst22], we will also write $\widehat{\text{Syn}}_E$ for $L_{\nu E} \text{Syn}_E$. In op. cit., objects of $\widehat{\text{Syn}}_E$ are called *hypercomplete*, stemming from their definition as sheaves; we will not use this name, and instead refer to them simply as νE -local objects. We warn the reader that, while this notation is convenient when working with a fixed E , it could lead to confusion when working with various E at once. In these notes, we will always work with a fixed E , so this confusion should not arise.

Warning 4.55. For general E , the unit in $\widehat{\text{Syn}}_E$ is not compact for general E , because the unit of $\mathcal{D}(\text{grComod}_{E_*E})$ is usually not compact. As a result, although $\widehat{\text{Syn}}_E$ is a symmetric monoidal deformation in its own right (obtained by νE -localising the functor ρ), the resulting right adjoint $\widehat{\text{Syn}}_E \rightarrow \text{FilSp}$ does *not* preserve colimits; see Theorem 3.88 (3).

At this point, we can begin to see the Adams spectral sequence appearing, at least its second page.

Example 4.56. Let X and Y be spectra. Combining Proposition 4.47 and Theorem 4.53, we learn that $C\tau$ -linear maps between synthetic analogues are computed by maps of comodules:

$$[\nu Y/\tau, \nu X/\tau]_{C\tau} \cong \text{Hom}_{E_*E}(E_*Y, E_*X).$$

For an integer k , let us denote the k -fold grading-shift functor on $\mathcal{D}(\text{grComod}_{E_*E})$ by $[k]$, and let us write Σ^k for the k -fold ∞ -categorical suspension as usual. By Proposition 4.47, the synthetic spectra $\nu Y/\tau$ and $\nu X/\tau$ are 0-truncated, which by Theorem 4.23 (b) in particular means they are νE -local. Using Theorem 4.53 and the definition $\mathbf{S}^{n,s} = \Sigma^{-s} \nu(\mathbf{S}^{n+s})$, it therefore follows that

$$[\Sigma^{n,s} \nu Y/\tau, \nu X/\tau]_{C\tau} \cong [(E_*Y)[n+s], \Sigma^s E_*X]_{\mathcal{D}(\text{grComod}_{E_*E})} = \text{Ext}_{E_*E}^{s, n+s}(E_*Y, E_*X).$$

In particular, we have

$$\pi_{n,s}(\nu X/\tau) = [\mathbf{S}^{n,s}, \nu X/\tau] \cong [C\tau \otimes \mathbf{S}^{n,s}, \nu X/\tau]_{C\tau} \cong \text{Ext}_{E_*E}^{s, n+s}(E_*, E_*X). \quad \blacktriangle$$

4.3.4 Synthetic lifts

Previously, we argued that the τ -inversion of a synthetic spectrum can be thought of as an ‘underlying spectrum’. We can also turn this question around, fixing a spectrum and asking how many synthetic spectra have this as their underlying spectrum. It is useful to introduce some terminology for this.

Definition 4.57. Let X be a spectrum. A **synthetic lift** of X is a synthetic spectrum S such that $S^{\tau=1} \cong X$.

Example 4.58. Theorem 4.41 says that ν provides a functorial synthetic lift. ▲

We think of a synthetic lift of a spectrum X as encoding a *modified Adams spectral sequence* for X . The synthetic analogue of X , from this perspective, is the standard synthetic lift; as we will see in Theorem 4.71, it encodes the ordinary Adams spectral sequence for X . For a further discussion of these ideas, see [CD24].

We end this section by discussing how to construct synthetic lifts out of old ones. As τ -inversion preserves colimits, taking colimits results in a synthetic lift of the colimit of the underlying spectra. More subtle is the use of limits, since τ -inversion does not preserve all limits. For instance, if S is a synthetic spectrum, then every term in its τ -adic tower

$$\cdots \longrightarrow S/\tau^3 \longrightarrow S/\tau^2 \longrightarrow S/\tau$$

vanishes upon τ -inversion; meanwhile, the τ -inversion of the limit is $(S_\tau^\wedge)[\tau^{-1}]$, which is nontrivial for many S (e.g., if $S = \nu X$ for X an E -nilpotent complete spectrum, by Theorem 4.71). Nevertheless, if the diagram is of a special form, then τ -inversion does preserve the limit.

Proposition 4.59. *Let $X: I \rightarrow \mathrm{Sp}$ be a diagram of spectra. Then we have an isomorphism*

$$\nu(\lim X) \cong \tau_{\geq 0}(\lim \nu(X)),$$

and the limit-comparison map

$$(\lim \nu(X))^{\tau=1} \longrightarrow \lim X$$

is an isomorphism of spectra. In particular, $\lim \nu(X)$ is a synthetic lift of $\lim X$.

Because ν is fully faithful, this in fact says that τ -inversion preserves the limit of any diagram that takes values in synthetic analogues.

The key input for the proof is the homological t-structure, particularly Remark 4.46 and the following lemma.

Lemma 4.60 ([Pst22], Lemma 4.35). *If S is a bounded above synthetic spectrum, then $S[\tau^{-1}]$ is zero. In particular, if S is any synthetic spectrum, then for every n , the map $\tau_{\geq n} S \rightarrow S$ becomes an isomorphism upon τ -inversion.*

Proof. For every integer s , the suspension $\Sigma^{0,s}$ decreases coconnectivity by s ; this follows directly from the coconnectivity criterion of Theorem 4.23 (b). Since νE -homology preserves filtered colimits, it follows that if S is bounded above, then the colimit

$$S[\tau^{-1}] = \mathrm{colim}(S \xrightarrow{\tau} \Sigma^{0,1} S \xrightarrow{\tau} \Sigma^{0,2} S \xrightarrow{\tau} \cdots),$$

is $(-\infty)$ -coconnective. Since the t-structure on Syn_E is right complete by Theorem 4.23 (f), it follows that $S[\tau^{-1}]$ vanishes.

The final claim follows from the fact that the cofibre of $\tau_{\geq n}S \rightarrow S$ is $\tau_{\leq n-1}S$, which in particular is bounded above, and that τ -inversion is an exact functor. ■

Proof of Proposition 4.59. Since ν takes values in connective synthetic spectra, we may consider $\nu \circ X$ as landing in $(\mathrm{Syn}_E)_{\geq 0}$. By Remark 4.46, the functor $\nu: \mathrm{Sp} \rightarrow (\mathrm{Syn}_E)_{\geq 0}$ is right adjoint to τ -inversion, implying that ν sends limits of spectra to limits in $(\mathrm{Syn}_E)_{\geq 0}$. A limit in the connective subcategory is computed as the connective cover of the limit in Syn_E , so we find that

$$\nu(\lim X) = \tau_{\geq 0}(\lim \nu(X)).$$

As τ -inversion is left inverse to ν , it follows that the right-hand side τ -inverts to $\lim X$. The claim now follows from Lemma 4.60. ■

4.4 The signature of a synthetic analogue

Previously in Section 2.5.2, we defined a (cosimplicial) model for the Adams spectral sequence, resulting in a functor $\mathrm{Sp} \rightarrow \mathrm{FilSp}$. The better and more modern definition is the following.

Definition 4.61. The *E*-based Adams filtration is the functor $\sigma \circ \nu_E: \mathrm{Sp} \rightarrow \mathrm{FilSp}$.

The goal of this section is to give a justification for this name: in Theorem 4.71, we show that this spectral sequence agrees with the décalage of the *E*-based Adams spectral sequence as defined in Definition 2.86. (For an explanation why the décalage appears, see Variant 2.87 and the discussion preceding it.) However, as we pointed out before, the point of this is not to let go of the synthetic origins of this functor, but rather to demonstrate that this recovers the correct notion.

One concrete reason for preferring this definition over the old one is the following.

Remark 4.62 (Lax monoidality; [PP23], Section 5.5). The functor $\sigma \circ \nu_E$ is a composite of two lax symmetric monoidal functors, making it a lax symmetric monoidal functor. This only requires *E* to be homotopy-associative. By contrast, to turn the classical definition of the *E*-Adams filtration into a lax symmetric monoidal functor, one would need an \mathbf{E}_∞ -structure on *E*. Such a structure does not always exist in cases of interest (e.g., BP or Morava K-theories), and the Adams spectral sequence does not depend on it, so it is not desirable to require these structures. For a further discussion, and a way to construct this for *E* having only a left-unital multiplication, see [PP23, Section 5.5].

Our proof strategy is to first show this comparison on resolution objects, which for the *E*-Adams spectral sequence are the homotopy *E*-modules. This relies on a computation of the synthetic homotopy groups of *E*-modules. The general case follows from this by descending from a resolution by such objects.

Remark 4.63 (History). Theorem 4.71 is not new and is well-known to experts, but has not been written down in this specific form. In [PP23, Proposition 5.56 and Theorem 5.60],

Patchkoria and Pstrągowski prove this result for the synthetic-like categories they construct therein; while formally these categories are different, the proofs follow the same ideas. Another closely-related result is [Pst25, Theorem 6.26], whose proof we follow closely in this section. Similar results in the nilpotent-complete case can be found in [Pst22, Remark 4.64] and [BHS23, Appendix A.1]. A proof of the nilpotent-complete case also appeared in [CDvN25, Section 1.4], which this section is an adaptation of.

The following holds for homotopy classes of maps $Y \rightarrow X$ between two spectra, but for simplicity we record it only for homotopy groups. In words, it says that the spectral sequence underlying νX is concentrated in nonnegative filtrations, and that the underlying spectrum of the filtration $\sigma(\nu X)$ is given by X .

Proposition 4.64 ([Pst22], Theorem 4.58). *Let X be a spectrum. Then for all $s \leq 0$ and all n , inverting τ induces a natural isomorphism*

$$\pi_{n,s}(\nu X) \xrightarrow{\cong} \pi_n X.$$

Phrased differently: inverting τ induces a natural isomorphism of bigraded $\mathbf{Z}[\tau]$ -modules

$$\pi_{*, \leq 0}(\nu X) \cong \pi_*(X)[\tau],$$

where $\pi_n X$ is placed in bidegree $(n, 0)$.

For the didactic value, we include Pstrągowski's proof.

Proof. The cofibre sequence

$$\nu X \xrightarrow{\tau} \Sigma^{0,1} \nu X \longrightarrow \Sigma^{0,1} \nu X / \tau$$

gives rise to a long exact sequence on bigraded homotopy groups; by Example 4.56, part of this reads

$$\mathrm{Ext}_{E_* E}^{s-2, n+s-1}(E_*, E_* X) \longrightarrow \pi_{n,s}(\nu X) \xrightarrow{\tau} \pi_{n,s-1}(\nu X) \longrightarrow \mathrm{Ext}_{E_* E}^{s-1, n+s-1}(E_*, E_* X).$$

Note that $\mathrm{Ext}_{E_* E}^{s,t}(E_*, E_* X) = 0$ whenever $s < 0$. Therefore if $s \leq 0$, we see that the two outer terms vanish, so that the map in the middle is an isomorphism. As a result, to prove the claim, we only have to show that τ -inversion induces an isomorphism

$$\pi_{n,0}(\nu X) = [\nu \mathbf{S}^n, \nu X] \longrightarrow [\mathbf{S}^n, X] = \pi_n X.$$

This follows from the fact that τ -inversion is a left inverse to ν (see Theorem 4.41) and that ν is fully faithful. ■

For a particularly nice class of spectra, this computes the entirety of their synthetic homotopy groups.

Proposition 4.65 ([Pst22], Proposition 4.60). *Let M be a spectrum admitting a homotopy E -module structure. Then inverting τ induces a natural isomorphism of bigraded $\mathbf{Z}[\tau]$ -modules*

$$\pi_{*,*}(\nu M) \xrightarrow{\cong} \pi_*(M)[\tau],$$

where $\pi_n M$ is placed in bidegree $(n, 0)$.

In words, this says that the signature spectral sequence for νM is concentrated in filtration zero, and as a result collapses without any differentials. Again we include Pstragowski's proof.

Proof. Using the previous result, we only have to show that $\pi_{n,s}(\nu M)$ vanishes when $s \geq 1$. We first show this for $s = 1$. Since M is a homotopy E -module, the Hurewicz homomorphism

$$E_*(-): \pi_n M \longrightarrow \mathrm{Hom}_{E_*E}(E_*[n], E_*M)$$

is an isomorphism; see [Pst22, Remark 3.18]. Under the isomorphism $\pi_n M \cong \pi_{n,0}(\nu M)$, the Hurewicz homomorphism is the right-most map in the exact sequence

$$\mathrm{Ext}_{E_*E}^{-1,n}(E_*, E_*M) \longrightarrow \pi_{n,1}(\nu M) \xrightarrow{\tau} \pi_{n,0}(\nu M) \longrightarrow \mathrm{Ext}_{E_*E}^{0,n}(E_*, E_*M).$$

As the Ext group on the left vanishes, we learn that $\pi_{n,1}(\nu M) = 0$ for all n .

Next, we consider the case $s > 1$. Since M is a homotopy E -module, it follows from [Pst22, Remark 3.18] that we have an isomorphism of graded comodules

$$E_*M \cong E_*E \otimes_{E_*} M_*,$$

implying that

$$\mathrm{Ext}_{E_*E}^{s,t}(E_*, E_*M) \cong \mathrm{Ext}_{E_*}^{s,t}(E_*, M_*).$$

In particular, we see that these Ext groups vanish whenever $s \geq 1$. By the long exact sequence, this means that multiplication by τ induces an isomorphism

$$\tau: \pi_{n,s+1}(\nu X) \xrightarrow{\cong} \pi_{n,s}(\nu X)$$

for all $s \geq 1$. We previously showed that $\pi_{*,1}(\nu M) = 0$, so we are done. ■

Example 4.66. Recall from Example 4.10 that for all spectra X , we have an isomorphism $\nu E \otimes \nu X \cong \nu(E \otimes X)$. Because $E \otimes X$ is a homotopy E -module, we learn from Proposition 4.65 that

$$\nu E_{*,*}(\nu X) = \pi_{*,*}(\nu E \otimes \nu X) \cong \pi_{*,*}(\nu(E \otimes X)) \cong E_*(X)[\tau].$$

In particular, by Theorem 4.23 (a), this shows that νX is connective in the homological t-structure. We now see why this is independent of the connectivity of the spectrum X : the connectivity of νX is about the collapse of the E -Adams spectral sequence for $\pi_*(E \otimes X)$.

We learn a number of things from this computation.

- (1) The shift $\Sigma^{0,s} \nu X$ for $s \neq 0$ is not in the essential image of ν (unless $E_*(X)$ vanishes).
 (2) The $\mathbf{Z}[\tau]$ -module $\nu E_{*,*}(\nu X)$ is τ -torsion free. As a result, we learn that

$$\nu E_{*,*}(C\tau \otimes \nu X) \cong (\nu E_{*,*}(\nu X))/\tau \cong E_*(X),$$

where we mean the quotient by τ in the (non-derived) algebraic sense. This explains (apart from the νE -locality) why $C\tau \otimes \nu X$ is 0-truncated in the homological t-structure; cf. Theorem 4.23 (b).

- (3) Inverting τ on $\nu E_{*,*}(\nu X)$ yields

$$\nu E_{*,*}(\mathcal{Y}X) = \nu E_{*,*}(\nu X[\tau^{-1}]) \cong E_*(X)[\tau^{\pm}].$$

This gives an indication of why $\nu X \rightarrow \mathcal{Y}X$ is a connective cover, and more generally, why the Whitehead tower of $\mathcal{Y}X$

$$\cdots \longrightarrow \tau_{\geq 1}(\mathcal{Y}X) \longrightarrow \tau_{\geq 0}(\mathcal{Y}X) \longrightarrow \tau_{\geq -1}(\mathcal{Y}X) \longrightarrow \cdots$$

can be identified with the multiplication-by- τ tower on νX

$$\cdots \xrightarrow{\tau} \Sigma^{0,-1} \nu X \xrightarrow{\tau} \nu X \xrightarrow{\tau} \Sigma^{0,1} \nu X \xrightarrow{\tau} \cdots \quad \blacktriangle$$

We can restate Proposition 4.65 in terms of the signature of νM .

Corollary 4.67. *Let M be a spectrum admitting a homotopy E -module structure. Then there is a natural isomorphism of filtered spectra*

$$\sigma(\nu M) \cong \text{Wh } M$$

which is naturally a symmetric monoidal natural transformation in M .

Proof. Because σ preserves τ -inversion by Proposition 4.31, we find that for any spectrum X , applying σ to the τ -inversion map $\nu X \rightarrow \nu X[\tau^{-1}]$ results in a natural (symmetric monoidal) transformation

$$\sigma(\nu X) \longrightarrow \sigma(\nu X)[\tau^{-1}] \cong \text{Const } X, \quad (4.68)$$

where we use the identification from Proposition 4.44. If now M is a homotopy E -module spectrum, then Proposition 4.65 implies that $\sigma(\nu M)$ is connective in the diagonal t-structure on filtered spectra. Indeed, combining Proposition 4.65 with Proposition 4.34 (2), we see that the group

$$\pi_{n,s}(\sigma(\nu M)) \cong \pi_{n,s-n}(\nu M)$$

vanishes whenever $s - n > 0$, that is, whenever $n < s$. As a result, the natural map (4.68) in the case $X = M$ factors through a natural map

$$\sigma(\nu M) \longrightarrow \tau_{\geq 0}^{\text{diag}}(\text{Const } M) = \text{Wh } M.$$

Moreover, this factorisation is through a symmetric monoidal transformation, because the diagonal t-structure on filtered spectra is monoidal (Proposition 2.23 (g)). To establish that it is an isomorphism, it suffices to show that each component $\sigma(\nu M)^s \rightarrow \tau_{\geq s} M$ is an isomorphism for all s . As this map is induced by τ -inversion, this is the other part of Proposition 4.65. ■

Although it follows directly from Proposition 4.65 that the signature spectral sequence of νM converges to M , a more refined argument even shows that νM is τ -complete in Syn_E itself. (If Syn_E is not cellular, then this is not automatic from completeness of $\sigma(\nu M)$.)

Corollary 4.69 ([BHS23], Lemma A.15). *Let M be a spectrum admitting a homotopy E -module structure. Then νM is τ -complete.*

Again for didactic value, we include the proof given by Burklund–Hahn–Senger.

Proof. We have to show that the limit (as $s \rightarrow \infty$) of

$$\dots \xrightarrow{\tau} \Sigma^{0,-s} \nu M \xrightarrow{\tau} \dots \xrightarrow{\tau} \Sigma^{0,-1} \nu M \xrightarrow{\tau} \nu M$$

vanishes. It is enough to check this on mapping spectra out of νP for all finite E -projective spectra P . As νP is dualisable with dual $\nu(P^\vee)$ by Proposition 4.8 (3), we find that

$$\text{map}(\nu P, \lim_s \Sigma^{0,-s} \nu M) \cong \lim_s \text{map}(\mathbf{S}^{0,s}, \nu(P^\vee) \otimes \nu M) \cong \lim_s \text{map}(\mathbf{S}^{0,s}, \nu(P^\vee \otimes M)),$$

where for the latter isomorphism we use Proposition 4.9 (3). The n -th homotopy group of the mapping spectrum at stage s on the right-hand side is given by

$$\pi_{n,s-n}(\nu(P^\vee \otimes M)).$$

Because $P^\vee \otimes M$ admits a homotopy E -module structure (since M does), these groups vanish when $n < s$ by Proposition 4.65. It follows that this mapping spectrum is s -connective, so that the limit over s is ∞ -connective, and therefore vanishes. ■

We are now ready to show that the spectral sequence underlying $\sigma(\nu X)$ is the E -based Adams spectral sequence for general X , or more precisely, that it captures its décalage. To show this, we may choose any preferred E -resolution of X to compute the E -based Adams spectral sequence; we use the one from Definition 2.86.

Construction 4.70. Since E is a homotopy ring spectrum, its unit map $\mathbf{S} \rightarrow E$ gives rise to a semicosimplicial spectrum $E^{[\bullet]}: \Delta_{\text{inj}} \rightarrow \text{Sp}$ which receives a map from \mathbf{S} . Tensoring this resulting diagram with a spectrum X , we obtain a diagram of spectra

$$X \longrightarrow E^{[\bullet]} \otimes X.$$

Applying $\sigma \circ \nu$ to this diagram, we obtain a diagram of filtered spectra

$$\sigma(\nu X) \longrightarrow \sigma(\nu(E^{[\bullet]} \otimes X)).$$

Note that for every $n \geq 1$, the spectrum $E^{\otimes n} \otimes X$ has the structure of a homotopy E -module. By Corollary 4.67, it follows that $\sigma(\nu(E^{\otimes n} \otimes X)) \cong \text{Wh}(E^{\otimes n} \otimes X)$. As such, the above diagram induces a map

$$\sigma(\nu X) \longrightarrow \text{Tot}(\text{Wh}(E^{[\bullet]} \otimes X)) = \text{Déc}^\Delta(E^{[\bullet]} \otimes X).$$

The target of this map is $\text{Déc}(\text{ASS}_E(X))$, as follows by combining Proposition 2.82 and Definition 2.86.

Theorem 4.71. *Let X be any spectrum.*

- (1) *The spectrum X is E -nilpotent complete if and only if νX is τ -complete.*
- (2) *The natural comparison map from Construction 4.70*

$$\sigma(\nu X) \longrightarrow \text{Déc}(\text{ASS}_E(X))$$

is completion (a.k.a. τ -completion) of filtered spectra, i.e., it is an isomorphism on associated graded and the target is complete (a.k.a. τ -complete).

- (3) *The above comparison map is an isomorphism of filtered spectra if X is E -nilpotent complete. If Syn_E is cellular, then the converse is true, i.e., it is an isomorphism if and only if X is E -nilpotent complete.*

Proof. Because σ preserves limits, the map from Construction 4.70 is obtained by applying σ to the natural map

$$\nu X \longrightarrow \text{Tot}(\nu(E^{[\bullet]} \otimes X)). \quad (4.72)$$

We begin by showing that this map is τ -completion, i.e., that it is an isomorphism after tensoring with $C\tau$ and that the target is τ -complete. By Corollary 4.69, the target is a limit of τ -complete objects, and is therefore τ -complete. Because $C\tau$ is finite, tensoring with it preserves limits, so we find that tensoring (4.72) with $C\tau$ results in a map in $\mathcal{D}(\text{grComod}_{E_*E}) \subseteq \text{Mod}_{C\tau}(\text{Syn}_E)$ of the form

$$E_*X \longrightarrow \text{Tot}((E_*E)^{\otimes[\bullet]} \otimes E_*X),$$

where the tensor products are over E_* . As E_*E is flat, we may take these tensor products to be underived. This map is an isomorphism in $\mathcal{D}(\text{grComod}_{E_*E})$, because for any $M \in \text{grComod}_{E_*E}$, the cobar complex

$$(E_*E)^{\otimes[\bullet]} \otimes M$$

constitutes a cosimplicial resolution of M by relative injectives, so that the map from M into its totalisation is a quasi-isomorphism. We conclude that (4.72) is indeed τ -completion in Syn_E .

Next, if we invert τ on (4.72), then by Proposition 4.59 we obtain the map of spectra

$$X \longrightarrow \text{Tot}(E^{[\bullet]} \otimes X).$$

This map is, by definition, an isomorphism if and only if X is E -nilpotent complete. Because any map of synthetic spectra is an isomorphism if and only if it so after inverting τ and after quotienting by τ (see Proposition 3.79), we find that (4.72) is an isomorphism if and only if X is E -nilpotent complete. As (4.72) is τ -completion, this proves item (1).

Item (2) follows immediately from the fact that σ preserves τ -completion; see Theorem 3.88 (2). If X is E -nilpotent complete, then (4.72) is an isomorphism, and hence so is the map from Construction 4.70. Finally, cellularity of Syn_E is equivalent to σ being conservative (Proposition 4.31), thereby showing the final claim of item (3). ■

Remark 4.73. The reason why completion appears in the comparison between $\sigma(\nu X)$ and the Adams filtration is due to our use of cosimplicial objects in defining the latter. Note that the colimit of the filtration $\sigma(\nu X)$ is always naturally isomorphic to X , but that this filtration may not be complete. On the other hand, the cosimplicial definition of $\text{ASS}_E(X)$ is always a complete filtration, but its colimit may not be X . These convergence problems are in fact the same, since the natural map from the former to the (d ecalage of the) latter is completion. We expect that $\sigma(\nu X)$ is isomorphic, as a filtered spectrum, to the d ecalage of a filtered Adams resolution of X (see Remark 2.90), so that Remark 2.90 would also explain this difference in packaging of the convergence problem. (On the other hand, incorporating monoidal structures as in the next remark would require working from the second page onward, for reasons explained by Remark 2.84.)

Remark 4.74 (Monoidal version). The left-hand side of the comparison map of Construction 4.70 is naturally a lax symmetric monoidal functor in X ; see Remark 4.62. Recall from Remark 2.79 that D ec^Δ is naturally lax symmetric monoidal. If E carries considerably more structure, then we can give the functor $X \mapsto E^{[\bullet]} \otimes X$ a lax symmetric monoidal structure as well, in which case the comparison map

$$\sigma(\nu X) \longrightarrow \text{D ec}^\Delta(E^{[\bullet]} \otimes X)$$

of Construction 4.70 matches up these monoidal structures, as follows.

- ◆ If E carries an \mathbf{E}_1 -structure, then the semicosimplicial spectrum $E^{[\bullet]}$ naturally extends to a cosimplicial spectrum $\Delta \rightarrow \text{Sp}$; see [MNN17, Construction 2.7]. If E carries an \mathbf{E}_n -structure for $1 \leq n \leq \infty$, then using Dunn additivity, this construction turns $E^{[\bullet]}$ into a cosimplicial \mathbf{E}_{n-1} -ring.
- ◆ Consequently, if E is \mathbf{E}_n for $1 \leq n \leq \infty$, then the functor

$$\text{Sp} \longrightarrow \text{Sp}^\Delta, \quad X \longmapsto E^{[\bullet]} \otimes X$$

is naturally a lax \mathbf{E}_{n-1} -monoidal functor. Postcomposing with the lax symmetric monoidal functor D ec^Δ , we obtain a lax \mathbf{E}_{n-1} -monoidal functor

$$\text{Sp} \longrightarrow \text{FilSp}, \quad X \longmapsto \text{D ec}^\Delta(E^{[\bullet]} \otimes X).$$

In this case, the comparison map of Construction 4.70 is an \mathbf{E}_{n-1} -monoidal natural transformation.

Said differently, by completing the filtered spectrum $\sigma(\nu_E(-))$, through Theorem 4.71 we obtain a lax symmetric monoidal structure on $\text{D}\acute{\text{e}}\text{c}(\text{ASS}_E(-))$, even if E is merely homotopy-associative. Only when E is \mathbf{E}_∞ can we identify this structure concretely in terms of E ; in general, we have to work at the level of synthetic spectra.

Remark 4.75. A slightly different way of showing that νX is τ -complete if and only if X is E -nilpotent complete is given by [BHS23, Proposition A.13]. They also show that this is equivalent to νX being νE -nilpotent complete. Note that this also follows from our proof: using Example 4.10, we see that (4.72) is precisely the νE -nilpotent completion of νX .

Remark 4.76. One could try to run the same argument for the Adams spectral sequence for maps $[Y, X]_*$ for a general spectrum Y . Due to the nature of the definition of Syn_E as in [Pst22], this is only sensible when Y is a filtered colimit of finite E -projective spectra. (The above proof breaks down in general because when taking (4.72) mod τ , we obtain a resolution only by *relative injectives*. This is also related to Syn_E being generated by the synthetic analogues of finite E -projectives; see Proposition 4.8 (4).) If $\pi_* E$ is a graded field, then this condition on Y is vacuous, but not in general. To obtain the Adams spectral sequence for general E and Y , one has to work with a different version of synthetic spectra, as in the work by Patchkoria–Pstragowski [PP23]. In Theorem 5.60 of op. cit., when working in this different version, they identify the spectral sequence underlying the filtered mapping spectrum (see Construction 3.81) from νY to νX with the Adams spectral sequence for $[Y, X]_*$.

4.5 The synthetic Omnibus Theorem

Previously in Section 4.3, we explained how the filtered Omnibus Theorem 3.62, as well as its truncated variants of Theorems 3.67 and 3.70, directly imply a synthetic version, up to a reindexing. For the convenience of the reader, we state this synthetic version here. We leave the analogous re-indexing of the truncated versions to the reader, or refer to [CDvN24, Theorems 2.21 and 2.28] for a recorded version.

Theorem 4.77 (Synthetic Omnibus). *Let X be a τ -complete synthetic spectrum, and assume that in its signature spectral sequence, we have $\text{RE}_\infty^{*,*} = 0$ (for instance, this happens if the spectral sequence converges strongly). Let $x \in E_2^{n,s} = \pi_{n,s}(X/\tau)$ be a nonzero class. Then the following are equivalent.*

- (1a) *The element x is a permanent cycle.*
- (1b) *The element $x \in \pi_{n,s}(X/\tau)$ lifts to an element of $\pi_{n,s} X$.*

For any such lift α to $\pi_{n,s} X$, the following are true.

- (2a) *If x survives to page r , then $\tau^{r-2} \cdot \alpha$ is nonzero.*
- (2b) *If x survives to page ∞ , then α maps to a nonzero element in $\pi_n X^{\tau=1}$, and this element is detected by x .*

Moreover, if x lifts to X , then there exists a lift α with either of the following additional properties.

- (3a) If x is the target of a d_r -differential, then $\tau^{r-1} \cdot \alpha = 0$.
- (3b) If $\theta \in \pi_n X^{\tau=1}$ is detected by x , then α is sent to θ under $\pi_{n,s} X \rightarrow \pi_n X^{\tau=1}$.

Finally, we have the following generation statement.

- (4) Let $\{\alpha_i\}$ be a collection of elements of $\pi_{n,*} X$ such that their mod τ reductions generate the permanent cycles in stem n . Then the τ -adic completion of the $\mathbf{Z}[\tau]$ -submodule of $\pi_{n,*} X$ generated by the $\{\alpha_i\}$ is equal to $\pi_{n,*} X$.

Proof. This follows by combining Theorem 3.62 with Proposition 4.34 and Theorem 3.88. ■

If the synthetic spectrum X in Theorem 4.77 is the synthetic analogue of a spectrum Y , then Theorem 4.71 tells us that the signature of νY is the E -Adams spectral sequence for Y . Moreover, in this case the convergence condition is precisely asking for the Adams spectral sequence for Y to be strongly convergent. (Indeed, because this filtration is left-concentrated in the sense of Definition 2.53, this follows from Theorem 2.55.) In this way, the filtered Omnibus Theorem together with the computation of $\sigma \circ \nu$ recovers the Omnibus Theorem for synthetic analogues of Burklund–Hahn–Senger [BHS23, Theorem 9.19].

This combination of Theorems 4.71 and 4.77 lets information flow both ways. On the one hand, we now see that if we use synthetic spectra to compute the homotopy of a synthetic analogue, this gives us new information about the underlying Adams spectral sequence. On the other hand, we can also use this to import existing knowledge about Adams spectral sequence into synthetic spectra, thereby giving us a starting point for new computations.

Remark 4.78 (Comparison of proofs). Our proof of the Omnibus Theorem is inspired by the one of Burklund–Hahn–Senger in [BHS23, Appendix A]. They identify the νE -Adams spectral sequence of νX with (to use our terminology) the τ -BSS for the E -ASS of X ; see Theorem A.8 of op. cit. for the precise meaning of this.^[4] Generalising this to an arbitrary synthetic spectrum requires finding suitable replacements for these three spectral sequences. In our approach, we view the signature spectral sequence as the appropriate replacement for the E -ASS, and we do away with the νE -ASS, going straight to the τ -BSS of the signature spectral sequence. As the signature spectral sequence is defined using filtered spectra, the proof naturally takes place there, so that the synthetic version is a special case of the filtered version of Section 3.5. The computation of $\sigma \circ \nu$ is then the result that gives this a concrete meaning.

As a consequence of Theorem 4.71, we learn that the strict filtration on $\pi_n X$ induced by

^[4]Note that they use second-page indexing, which can be obtained from ours using Remark 3.55. See also Variant 4.39.

$\pi_{n,*} \nu X$ coincides with the E -Adams filtration, recovering [BHS23, Corollary 9.21]. See Definition 2.101 for the definition of the (algebraic) Adams filtration.

Corollary 4.79 (Geometric Adams filtration). *Let X be a spectrum, and let n be an integer. Let $f: \mathbf{S}^n \rightarrow X$ be a map, and let $s \geq 0$. Then f has (algebraic) E -Adams filtration (see Definition 2.101) at least s if and only if there exists a factorisation*

$$\begin{array}{ccc} \nu \mathbf{S}^n = \mathbf{S}^{n,0} & \xrightarrow{\nu f} & \nu X. \\ \tau^s \downarrow & \nearrow \text{---} & \\ \Sigma^{0,s} \nu \mathbf{S}^n = \mathbf{S}^{n,s} & & \end{array}$$

In other words, f is of (algebraic) E -Adams filtration at least s if and only if νf is divisible by τ^s .

In particular, the (algebraic) E -Adams filtration on $\pi_n X$ coincides with

$$F^s \pi_n X = \text{im} \left([\mathbf{S}^{n,s}, \nu X] \xrightarrow{\tau=1} [\mathbf{S}^n, X] \right).$$

Proof. Combine Propositions 2.102 and 2.104 with Theorem 4.71. ■

Phrased differently, the functor ν sends a map $f: \mathbf{S}^n \rightarrow X$ of spectra to the map νf , which we think of as f placed in Adams filtration 0. Then a τ^s -division of νf (should it exist) is a witness that f has Adams filtration at least s .

More generally, the analogous result holds when \mathbf{S}^n is replaced by a spectrum Y that can be written as a filtered colimit of finite E -projective spectra. For a general spectrum Y , one has to work with the different construction of synthetic spectra from [PP23], as explained in Remark 4.76; the general result in this case is [PP23, Theorem 5.60 (2)].

Notation 4.80. If Y is a spectrum and $\alpha \in \pi_n Y$ is an element, then there are two conventions one can take for naming elements of $\pi_{*,*} \nu Y$.

- (1) One could denote $\nu(\alpha) \in \pi_{n,0} \nu Y$ by the same symbol α again, and write $\tilde{\alpha}$ for (a choice of) the maximal τ -division thereof. This is the convention used by Burklund–Hahn–Senger in [BHS23].
- (2) One could denote (a choice of) the maximal τ -division of $\nu(\alpha)$ by the same symbol α . This is used, e.g., by Burklund in [Bur21].

The second convention is the one we will use. Conceptually, it emphasises the Adams filtration of an element in $\pi_n Y$ as an intrinsic invariant, so that we should be thinking of its maximal τ -division as its ‘true origin’.

Remark 4.81 (Uniqueness). Note that the maximal τ -division of $\nu(\alpha)$ may not be uniquely defined, as $\pi_{n,*} \nu Y$ might contain τ -power torsion. Specifically, if α has filtration s , then the maximal τ -division of $\nu(\alpha)$ is defined up to τ^s -power torsion in $\pi_{n,s} \nu Y$. Via the

Omnibus Theorem (if the convergence conditions are met), this is more or less saying that it is defined up to permanent cycles in filtrations $k \geq s$ that are hit by a differential of length $\leq k - 1$.

Example 4.82. Consider the 1-stem of the MU-synthetic sphere. In the ordinary sphere, the element $\eta \in \pi_1 \mathbf{S}$ has Adams–Novikov filtration 1. Translated in synthetic terms, we have $\pi_{1,s} C\tau = 0$ for $s \neq 1$, and $\pi_{1,1} C\tau$ is isomorphic to $\mathbf{Z}/2$ with generator h_0 . Since the 1-stem does not receive or support any differentials (for degree reasons), it follows from Corollary 3.64 that $\pi_{1,*} \mathbf{S}$ is τ -power torsion free. It follows that τ -inversion is an injection, and as $\pi_1 \mathbf{S} \cong \mathbf{Z}/2 \cdot \eta$, we obtain an isomorphism

$$\pi_{1,*} \mathbf{S} \cong \mathbf{Z}/2[\tau] \cdot x,$$

where x is in bidegree $(1, 1)$, where x maps to h_0 under mod τ reduction, and where $\tau \cdot x$ is the synthetic analogue of the map $\eta \in \pi_1 \mathbf{S}$.

The second convention above would be to use the symbol η to denote the generator x . ▲

As an another example, in Example 3.12, we studied essentially $\pi_{n,s} \mathbf{S}$ of the (2-local) MU-synthetic sphere for $n \leq 3$, and also followed the second notational convention mentioned above.

Chapter 5

Variants of synthetic spectra

In this chapter, we discuss certain variants and modifications of synthetic spectra. The first main result is that, under a mild condition, the ∞ -category of synthetic spectra is equivalent to modules in FilSp over a filtered ring spectrum; see Section 5.2. The second is that MU-synthetic spectra are equivalent to (a small subcategory of) \mathbf{C} -motivic spectra; see Section 5.4. The first of these results requires a discussion of cellularity, and the second requires a discussion of evenness; these are discussed in Section 5.1 and Section 5.3, respectively.

Most of the results in this section are not new, but are presented slightly differently compared to the literature. In particular, our discussion of filtered models is once again focussed on the signature adjunction. One minor new result is that our discussion of evenness shows that even synthetic spectra are closed under limits (Corollary 5.35), but our approach has the downside of only applying to the cellular case. Our discussion of the relation to motivic spectra is no more than a quick review, and will not be used heavily.

5.1 Cellularity

Definition 5.1. Let \mathcal{C} be a monoidal deformation. The **cellular subcategory** $\mathcal{C}^{\mathrm{cell}}$ of \mathcal{C} is the smallest subcategory closed under colimits that contains $\rho(\mathbf{S}^{n,s})$ for all n and s . We say that \mathcal{C} is **cellular** if $\mathcal{C}^{\mathrm{cell}} = \mathcal{C}$.

Remark 5.2. A different way of stating the definition of cellularity is that the object $\mathbf{1}_{\mathcal{C}}$ of \mathcal{C} generates \mathcal{C} as a FilSp -linear category under colimits. Phrased in this way, the definition can be extended to a general deformation, where one has to provide a choice of an object to play the role of $\mathbf{1}_{\mathcal{C}}$. In this case, one might speak of the deformation being a *monogenic* FilSp -linear ∞ -category.

Proposition 5.3. *Let \mathcal{C} be a monoidal deformation.*

- (1) *The ∞ -category $\mathcal{C}^{\mathrm{cell}}$ is a presentable stable ∞ -category.*

(2) The subcategory $\mathcal{C}^{\text{cell}} \subseteq \mathcal{C}$ is closed under colimits and tensor products.

Consequently, $\mathcal{C}^{\text{cell}}$ inherits the structure of a presentably monoidal ∞ -category from \mathcal{C} , and we have a colocalisation adjunction

$$\mathcal{C}^{\text{cell}} \begin{array}{c} \xrightarrow{\quad} \\ \xleftarrow{(-)^{\text{cell}}} \end{array} \mathcal{C}$$

with a monoidal left adjoint. If \mathcal{C} is a symmetric monoidal deformation, then $\mathcal{C}^{\text{cell}}$ is a symmetric monoidal subcategory and the inclusion is a symmetric monoidal functor.

Proof. Presentability follows because it is generated under small colimits by a small set of objects. Closure under colimits is clear, and closure under tensor products follows because the tensor product of \mathcal{C} preserves colimits in each variable separately. The Adjoint Functor Theorem then yields the desired adjunction. ■

We refer to the right adjoint $\mathcal{C} \rightarrow \mathcal{C}^{\text{cell}}$ as the **cellularisation functor**, which is canonically lax monoidal (symmetric if \mathcal{C} is a symmetric deformation). It follows from this adjunction that for every $X \in \mathcal{C}$, the counit $X^{\text{cell}} \rightarrow X$ induces an isomorphism

$$\pi_{*,*} X \cong \pi_{*,*}(X^{\text{cell}}).$$

In this sense, one may think of the passage from \mathcal{C} to $\mathcal{C}^{\text{cell}}$ as similar to the passage from topological spaces up to homotopy equivalence, to topological spaces up to weak homotopy equivalence.

This is also indicative of a larger phenomenon: when using a deformation to study the resulting signature spectral sequences, we may as well work with its cellularisation.

Proposition 5.4. *Let \mathcal{C} be a (symmetric) monoidal deformation.*

- (1) The functor σ is conservative if and only if \mathcal{C} is cellular.
- (2) We have natural commutative diagrams of lax (symmetric) monoidal functors

$$\begin{array}{ccc} & \mathcal{C}^{\text{cell}} & \\ \rho \nearrow & & \searrow \\ \text{FilSp} & \xrightarrow{\quad \rho \quad} & \mathcal{C} \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{C} & \xrightarrow{\quad \sigma \quad} & \text{FilSp} \\ \downarrow (-)^{\text{cell}} & & \uparrow \sigma \\ & \mathcal{C}^{\text{cell}} & \end{array}$$

- (3) The adjunction $\rho \dashv \sigma$ restricts to an adjunction between FilSp and $\mathcal{C}^{\text{cell}}$, which we denote by

$$\text{FilSp} \begin{array}{c} \xrightarrow{\rho^{\text{cell}}} \\ \xleftarrow{\sigma^{\text{cell}}} \end{array} \mathcal{C}^{\text{cell}}.$$

Proof. The first claim is a restatement of Theorem 3.88 (4). The functor ρ lands in $\mathcal{C}^{\text{cell}}$ because it preserves colimits and because FilSp is generated under colimits by the bigraded spheres. To see that σ factors over cellularisation, we have to show that σ sends the counit $X^{\text{cell}} \rightarrow X$ to an isomorphism for every $X \in \mathcal{C}$. By adjunction, this follows from ρ landing in the cellular subcategory. The final claim follows immediately from this. ■

Proposition 5.5. *Let \mathcal{C} be a monoidal deformation. Then \mathcal{C} is cellular if and only if the generic fibre of \mathcal{C} is generated under colimits by $\rho(\mathbf{S}^{n,0}[\tau^{-1}])$ for $n \in \mathbf{Z}$ and the special fibre is generated under colimits by $\rho(C\tau \otimes \mathbf{S}^{n,s})$ for $n, s \in \mathbf{Z}$.*

Proof. The deformation \mathcal{C} is cellular if and only if the counit $X^{\text{cell}} \rightarrow X$ is an isomorphism for all $X \in \mathcal{C}$. We may check whether this map is an isomorphism after inverting and after modding out by τ ; see Proposition 3.79. One checks that the cellularisation adjunction reduces to the analogous adjunction for generation by the τ -invertible or the mod τ bigraded spheres, respectively. Since the filtered spheres $\mathbf{S}^{n,s}$ and $\mathbf{S}^{n,t}$ become isomorphic upon τ -inversion, the case of the generic fibre reduces to the stated claim. ■

Next, we turn to the case of synthetic spectra. To the author's knowledge, it is not known if for every E of Adams type, the deformation Syn_E is cellular. It is known to hold in many cases. In the case $E = \mathbf{F}_p$, it is relatively easy to see that it is cellular; see [CD24, Lemma 2.3]. Pstragowski proved in his foundational work that MU-synthetic spectra are cellular; see [Pst22, Theorem 6.2]. Later, Lawson generalised this to arbitrary connective E .

Theorem 5.6 (Lawson [Law24b]). *Let E be a homotopy-associative ring spectrum of Adams type. If E is connective, then Syn_E is cellular.*

In joint work with Barkan, we prove cellularity in an important nonconnective case.

Theorem 5.7 ([BvN25]). *Let E be a Morava E -theory at an arbitrary height and prime. Then Syn_E is cellular.*

Note that Theorem 5.6 does not imply Theorem 5.7. Indeed, if E denotes Morava E -theory, then (except at height 0) we have that \mathbf{F}_p is E -acyclic for all p , while if R is a connective homotopy-ring spectrum, then there is at least one p for which \mathbf{F}_p is not R -acyclic. It follows that Syn_E is not equivalent to Syn_R for any connective homotopy-ring spectrum R .

5.2 Filtered models

Under certain conditions, there is a description of Syn_E as modules in FilSp over a certain filtered ring spectrum. These results have appeared before [BHS22, Appendix C] [Law24b] [Pst25, Sections 3.3 and 3.5]. We include a discussion here to highlight its connection to the signature adjunction. We learned this approach from Shaul Barkan. See also [BHS22, Appendix A].

Theorem 5.8. *Let \mathcal{C} be a monoidal deformation. Suppose that \mathcal{C} satisfies the following conditions.*

- (a) *The monoidal unit $\mathbf{1}_{\mathcal{C}}$ of \mathcal{C} is compact.*
- (b) *The deformation \mathcal{C} is cellular.*

Then the adjunction

$$\mathrm{FilSp} \begin{array}{c} \xrightarrow{\rho} \\ \xleftarrow{\sigma} \end{array} \mathcal{C}$$

is a monadic adjunction that identifies the ∞ -category underlying \mathcal{C} with $\mathrm{Mod}_{\sigma(\mathbf{1}_{\mathcal{C}})}(\mathrm{FilSp})$. If \mathcal{C} is a symmetric monoidal deformation, then $\sigma(\mathbf{1}_{\mathcal{C}})$ is a filtered \mathbf{E}_{∞} -ring, and the identification of \mathcal{C} with $\mathrm{Mod}_{\sigma(\mathbf{1}_{\mathcal{C}})}(\mathrm{FilSp})$ is naturally one of symmetric monoidal ∞ -categories.

Proof. It is enough to check the three conditions of [MNN17, Proposition 5.29]. The second, that σ preserves colimits, follows from Theorem 3.88 (3). The first, that the adjunction satisfies the projection formula, follows from σ preserving colimits and Lemma 3.87. Finally, the third condition, that σ is conservative, follows from Theorem 3.88 (4). ■

Remark 5.9. Alternatively, as explained by Pstragowski, the hypotheses on \mathcal{C} can be interpreted as saying that it has a single compact filtered generator, so that a filtered version of Schwede–Shipley applies; see [Pst22, Proposition 3.16].

Specialising to the synthetic setting, we immediately obtain the following.

Corollary 5.10. *Let E be a homotopy-associative ring spectrum of Adams type. Then the adjunction*

$$\mathrm{FilSp} \begin{array}{c} \xrightarrow{\rho^{\mathrm{cell}}} \\ \xleftarrow{\sigma^{\mathrm{cell}}} \end{array} \mathrm{Syn}_E^{\mathrm{cell}}$$

is a monadic adjunction that identifies $\mathrm{Syn}_E^{\mathrm{cell}}$ with $\mathrm{Mod}_{\sigma(\nu\mathbf{S})}(\mathrm{FilSp})$ as a symmetric monoidal ∞ -category. Under this equivalence, the τ -completion of $(\nu X)^{\mathrm{cell}}$ (where X is a spectrum) is sent to $\mathrm{Tot}(\mathrm{Wh}(E^{[\bullet]} \otimes X))$.

Proof. This follows by combining Theorem 5.8 and Theorem 4.71. ■

Note, however, that this does not yet give an alternative description of Syn_E (even after forgetting the monoidal structure): we have not yet identified the filtered spectrum $\sigma(\nu\mathbf{S})$ as a filtered \mathbf{E}_1 -ring (let alone as a filtered \mathbf{E}_{∞} -ring) without reference to Syn_E . We can do this using Theorem 4.71, but only when E is sufficiently structured; see Remark 4.74. However, this result only identifies the *completion* of $\sigma(\nu\mathbf{S})$; by the same theorem, $\sigma(\nu\mathbf{S})$ is complete if and only if \mathbf{S} is E -nilpotent complete.

The prime example of such a setting is $E = \mathrm{MU}$, which is an \mathbf{E}_{∞} -ring, and the sphere is MU -nilpotent complete (see Example 2.111). As mentioned before, $\mathrm{Syn}_{\mathrm{MU}}$ is cellular, resulting in the following.

Corollary 5.11. *There is a symmetric monoidal equivalence*

$$\mathrm{Syn}_{\mathrm{MU}} \simeq \mathrm{Mod}_{\mathrm{Tot}(\mathrm{Wh}(\mathrm{MU}^{[\bullet]}))}(\mathrm{FilSp})$$

under which σ becomes the forgetful functor (as a lax symmetric monoidal functor). Moreover, on MU-nilpotent complete spectra, the functor v is identified with the functor sending X to $\mathrm{Tot}(\mathrm{Wh}(\mathrm{MU}^{[\bullet]} \otimes X))$ (as a lax symmetric monoidal functor).

Remark 5.12 ([BHS22], Appendix C). Because we can identify the τ -completion of $\sigma(v\mathbf{S})$ by Theorem 4.71, we can ignore the nilpotent-completeness issues by passing to a slight modification of synthetic spectra. Namely, if E is an \mathbf{E}_1 -ring of Adams type, then the functor σ also induces an equivalence

$$\mathrm{Mod}_{(v\mathbf{S})_{\hat{\tau}}}(\mathrm{Syn}_E^{\mathrm{cell}}) \xrightarrow{\simeq} \mathrm{Mod}_{\mathrm{Tot}(\mathrm{Wh}(E^{[\bullet]}))}(\mathrm{FilSp}),$$

which is symmetric monoidal if E is an \mathbf{E}_∞ -ring. This follows by applying Theorem 5.8 to the symmetric monoidal deformation $\mathrm{Mod}_{(v\mathbf{S})_{\hat{\tau}}}(\mathrm{Syn}_E)$ and using that σ preserves τ -completion (Theorem 3.88 (2)). This is the formulation of the result given by Burklund–Hahn–Senger in [BHS22, Proposition C.22]. (Of course, the same reasoning holds in the context of Theorem 5.8, yielding a filtered model for modules over the τ -completion of the unit of \mathcal{C} .)

A more elementary but curious-looking example is the case where $E = \mathbf{S}$. In this case, all nilpotent completeness conditions become vacuous, and we learn the following.

Corollary 5.13. *There is a symmetric monoidal equivalence*

$$\mathrm{Syn}_{\mathbf{S}} \simeq \mathrm{Mod}_{\mathrm{Wh}\mathbf{S}}(\mathrm{FilSp})$$

under which σ becomes the forgetful functor and under which v becomes the Whitehead filtration functor (and these identifications are as lax symmetric monoidal functors), and under which the homological t -structure is identified with the diagonal t -structure.

One can check that for every Adams type E , the functor $\sigma: \mathrm{Syn}_E \rightarrow \mathrm{FilSp}$ factors through modules over $\mathrm{Wh}\mathbf{S}$, so that this is in a sense the minimal structure present on E -synthetic spectra as E varies.

5.3 Evenness

The goal of this section is to discuss *even synthetic spectra* introduced in [Pst22, Section 5.2], also studied in a slightly different form in [Pst25]. For spectral sequences, *evenness* is nothing more than the condition that the starting page is concentrated in certain even degrees (where the specific meaning of this depends on the indexing system being used). On the synthetic side, even objects turn out to have geometric interpretations: the special fibre of even MU-synthetic is the derived ∞ -category of quasi-coherent sheaves over $\mathfrak{M}_{\mathrm{fg}}$,

see Example 5.37. Moreover, in the next section, we will see that even MU-synthetic spectra can be viewed as *motivic spectra*.

We will reinterpret the notion of evenness defined by Pstragowski through the language of deformations. Accordingly, we begin by discussing this in the setting of filtered spectra.

5.3.1 Even filtered spectra

We remind the reader that we use first-page indexing on filtered spectra.

Lemma 5.14. *Let X be a filtered spectrum. Then the following conditions are equivalent.*

- (a) *For every s , the transition map*

$$X^{2s} \longrightarrow X^{2s-1}$$

is an isomorphism. In other words, for every n and s , the map

$$\tau: \pi_{n,2s} X \longrightarrow \pi_{n,2s-1} X$$

is an isomorphism.

- (b) *For every odd integer s , the associated graded spectrum $\mathrm{Gr}^s X$ vanishes. In other words, the homotopy groups $\pi_{n,s}(X/\tau)$ vanish whenever s is odd.*
- (c) *The filtered spectrum X belongs to the smallest subcategory of FilSp generated under colimits by the filtered spheres $\mathbf{S}^{n,2s}$ for all n and s .*

Proof. By exactness, it follows that (a) is equivalent to (b). We prove that (a) is equivalent to (c). Write $\mathrm{FilSp}^{\mathrm{ev}}$ for the full subcategory of those filtered spectra X satisfying (a), and write \mathcal{C} for the full subcategory of FilSp generated by $\mathbf{S}^{n,2s}$ for all n and s . Clearly $\mathrm{FilSp}^{\mathrm{ev}}$ is closed under colimits, so it follows that $\mathcal{C} \subseteq \mathrm{FilSp}^{\mathrm{ev}}$. To prove the other inclusion, by [Yan22, Corollary 2.5], it is enough to show that the functors $\pi_{n,2s}$ jointly detect isomorphisms on $\mathrm{FilSp}^{\mathrm{ev}}$. This follows from the fact that the functors $\pi_{n,s}$ jointly detect isomorphisms on FilSp , and that when restricted to $\mathrm{FilSp}^{\mathrm{ev}}$, the functors $\pi_{n,2s}$ and $\pi_{n,2s-1}$ are naturally isomorphic (induced by τ). ■

Definition 5.15. We say that a filtered spectrum is **even** if it satisfies the equivalent conditions of Lemma 5.14. We write $\mathrm{FilSp}^{\mathrm{ev}}$ for the full subcategory of FilSp on the even filtered spectra.

Evenness is a notion that is detectable on the spectral sequence: in first-page indexing on filtered spectra, it says that the first page vanishes in odd filtrations.

Example 5.16.

- (1) A filtered sphere $\mathbf{S}^{n,s}$ is even if and only if s is even. We may refer to a sphere of this form as an *even filtered sphere*; note that this places no restrictions on the variable n .

- (2) Let X be a spectrum. Then its Whitehead filtration $\text{Wh } X$ is even if and only if $\pi_* X$ is concentrated in even degrees, and likewise for $\text{Post } X$. \blacktriangle

Proposition 5.17.

- (1) The ∞ -category FilSp^{ev} is a presentable stable ∞ -category.
 (2) The subcategory $\text{FilSp}^{\text{ev}} \subseteq \text{FilSp}$ is closed under limits, colimits, and tensor products.

Consequently, FilSp^{ev} inherits a presentably symmetric monoidal structure from FilSp , and the inclusion functor admits both a left and a right adjoint and is a symmetric monoidal functor.

Proof. Presentability follows from characterisation (c) of Lemma 5.14. Closure under limits and colimits follows from characterisation (a), which in particular implies it is a stable subcategory. Closure under tensor products follows from characterisation (b), using that $\text{Gr}: \text{FilSp} \rightarrow \text{grSp}$ is symmetric monoidal (see Remark 3.26), and that $\text{grSp}^{\text{ev}} \subseteq \text{grSp}$ is obviously closed under tensor products. The inclusion admits a right adjoint because it preserves colimits. Because the even filtered spheres are compact, the inclusion also preserves compact objects, and hence admits a left adjoint. \blacksquare

Remark 5.18. One may check that the left adjoint to the inclusion $\text{FilSp}^{\text{ev}} \subseteq \text{FilSp}$ sends a filtered spectrum X to

$$\dots \longrightarrow X^1 \rightrightarrows X^1 \longrightarrow X^{-1} \rightrightarrows X^{-1} \longrightarrow X^{-3} \rightrightarrows \dots$$

with X^{-1} in filtrations 0 and -1 , and the right adjoint sends X to

$$\dots \longrightarrow X^2 \rightrightarrows X^2 \longrightarrow X^0 \rightrightarrows X^0 \longrightarrow X^{-2} \rightrightarrows \dots$$

with X^0 in filtrations 0 and -1 .

Remark 5.19. There is a symmetric monoidal equivalence $p: \text{FilSp}^{\text{ev}} \xrightarrow{\simeq} \text{FilSp}$ that ‘pinches’ the transition maps together: it sends an even filtered spectrum X to

$$\dots \longrightarrow X^2 \longrightarrow X^0 \longrightarrow X^{-2} \longrightarrow \dots$$

For instance, this sends $\mathbf{S}^{n,2s}$ to $\mathbf{S}^{n,s}$. Formally, precomposition with $2: \mathbf{Z}^{\text{op}} \rightarrow \mathbf{Z}^{\text{op}}$ results in a symmetric monoidal functor $\text{FilSp} \rightarrow \text{FilSp}$. Restricting the domain to the even subcategory results in the claimed symmetric monoidal equivalence p . Indeed, this restricted functor has a two-sided inverse given by d^* , where $d: \mathbf{Z}^{\text{op}} \rightarrow \mathbf{Z}^{\text{op}}$ is the map of posets that sends $2s$ and $2s - 1$ to s .

Remark 5.20. We can regard FilSp^{ev} as a symmetric monoidal deformation in its own right, with deformation parameter τ^2 . Formally, we have a symmetric monoidal functor $2: \mathbf{Z} \rightarrow \mathbf{Z}$. Writing $i: \mathbf{Z} \rightarrow \text{FilSp}$ for the functor from Definition 2.17, we see that

$i \circ 2: \mathbf{Z} \rightarrow \text{FilSp}$ is a symmetric monoidal functor that lands in the subcategory FilSp^{ev} . The underlying functor can be depicted by the diagram

$$\dots \xrightarrow{\tau^2} \mathbf{S}^{0,-2} \xrightarrow{\tau^2} \mathbf{S} \xrightarrow{\tau^2} \mathbf{S}^{0,2} \xrightarrow{\tau^2} \dots$$

Via Notation 3.83, this gives FilSp^{ev} the structure of a symmetric monoidal deformation. The resulting right adjoint $\text{FilSp}^{\text{ev}} \rightarrow \text{FilSp}$ is in fact the functor p from Remark 5.19.

5.3.2 Even deformations

We focus on the case of monoidal deformations. We have a choice of using either characterisation (b) or (c) from Lemma 5.14 as the definition of evenness in a monoidal deformation. In general these notions differ, but we will show that they agree under suitable hypotheses. It is an arbitrary choice which of the two characterisations we regard as the definition.

Definition 5.21. Let \mathcal{C} be a monoidal deformation. The **even subcategory** \mathcal{C}^{ev} of \mathcal{C} is the smallest subcategory closed under colimits that contains $\rho(\mathbf{S}^{n,2s})$ for all n and s . We say that an object $X \in \mathcal{C}$ is **even** if it belongs to \mathcal{C}^{ev} .

Because our definition uses only the filtered spheres, it is only applicable in the cellular setting. More precisely, the above definition implies that every even object is cellular.

Proposition 5.22. *Let \mathcal{C} be a monoidal deformation.*

- (1) *The ∞ -category \mathcal{C}^{ev} is a presentable stable ∞ -category.*
- (2) *The subcategory $\mathcal{C}^{\text{ev}} \subseteq \mathcal{C}^{\text{cell}}$ is closed under colimits and tensor products.*

Only under the following hypotheses on \mathcal{C} do we obtain a usable notion of evenness.

Proposition 5.23. *Let \mathcal{C} be a cellular monoidal deformation with a compact unit. Suppose that the filtered spectrum $\sigma(\mathbf{1}_{\mathcal{C}})$ is even. Then for every $X \in \mathcal{C}$, the following are equivalent.*

- (a) *The object X belongs to \mathcal{C}^{ev} .*
- (b) *The homotopy groups $\pi_{n,s}(X/\tau)$ vanish whenever s is odd.*
- (c) *The filtered spectrum σX is even.*

Moreover, the adjunction $\rho \dashv \sigma$ restricts to an adjunction between FilSp^{ev} and \mathcal{C}^{ev} , which we denote by

$$\text{FilSp}^{\text{ev}} \begin{array}{c} \xrightarrow{\rho^{\text{ev}}} \\ \xleftarrow{\sigma^{\text{ev}}} \end{array} \mathcal{C}^{\text{ev}}.$$

Proof. Clearly (b) is equivalent to (c), because this is the case for filtered spectra by Lemma 5.14. Using Lemma 3.87 (1), for all n and s we have an isomorphism

$$\sigma(\rho(\mathbf{S}^{n,s})) \cong \Sigma^{n,s} \sigma(\mathbf{1}_{\mathcal{C}}).$$

It follows that if $\sigma(\mathbf{1}_{\mathcal{C}})$ is even, then $\sigma(\rho(\mathbf{S}^{n,2s}))$ is even for all n and s . Because the unit of \mathcal{C} is compact, the functor $\sigma: \mathcal{C} \rightarrow \text{FilSp}$ preserves colimits; see Theorem 3.88 (3). Using this, we see that (a) implies (c). To see that (c) implies (a), one either reasons as in the proof of Lemma 5.14, or one deduces it from the filtered case using Theorem 5.8. Finally, these conditions show that when restricting σ to \mathcal{C}^{ev} , it lands in FilSp^{ev} , so that the adjunction restricts as indicated. ■

In this case, evenness is preserved under limits too, which is not obvious from the definition of \mathcal{C}^{ev} from Definition 5.21.

Corollary 5.24. *Let \mathcal{C} be a monoidal deformation satisfying the conditions of Proposition 5.23. Then the subcategory $\mathcal{C}^{\text{ev}} \subseteq \mathcal{C}$ is closed under limits.*

Proof. Because σ preserves limits, this follows from Proposition 5.17 and from characterisation (c) in Proposition 5.23. ■

Remark 5.25. The filtered model of Theorem 5.8 can be adapted to the even case as well. If the conditions of Proposition 5.23 hold, then \mathcal{C} is equivalent to modules over $\sigma(\mathbf{1}_{\mathcal{C}})$ by Theorem 5.8, which by assumption is an \mathbf{E}_{∞} -algebra in FilSp^{ev} . As FilSp^{ev} is a monoidal subcategory of FilSp , it follows that σ induces an equivalence of ∞ -categories

$$\mathcal{C}^{\text{ev}} \xrightarrow{\simeq} \text{Mod}_{\sigma(\mathbf{1}_{\mathcal{C}})}(\text{FilSp}^{\text{ev}}).$$

Postcomposing this with the (symmetric monoidal) equivalence $p: \text{FilSp}^{\text{ev}} \xrightarrow{\simeq} \text{FilSp}$ from Remark 5.19, we obtain an equivalence

$$\mathcal{C}^{\text{ev}} \xrightarrow{\simeq} \text{Mod}_{p(\sigma(\mathbf{1}_{\mathcal{C}}))}(\text{FilSp}).$$

If \mathcal{C} is symmetric monoidal, then these equivalences are of symmetric monoidal ∞ -categories.

5.3.3 Even synthetic spectra

We now specialise to the case of synthetic spectra. We do not get a good notion of evenness in Syn_E for every E , because the signature of the synthetic sphere may not be even. The following condition on E will ensure this.

Definition 5.26 ([Pst22], Definition 5.8). Let E be a homotopy associative ring spectrum.

- (1) A finite spectrum P is called **finite even E -projective** if E_*P is a projective E_* -module and is concentrated in even degrees.
- (2) We say that E is of **even Adams type** if it is of Adams type (see Definition 2.91) and can be written as a filtered colimit of finite even E -projective spectra.

If E is of even Adams type, then it follows that E_*E is also concentrated in even degrees. Beware that asking E to be of even Adams type is stronger than asking for E_* to be

concentrated in even degrees and for E to be of Adams type, as the following example illustrates.

Example 5.27.

- (1) The sphere spectrum \mathbf{S} is of even Adams type.
- (2) The ring spectrum \mathbf{MU} is of even Adams type; see [Pst22, Example 5.9].
- (3) The ring spectrum \mathbf{F}_p is not of even Adams type, even though $\pi_*\mathbf{F}_p$ is even and \mathbf{F}_p is of Adams type. Indeed, $\pi_*(\mathbf{F}_p \otimes \mathbf{F}_p)$ is the (p -primary) dual Steenrod algebra, which is not concentrated in even degrees.
- (4) On the other hand, if E is Landweber exact and E_* is concentrated in even degrees, then E is in fact of even Adams type; see [Pst22, Example 5.9]. ▲

Lemma 5.28. *Let (A, Γ) be a graded Hopf algebroid, with A and Γ concentrated in even degrees. Write*

$$\mathrm{grComod}_{(A, \Gamma)}^{\mathrm{ev}} \quad \text{and} \quad \mathrm{grComod}_{(A, \Gamma)}^{\mathrm{odd}}$$

for the full subcategories of $\mathrm{grComod}_{(A, \Gamma)}$ on those graded comodules that are concentrated in even, respectively odd, degrees. Then we have a symmetric monoidal equivalence

$$\mathrm{grComod}_{(A, \Gamma)} \simeq \mathrm{grComod}_{(A, \Gamma)}^{\mathrm{ev}} \times \mathrm{grComod}_{(A, \Gamma)}^{\mathrm{odd}}.$$

Proof. This is immediate. ■

The characterisation of evenness looks slightly different in the synthetic setting, due to the reindexing when passing from filtered to synthetic spheres from Proposition 4.34. Namely, because of the relation

$$\rho(\mathbf{S}_{\mathrm{fil}}^{n,s}) = \mathbf{S}_{\mathrm{syn}}^{n,s-n},$$

we should think of a synthetic sphere $\mathbf{S}_{\mathrm{syn}}^{n,s}$ as being even when $n + s$ is even.

Lemma 5.29. *Let E be a homotopy-associative ring spectrum of even Adams type. Then $\sigma(\mathbf{S}_{\mathrm{syn}})$ is even.*

Proof. Recall from Proposition 4.34 that we have an isomorphism

$$\pi_{n,s}^{\mathrm{syn}}(-) \cong \pi_{n,s+n}^{\mathrm{fil}}(\sigma(-)).$$

It follows that σX is even if and only if $\pi_{n,s}^{\mathrm{syn}}(X/\tau)$ vanishes whenever $n + s$ is odd. By Example 4.56, we have an isomorphism

$$\pi_{n,s}(\mathbf{S}_{\mathrm{syn}}/\tau) \cong \mathrm{Ext}_{E_*E}^{s,n+s}(E_*, E_*) = \mathrm{Ext}_{E_*E}^s(E_*[n+s], E_*).$$

Using the splitting of $\mathrm{grComod}_{E_*E}$ from Lemma 5.28, it follows that these groups vanish whenever $n + s$ is odd, proving that $\sigma(\mathbf{S}_{\mathrm{syn}})$ is even. ■

Translated to the synthetic setting, Proposition 5.23 takes the following form.

Proposition 5.30. *Let E be a homotopy-associative ring spectrum of even Adams type. Let X be a cellular E -synthetic spectrum. Then the following are equivalent.*

- (a) *The synthetic spectrum X belongs to the smallest subcategory of $\text{Syn}_E^{\text{cell}}$ generated under colimits by the synthetic spheres $\mathbf{S}^{n,s}$ where $n + s$ is even.*
- (b) *The homotopy groups $\pi_{n,s}(X/\tau)$ vanish whenever $n + s$ is odd; in other words, the second page of the signature spectral sequence of X is concentrated in even total degree.*
- (c) *The filtered spectrum σX is even in the sense of Definition 5.15.*

If this is the case, then we say that X is *even*. Moreover, the adjunction $\rho \dashv \sigma$ restricts to an adjunction

$$\text{FilSp}^{\text{ev}} \begin{array}{c} \xrightarrow{\rho^{\text{ev}}} \\ \xleftarrow{\sigma^{\text{ev}}} \end{array} (\text{Syn}_E^{\text{cell}})^{\text{ev}}.$$

Proof. Using Lemma 5.29, it follows that Proposition 5.23 applies. ■

Remark 5.31. The restriction to the cellular case is an artefact of our use of filtered spectra to define the notion of evenness. In the case of synthetic spectra, one can define a notion of evenness that does not require cellularity assumptions: see [Pst22, Section 5.2]. It follows from Theorem 5.13 of op. cit. that if Syn_E is cellular, then these two notions of evenness coincide. Moreover, Pstragowski proves that the homological t-structure on Syn_E restricts to a t-structure on Syn_E^{ev} whose heart is equivalent to $\text{grComod}_{E_*E}^{\text{ev}}$; see Remark 5.11 of op. cit.

Example 5.32. To see why we have to restrict to E of even Adams type, consider the case $E = \mathbf{F}_2$. Then for every $s \geq 0$, we have an isomorphism $\pi_{0,s}(C\tau) \cong \mathbf{F}_2 \cdot h_0^s$. It follows that $\sigma(\mathbf{S}_{\text{syn}})$ is not an even filtered spectrum. As a result, the subcategory of $\text{Syn}_{\mathbf{F}_2}$ generated under colimits by the spheres $\mathbf{S}^{n,s}$ for $n + s$ even is not related to FilSp^{ev} in a natural way. ▲

Corollary 5.33. *Let E be a homotopy-associative ring spectrum of even Adams type. Let X be a spectrum. Then $(\nu X)^{\text{cell}}$ is even if and only if E_*X is concentrated in even degrees.*

Proof. This follows immediately from the isomorphism from Example 4.56:

$$\pi_{n,s}(\nu X/\tau) \cong \text{Ext}_{E_*E}^{s,n+s}(E_*, E_*X) = \text{Ext}_{E_*E}^s(E_*[n+s], E_*X),$$

combined with the splitting of Lemma 5.28. ■

Example 5.34. By Corollary 5.33, the synthetic sphere $\mathbf{S}^{n,s}$ is even if and only if $n + s$ is even. In particular, the sphere $\mathbf{S}^{0,-1}$ is *not* even, and as a result, the map τ does not live in Syn_E^{ev} . However, the map $\tau^2: \mathbf{S}^{0,-2} \rightarrow \mathbf{S}$ does live in this subcategory.

If we restrict our attention to only the even synthetic spectra, then it would be more useful to use the letter τ for what we would otherwise denote by τ^2 ; in other words, to regard τ^2 as the deformation parameter. (The resulting deformation adjunction would then factor as the deformation structure on $\mathrm{FilSp}^{\mathrm{ev}}$ from Remark 5.20 followed by the adjunction $\rho^{\mathrm{ev}} \dashv \sigma^{\mathrm{ev}}$.) From our perspective, this is what happens in the literature when *motivic spectra* are used to study the stable stems; see Section 5.4, particularly Theorem 5.40, for more information. \blacktriangle

Our approach naturally gives us limit-closure of even synthetic spectra, at least up to cellularisation. Note that the limit-closure of even synthetic spectra is absent from the discussion in [Pst22, Section 5.2].

Corollary 5.35. *Let E be a homotopy-associative ring spectrum of even Adams type. Then the subcategory $(\mathrm{Syn}_E^{\mathrm{cell}})^{\mathrm{ev}} \subseteq \mathrm{Syn}_E^{\mathrm{cell}}$ is closed under limits. In particular, if Syn_E is cellular, then Pstragowski’s subcategory $\mathrm{Syn}_E^{\mathrm{ev}} \subseteq \mathrm{Syn}_E$ is closed under limits.*

Proof. This follows from Corollary 5.24. \blacksquare

Remark 5.36. Via Remark 5.25, the filtered models for Syn_E from Section 5.2 also carry over to $\mathrm{Syn}_E^{\mathrm{ev}}$. For instance, we obtain a symmetric monoidal equivalence

$$\mathrm{Mod}_{(\nu\mathcal{S})_{\tau}^{\wedge}}(\mathrm{Syn}_E^{\mathrm{ev}}) \simeq \mathrm{Mod}_{\mathrm{Tot}(\tau_{\geq 2*}(E^{\bullet}))}(\mathrm{FilSp}).$$

Example 5.37. The most relevant case for us is where $E = \mathrm{MU}$. Since MU -synthetic spectra are cellular, the above discussion applies. Recall that there is an equivalence

$$\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}^{\mathrm{ev}} \simeq \mathrm{QCoh}(\mathfrak{M}_{\mathrm{fg}}).$$

As a result, the special fibres of $\mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}}$ and $\widehat{\mathrm{Syn}}_{\mathrm{MU}}^{\mathrm{ev}}$ are, respectively,

$$\mathrm{IndCoh}(\mathfrak{M}_{\mathrm{fg}}) \quad \text{and} \quad \mathcal{D}(\mathrm{QCoh}(\mathfrak{M}_{\mathrm{fg}})).$$

From the geometric perspective, the category of all graded $\mathrm{MU}_*\mathrm{MU}$ is a little more awkward, where we have to work with two sheaves (the even and odd parts) separately, so that only even MU -synthetic spectra have this more elegant geometric description. \blacktriangle

5.4 Motivic spectra

We assume basic familiarity with motivic spectra.

Notation 5.38. We write $\mathrm{Sp}_{\mathbf{C}}$ for Morel–Voevodsky’s symmetric monoidal ∞ -category of *motivic spectra* over $\mathrm{Spec} \mathbf{C}$. We write $\mathrm{Sp}_{\mathbf{C}}^{\mathrm{cell}}$ for the smallest stable subcategory of $\mathrm{Sp}_{\mathbf{C}}$ that is closed under colimits and that contains both \mathbf{G}_m and the (categorical) suspension of the unit.

The objects \mathbf{G}_m and the suspension of the unit are the two ‘circles’ of motivic homotopy theory, making it have a notion of bigraded homotopy, and explaining the above notion of cellularity.

Motivic homotopy theory is by nature an algebro-geometric homotopy theory. It is very surprising then that over \mathbf{C} , it turns out to have a deep connection to the Adams–Novikov spectral sequence for ordinary spectra.

- ◆ Levine [Lev15] shows that the Betti realisation of the slice spectral sequence for the \mathbf{C} -motivic sphere spectrum is isomorphic to the (décalage of the double speed) Adams–Novikov filtration for the sphere spectrum (more generally, Levine shows this over an algebraically closed field of characteristic zero).
- ◆ Hu–Kriz–Ormsby [HKO11] show that at the prime 2, differentials in the \mathbf{C} -motivic Adams–Novikov spectral sequence for the motivic sphere can be deduced formally from differentials in the ordinary Adams–Novikov spectral sequence for the sphere. (More generally, they work over an algebraically closed field of characteristic zero.) Stahn [Sta21] showed the analogous odd-primary version of this.
- ◆ There is a twisted endomorphism τ of the motivic sphere spectrum, and its cofibre admits an \mathbf{E}_∞ -structure. Gheorghe–Wang–Xu [GWX21] show that modules over $\mathbf{C}\tau$ in $(\mathrm{Sp}_{\mathbf{C}}^{\mathrm{cell}})_p^\wedge$ is equivalent to the derived ∞ -category of $\mathrm{BP}_*\mathrm{BP}$ -comodules.

Remark 5.39. Bachmann–Kong–Wang–Xu [Bac+22] have generalised the results by Gheorghe–Wang–Xu on the structure of \mathbf{C} -motivic spectra to an arbitrary base field. The answer is more complicated; very roughly speaking, the arithmetic of the base field starts to play a big role (which was not visible in the case of \mathbf{C} because it is algebraically closed).

The above suggests that the \mathbf{C} -motivic category should have a close relation to the MU-synthetic (or BP-synthetic) category. This is true in a very strong sense: up to p -completion, MU-synthetic spectra form a full subcategory of \mathbf{C} -motivic spectra.

Theorem 5.40 (Gheorghe–Isaksen–Krause–Ricka, Pstrągowski). *Let p be an arbitrary prime. There exists a symmetric monoidal equivalence*

$$(\mathrm{Sp}_{\mathbf{C}}^{\mathrm{cell}})_p^\wedge \xrightarrow{\cong} (\mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}})_p^\wedge$$

that sends τ to τ^2 . Under this equivalence, Betti realisation corresponds to τ -inversion.

Proof. See [GIKR21, Theorem 6.12] or [Pst22, Theorem 7.34]. These two equivalences are related via the equivalence of Corollary 5.11 (or rather, the even version of it; see Remark 5.36). ■

The aforementioned connections between \mathbf{C} -motivic spectra and Adams–Novikov spectral sequences now match up with synthetic notions we previously introduced. For instance, the result of [HKO11] now corresponds to the Omnibus Theorem; see also [BHS23, Remark A.2].

Warning 5.41. Unlike what the case of synthetic or ordinary spectra might suggest, the two modifications that we have to do on $\mathrm{Sp}_{\mathbb{C}}$ are quite drastic.

- ◆ Most varieties are not cellular, so that $\mathrm{Sp}_{\mathbb{C}}^{\mathrm{cell}}$ does not see a lot of the algebraic geometry contained in that category.
- ◆ Unlike in the case of ordinary spectra, rational motivic spectra are highly non-trivial and contain a lot of interesting information. Rational objects are killed by p -completion, so we lose a lot of information by passing to p -complete objects.

As explained in Chapter 1, cellular motivic spectra were used in computational stable homotopy theory in various ways. Anachronistically, the equivalence of Theorem 5.40 explains why motivic spectra in particular were so useful in this regard. More seriously, as explained in [IWX23, Section 1.1.2], this provides a way to make these computational arguments using a much lighter technical setup, so that it does not logically depend on the setup of motivic homotopy theory. This is useful as this usage of motivic homotopy theory relies on deep results, such as Voevodsky’s computation of the motivic dual Steenrod algebra [Voe03; Voe10]. The analogous computation in the synthetic setting is a more straightforward adaptation of the spectral one; see [Pst22, Section 6.2]. (One cannot combine this with the equivalence of Theorem 5.40 to obtain a new proof of these computations however, as Voevodsky’s result is input to the proof of Theorem 5.40.)

Example 5.42. Gheorghe–Isaksen–Krause–Ricka [GIKR21, Definition 3.2] define a functor $\Gamma: \mathrm{Sp} \rightarrow \mathrm{FilSp}$ given by sending X to

$$\mathrm{Tot}(\tau_{\geq 2*}(\mathrm{MU}^{\bullet} \otimes X)).$$

Observe that this is a cosimplicial décalage (Definition 2.78) using the double-speed Whitehead filtration instead of the ordinary one. We may identify Γ with the functor ν , up to some slight caveats due to working with even objects and with cosimplicial objects (which can only capture completions). Specifically, under the equivalence

$$\mathrm{Mod}_{\mathrm{Tot}(\tau_{\geq 2*}(\mathrm{MU}^{\bullet}))}(\mathrm{FilSp}) \simeq \mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}}$$

from Remark 5.36, we see that on the full subcategory of Sp on those spectra with even MU-homology, the functor Γ can be identified with the τ -completion of ν . On general spectra, it identifies Γ with the right adjoint to the inclusion $\mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}} \subseteq \mathrm{Syn}_{\mathrm{MU}}$ applied to τ -completion of ν . (One might refer to this right adjoint composed with ν as the *even synthetic analogue*.) ▲

Under the equivalence of Theorem 5.40, the functor ν (or equivalently Γ) is useful because, in general motivic homotopy theory, there is no ‘motivic analogue’ functor. At least over \mathbb{C} , this allows for the construction of new (cellular) motivic spectra.

Example 5.43. In [GIKR21, Section 5], Gheorghe–Isaksen–Krause–Ricka use Theorem 5.40 to define what they call *motivic modular forms*. Namely, they consider $\Gamma(\mathrm{tmf})$, which after

p -completion defines a \mathbf{C} -motivic spectrum. It plays an important part in the computation of Isaksen–Wang–Xu [IWX23].

We prefer to think of this as *synthetic modular forms*, as its construction is inherently synthetic. Phrased in synthetic terms via Example 5.42, this definition is given by

$$\mathbf{smf} := \nu \mathbf{tmf}.$$

This uses that the MU-homology of \mathbf{tmf} is even. (Note that, as \mathbf{tmf} is connective, it is MU-nilpotent complete, so its synthetic analogue is τ -complete.) Note also that these synthetic modular forms are different from the nonconnective versions \mathbf{Smf} and \mathbf{SMF} introduced in joint work of Carrick, Davies and the author in [CDvN25]. \blacktriangle

Recall from Remark 4.20 that there is an alternative grading convention for the synthetic spheres known as *motivic grading*. As the name suggests, this fits better with grading conventions of (stable) motivic homotopy theory.

Remark 5.44 (Indexing conventions). Let us write $\mathbf{S}_{\text{mot}}^{t,w}$ for the motivic bigraded (t, w) -sphere. As explained in [Pst22, Section 7.1], this amounts to

$$\mathbf{S}_{\text{mot}}^{0,0} = \Sigma_+^\infty \mathbf{A}^0 \quad \text{and} \quad \mathbf{S}_{\text{mot}}^{2,1} = \Sigma_+^\infty \mathbf{P}^1.$$

Let us for the moment use the motivic grading on synthetic spectra, writing $\mathbf{S}_{\text{syn}}^{t,w}$ for the synthetic (t, w) -sphere in the sense of Remark 4.20. Then $\mathbf{S}_{\text{syn}}^{t,w}$ is even in the sense of Proposition 5.30 if and only if w is even; see Example 5.34. Under the equivalence of Theorem 5.40, we have the correspondence between

$$\mathbf{S}_{\text{mot}}^{t,w} \quad \text{and} \quad \mathbf{S}_{\text{syn}}^{t,2w}.$$

If we instead use Adams grading on synthetic spectra, then this correspondence is between

$$\mathbf{S}_{\text{mot}}^{t,w} \quad \text{and} \quad \mathbf{S}_{\text{syn}}^{t,2w-t}.$$

Appendix A

Informal introduction to spectral sequences

This chapter is meant as an informal introduction to spectral sequences and the role of τ . Although it is certainly possible to take this chapter as a first introduction to spectral sequences, we particularly have in mind two kinds of readers: one who wants to look at spectral sequences for the second time, and one who is familiar with these, but wants to learn about the τ -formalism specifically.

For the rest of this chapter, we fix a filtered spectrum $X: \mathbf{Z}^{\text{op}} \rightarrow \text{Sp}$. Recall from Chapter 2 that we regard X as a tool to understand its colimit $X^{-\infty}$. For simplicity, and since this covers most of our use cases, throughout this chapter we assume that X is constant from degree 0 onwards:

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

As a result, we will 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 strict 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, to determine which elements have filtration at least 1.

A.1 The reconstruction problem

We have a cofibre sequence

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

leading to a long exact sequence

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

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 $\pi_n X^0 \rightarrow \pi_n \text{Gr}^0 X$. This pattern continues: if $\alpha \in \pi_n X$ has filtration at least 1, we then ask if its filtration is at least 2. Choosing a lift to $\pi_n X^1$, we look at the associated graded $\text{Gr}^1 X$, whose homotopy sits in a long exact sequence

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

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 practice, 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 $\pi_n \text{Gr}^s X$ 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 $\pi_n \text{Gr}^s X$ comes from $\pi_n X^s$ (in other words, there are “fake elements”),
- (2) even if an element in $\pi_n \text{Gr}^s X$ 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 $\pi_* X^0$.

A.2 Differentials: obstructions to lifting

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

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

which tells us that an element in $\pi_n \text{Gr}^s X$ 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

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X^2 & \longrightarrow & X^1 & \longrightarrow & X^0 \\ & & & & & & \downarrow \\ & & & & & & \text{Gr}^0 X, \\ & & & & \swarrow \text{dashed} & & \\ & & & & \partial & & \end{array}$$

where the dashed arrow indicates that the map is of degree 1: it is the boundary map $\partial: \text{Gr}^0 X \rightarrow \Sigma X^1$ of a cofibre sequence. By exactness, an element $x \in \pi_n \text{Gr}^0 X$ comes from $\pi_n X^0$ if and only if its image in $\pi_{n-1} X^1$ is zero. However, as 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 $\partial x \in \pi_{n-1} X^1$ being zero, we first ask if its image in the associated graded $X^1 \rightarrow \text{Gr}^1 X$ is zero:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X^2 & \longrightarrow & X^1 & \longrightarrow & X^0 \\ & & & & \downarrow & \swarrow \text{dashed } \partial & \downarrow \\ & & & & \text{Gr}^1 X & & \text{Gr}^0 X. \end{array}$$

Write $d_1(x)$ for the image of ∂x in $\pi_{n-1} \text{Gr}^1 X$. If $d_1(x) \neq 0$, then $\partial x \neq 0$ as well, so in particular we learn that x is not in the image of $\pi_n X^0$. If on the other hand $d_1(x) = 0$, then we are not yet done: all we learn is that $\partial x \in \pi_{n-1} X^1$ 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 $\pi_{n-1} \text{Gr}^2 X$:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X^2 & \longrightarrow & X^1 & \longrightarrow & X^0 \\ & & \downarrow & & \swarrow \text{dashed } \partial & & \downarrow \\ & & \text{Gr}^2 X & & & & \text{Gr}^0 X. \end{array}$$

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

We obtain inductively defined elements $d_r(x)$ for $r \geq 1$. If they all vanish, then our class x lifts (possibly not uniquely) to an element of the limit $X^\infty = \lim_s X^s$. This gets us into convergence issues. In good situations, this limit vanishes; let us assume that this is the case. This is good news: it means that we can detect whether $\partial x \in \pi_{n-1} X^1$ is zero by checking if the $d_r(x)$ are zero for all $r \geq 1$. This, in turn, means that we can answer the question whether $x \in \pi_n \text{Gr}^0 X$ comes from $\pi_n X$.

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

A.3 Differentials: kernels of transition maps

On to issue (2), which is asking what the kernel of $\pi_n X^s \rightarrow \pi_n X^0$ is. Because the map $X^s \rightarrow X^0$ is a composite of s maps, we can focus on the map $X^s \rightarrow X^{s-1}$ 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 x for its image in $\pi_n \text{Gr}^s X$. Our aim is to understand whether α maps to zero in $\pi_n X^{s-1}$. We have a long exact sequence

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

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 $\pi_{n+1} \text{Gr}^s X \rightarrow \pi_n X^s$. Notice that in terms of x , this is equivalent to the existence of an element $y \in \pi_{n+1} \text{Gr}^{s-1}$ such that $d_1(y) = x$. Iterating this procedure, we find that the element $\alpha \in \pi_n X^s$ maps to a nonzero element in $\pi_n X$ if and only if x is not in the image of d_1, \dots, d_s .

Remark A.1. It might appear there is an asymmetry in the above: to resolve issue (1), we had to check that infinitely many differentials on x 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.

Phrasing the previous story 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 $x \in \pi_n \text{Gr}^s X$. If all differentials on x vanish, and moreover x is not the target of a differential, then any lift of x to $\pi_n X^s$ maps to a nonzero element in $\pi_n X^0$. However, if $d_r(y) = x$ for some $r \leq s$ and some y , then we only learn that *there exists* a lift of x 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 x , then for any $\beta \in \pi_n X^s$ that comes from $\pi_n X^{s+1}$, the element $\alpha + \beta$ also lifts x . But the associated graded has no control over β : it maps to zero in $\pi_n \text{Gr}^s X$. This is a matter we cannot ignore, since β need not even map to zero in $\pi_n X^0$. The summary then is that the associated graded $\pi_n \text{Gr}^s X$ 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 practice, one might be able to bootstrap this together by comparing different spectral sequences: in one spectral sequence, 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 sequences where it is not clear how to choose this lift.

A.4 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

$$E_1^{n,s} := \pi_n \operatorname{Gr}^s X,$$

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

$$d_1 : E_1^{n,s} \longrightarrow E_1^{n-1,s+1}$$

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

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.

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, \dots, d_{r-1} . We therefore inductively define, for $r \geq 2$,

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

Roughly speaking, doing this process infinitely many times results in page ∞ , denoted $E_\infty^{n,s}$. In good cases (the other part of convergence issues), this is isomorphic to the associated graded of the induced strict 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 general, without the simplifying assumption that $X^0 \cong X^{-\infty}$, this would instead be the associated graded of the filtration on $\pi_n X^{-\infty}$.

Remark A.2. The reason the differential goes r units up is because we are using a *cohomological* indexing on the filtration. If we instead indexed X to be a functor $\mathbf{Z} \rightarrow \operatorname{Sp}$, which is a *homological* indexing on filtration, then the differential would decrease the filtration.

In summary: 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 $\pi_n X^s \rightarrow \pi_n X^{s-1}$. Taking homology for a d_r -differential removes both the fake elements, and kills elements that die under $X^s \rightarrow X^{s-r}$. Letting all differentials run brings us to the associated graded of the filtration we were trying to understand.

A.5 Reformulation in terms of τ

It would be useful to introduce some notation to make it easier to describe transition maps. This is what τ is designed to do. If X is a filtered spectrum, let us define

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

Let $\mathbf{Z}[\tau]$ denote the bigraded ring where τ has bidegree $(0, -1)$. We turn $\pi_{*,*} X$ into a bigraded $\mathbf{Z}[\tau]$ -module by letting τ act as the transition maps.

Practically, all this means is that if $\alpha \in \pi_{n,s} X$, then we write $\tau \cdot \alpha$ for the image of α under the transition map $X^s \rightarrow X^{s-1}$. This is helpful as we do not have to give each transition map its own name, which would become quite cumbersome when expressing more involved relations.

The previous discussion then takes the following form. If $x \in \pi_n \text{Gr}^s X$ is a class that is hit by a d_r -differential, then there exists a lift $\alpha \in \pi_{n,s} X$ of it such that

$$\tau^r \cdot \alpha = 0.$$

Further, passing from $\pi_{n,s} X$ to $\pi_n X^{-\infty}$ is given by the colimit along the transition maps, which in this language is given by inverting τ . In this way, we see the fundamental difference between $\pi_{*,*} X$ and $\pi_* X^{-\infty}$: while differentials in the spectral sequence are responsible for killing elements in the latter, in the former they only kill τ -power multiples of it, and thereby leave a trace of their existence.

Bibliography

- [Ada58] J. F. Adams. ‘On the Structure and Applications of the Steenrod Algebra’. *Commentarii Mathematici Helvetici* 32.1 (Dec. 1958), pp. 180–214. DOI: 10.1007/BF02564578.
- [Ada95] J. F. Adams. *Stable Homotopy and Generalised Homology*. Chicago Lectures in Mathematics. Univ. of Chicago Press, 1995. 373 pp.
- [AP25] T. Annala and P. Pstragowski. *A Note on Weight Filtrations at the Characteristic*. 2nd Oct. 2025. arXiv: 2502.19626v3 [math]. Preprint.
- [Ant24] B. Antieau. *Spectral sequences, décalage, and the Beilinson t -structure*. 14th Nov. 2024. arXiv: 2411.09115 [math]. Preprint.
- [Ari21] S. Ariotta. *Coherent cochain complexes and Beilinson t -structures, with an appendix by Achim Krause*. 2nd Sept. 2021. arXiv: 2109.01017v1 [math]. Preprint.
- [Bac+22] T. Bachmann, H. J. Kong, G. Wang and Z. Xu. ‘The Chow t -structure on the ∞ -category of motivic spectra’. *Annals of Mathematics* 195.2 (1st Mar. 2022). DOI: 10.4007/annals.2022.195.2.5.
- [BL01] A. Baker and A. Lazarev. ‘On the Adams spectral sequence for R -modules’. *Algebraic & Geometric Topology* 1.1 (7th Apr. 2001), pp. 173–199. DOI: 10.2140/agt.2001.1.173.
- [BDS15] P. Balmer, I. Dell’Ambrogio and B. Sanders. ‘Restriction to Finite-Index Subgroups as Étale Extensions in Topology, KK-Theory and Geometry’. *Algebraic & Geometric Topology* 15.5 (12th Nov. 2015), pp. 3025–3047. DOI: 10.2140/agt.2015.15.3025.
- [Bar23] S. Barkan. *Chromatic Homotopy Is Monoidally Algebraic at Large Primes*. 5th June 2023. arXiv: 2304.14457v2 [math]. Preprint.
- [BvN25] S. Barkan and S. van Nigtevecht. *Cellularity of Chromatic Synthetic Spectra*. 2nd May 2025. arXiv: 2505.01086v1 [math]. Preprint.

- [Bar+21] T. Barthel, N. Castellana, D. Heard and G. Valenzuela. ‘Local Gorenstein Duality for Cochains on Spaces’. *Journal of Pure and Applied Algebra* 225.2 (Feb. 2021), p. 106495. DOI: 10.1016/j.jpaa.2020.106495.
- [BP23] T. Barthel and P. Pstragowski. ‘Morava K-theory and Filtration by Powers’. *Journal of the Institute of Mathematics of Jussieu* (4th July 2023), pp. 1–77. DOI: 10.1017/S1474748023000233.
- [Boa99] J. M. Boardman. ‘Conditionally Convergent Spectral Sequences’. In: *Contemporary Mathematics*. Ed. by J.-P. Meyer, J. Morava and W. S. Wilson. Vol. 239. American Mathematical Society, 1999, pp. 49–84. DOI: 10.1090/conm/239/03597.
- [Bou79] A. Bousfield. ‘The Localization of Spectra with Respect to Homology’. *Topology* 18.4 (1979), pp. 257–281. DOI: 10.1016/0040-9383(79)90018-1.
- [Bro61] W. Browder. ‘Torsion in H-Spaces’. *The Annals of Mathematics* 74.1 (July 1961), p. 24. DOI: 10.2307/1970305.
- [BR22] R. Bruner and J. Rognes. *The Cohomology of the Mod 2 Steenrod Algebra*. NIRD RDA, 2022. DOI: 10.11582/2022.00015.
- [Bur21] R. Burklund. ‘An Extension in the Adams Spectral Sequence in Dimension 54’. *Bulletin of the London Mathematical Society* 53.2 (Apr. 2021), pp. 404–407. DOI: 10.1112/blms.12428.
- [Bur22] R. Burklund. ‘Synthetic Cookware v0.2’. 17th Mar. 2022. URL: <https://burklund.github.io/cookware-v0.2.pdf>. Draft.
- [BHS22] R. Burklund, J. Hahn and A. Senger. *Galois reconstruction of Artin-Tate \mathbb{R} -motivic spectra*. 30th June 2022. arXiv: 2010.10325v2 [math]. Preprint.
- [BHS23] R. Burklund, J. Hahn and A. Senger. ‘On the Boundaries of Highly Connected, Almost Closed Manifolds’. *Acta Mathematica* 231.2 (2023), pp. 205–344. DOI: 10.4310/ACTA.2023.v231.n2.a1.
- [BIX25] R. Burklund, D. C. Isaksen and Z. Xu. ‘Classical Stable Homotopy Groups of Spheres via \mathbb{F}_2 -Synthetic Methods’. *Peking Mathematical Journal* (11th Apr. 2025). DOI: 10.1007/s42543-025-00098-y.
- [BLP22] R. Burklund, I. Levy and P. Pstragowski. ‘Adams-Type Maps Are Not Stable under Composition’. *Proceedings of the American Mathematical Society, Series B* 9.35 (30th Aug. 2022), pp. 373–376. DOI: 10.1090/bproc/137.
- [BP25] R. Burklund and P. Pstragowski. ‘Quivers and the Adams Spectral Sequence’. *Advances in Mathematics* 471 (June 2025), p. 110270. DOI: 10.1016/j.aim.2025.110270.

- [CD24] C. Carrick and J. M. Davies. *A synthetic approach to detecting v_1 -periodic families*. 29th Jan. 2024. arXiv: 2401.16508 [math]. Preprint.
- [CDvN24] C. Carrick, J. M. Davies and S. van Nigtevecht. *The Descent Spectral Sequence for Topological Modular Forms*. 2nd Dec. 2024. arXiv: 2412.01640v1 [math]. Preprint.
- [CDvN25] C. Carrick, J. M. Davies and S. van Nigtevecht. ‘Descent Spectral Sequences Through Synthetic Spectra’. *International Mathematics Research Notices* 2025.15 (7th Aug. 2025). DOI: 10.1093/imrn/rnaf216.
- [CE56] H. Cartan and S. Eilenberg. *Homological Algebra*. Princeton Mathematical Series 19. Princeton University Press, 1956. DOI: 10.1515/9781400883844.
- [Chu22] D. Chua. *The E_3 page of the Adams spectral sequence*. 1st Apr. 2022. arXiv: 2105.07628v2 [math]. Preprint.
- [Cre25] J. Creemers. ‘Spectral Sequences Obtained from Towers in the Infinity Category of Spectra: Décalage, Exact Couples and Massey Products’. MA thesis. Utrecht University, 22nd Jan. 2025. URL: <https://studenttheses.uu.nl/handle/20.500.12932/48571>.
- [Del71] P. Deligne. ‘Théorie de Hodge: II’. *Institut des Hautes Études Scientifiques* 40 (1971), pp. 5–57. URL: https://www.numdam.org/item/PMIHES_1971__40__5_0/.
- [Dev97] E. S. Devinatz. ‘Morava Modules and Brown-Comenetz Duality’. *American Journal of Mathematics* 119 (1997), pp. 741–770. URL: <http://www.jstor.org/stable/25098553>.
- [Dug14] D. Dugger. ‘Coherence for Invertible Objects and Multigraded Homotopy Rings’. *Algebraic & Geometric Topology* 14.2 (21st Mar. 2014), pp. 1055–1106. DOI: 10.2140/agt.2014.14.1055.
- [Dug17] D. Dugger. ‘A Primer on Homotopy Colimits’. Jan. 2017. URL: <https://pages.uoregon.edu/ddugger/hocolim.pdf>.
- [Dug+24] D. Dugger, B. I. Dundas, D. C. Isaksen and P. A. Østvær. ‘The Multiplicative Structures on Motivic Homotopy Groups’. *Algebraic & Geometric Topology* 24.3 (28th June 2024), pp. 1781–1786. DOI: 10.2140/agt.2024.24.1781.
- [EH66] B. Eckmann and P. Hilton. ‘Exact Couples in an Abelian Category’. *Journal of Algebra* 3.1 (Jan. 1966), pp. 38–87. DOI: 10.1016/0021-8693(66)90019-6.
- [EKMM] A. D. Elmendorf, I. Kříž, M. A. Mandell and J. P. May. *Rings, modules, and algebras in stable homotopy theory*. Mathematical surveys and monographs 47. American Mathematical Society, 2007.

- [GGN16] D. Gepner, M. Groth and T. Nikolaus. ‘Universality of Multiplicative Infinite Loop Space Machines’. *Algebraic & Geometric Topology* 15.6 (12th Jan. 2016), pp. 3107–3153. DOI: 10.2140/agt.2015.15.3107.
- [GIKR21] B. Gheorghe, D. C. Isaksen, A. Krause and N. Ricka. ‘C-motivic modular forms’. *Journal of the European Mathematical Society* 24.10 (9th Nov. 2021), pp. 3597–3628. DOI: 10.4171/jems/1171.
- [GWX21] B. Gheorghe, G. Wang and Z. Xu. ‘The Special Fiber of the Motivic Deformation of the Stable Homotopy Category Is Algebraic’. *Acta Mathematica* 226.2 (2021), pp. 319–407. DOI: 10.4310/ACTA.2021.v226.n2.a2.
- [GJ09] P. G. Goerss and J. F. Jardine. *Simplicial Homotopy Theory*. Modern Birkhäuser Classics Ser. Springer Basel AG, 2009. 1 p. DOI: 10.1007/978-3-0346-0189-4.
- [GM92] J. Greenlees and J. May. ‘Derived Functors of I-adic Completion and Local Homology’. *Journal of Algebra* 149.2 (July 1992), pp. 438–453. DOI: 10.1016/0021-8693(92)90026-I.
- [Hed20] A. P. Hedenlund. ‘Multiplicative Tate Spectral Sequences’. PhD thesis. University of Oslo, 2020. URL: <https://www.mn.uio.no/math/personer/vit/rognesthesis/hedenlund-thesis.pdf>.
- [HR19] G. O. Helle and J. Rognes. ‘Boardman’s Whole-Plane Obstruction Group for Cartan-Eilenberg Systems’. *Documenta Mathematica* 24 (2019), pp. 1855–1878. DOI: 10.4171/dm/717.
- [HL17] M. J. Hopkins and J. Lurie. ‘On Brauer Groups of Lubin-Tate Spectra I’. Mar. 2017. URL: <https://www.math.ias.edu/~lurie/papers/Brauer.pdf>.
- [Hov04] M. Hovey. ‘Homotopy theory of comodules over a Hopf algebroid’. In: *Homotopy theory: relations with algebraic geometry, group cohomology, and algebraic K-theory. Papers from the international conference on algebraic topology, Northwestern University, Evanston, IL, USA, March 24–28, 2002*. 2004, pp. 261–304. DOI: 10.1090/conm/346.
- [HS99] M. Hovey and N. P. Strickland. ‘Morava K-theories and localisation’. *Memoirs of the American Mathematical Society* 139.666 (1999). DOI: 10.1090/memo/0666.
- [HKO11] P. Hu, I. Kriz and K. Ormsby. ‘Remarks on Motivic Homotopy Theory over Algebraically Closed Fields’. *Journal of K-Theory* 7.1 (Feb. 2011), pp. 55–89. DOI: 10.1017/is010001012jkt098.
- [Isa19] D. Isaksen. ‘Stable Stems’. *Memoirs of the American Mathematical Society* 262.1269 (Nov. 2019). DOI: 10.1090/memo/1269.

- [IWX20] D. C. Isaksen, G. Wang and Z. Xu. ‘Stable Homotopy Groups of Spheres’. *Proceedings of the National Academy of Sciences* 117.40 (6th Oct. 2020), pp. 24757–24763. DOI: 10.1073/pnas.2012335117.
- [IWX23] D. C. Isaksen, G. Wang and Z. Xu. ‘Stable Homotopy Groups of Spheres: From Dimension 0 to 90’. *Publications mathématiques de l’IHÉS* 137.1 (June 2023), pp. 107–243. DOI: 10.1007/s10240-023-00139-1.
- [Isa+24] D. C. Isaksen et al. ‘The \mathbb{C} -motivic Adams-Novikov spectral sequence for topological modular forms’. *Advances in Mathematics* 458 (Oct. 2024), p. 109966. DOI: 10.1016/j.aim.2024.109966.
- [Law24a] T. Lawson. *Filtered Spaces, Filtered Objects*. 10th Oct. 2024. arXiv: 2410.08348v1 [math]. Preprint.
- [Law24b] T. Lawson. *Synthetic Spectra Are (Usually) Cellular*. 6th Feb. 2024. arXiv: 2402.03257v2 [math]. Preprint.
- [Lev15] M. Levine. ‘The Adams–Novikov Spectral Sequence and Voevodsky’s Slice Tower’. *Geometry & Topology* 19.5 (20th Oct. 2015), pp. 2691–2740. DOI: 10.2140/gt.2015.19.2691.
- [Lin23] W. Lin. ‘Noncommutative Gröbner Bases and Ext Groups’. *Peking Mathematical Journal* (6th Nov. 2023). DOI: 10.1007/s42543-023-00080-6.
- [LWX25] W. Lin, G. Wang and Z. Xu. *On the Last Kervaire Invariant Problem*. 22nd Feb. 2025. arXiv: 2412.10879 [math]. Preprint.
- [HTT] J. Lurie. *Higher Topos Theory*. Annals of Mathematics Studies no. 170. Princeton University Press, 2009.
- [Lur15] J. Lurie. ‘Rotation Invariance in Algebraic K-theory’. Sept. 2015. URL: <https://www.math.ias.edu/~lurie/papers/Waldhaus.pdf>.
- [HA] J. Lurie. ‘Higher Algebra’. 18th Sept. 2017. URL: <https://www.math.ias.edu/~lurie/papers/HA.pdf>.
- [SAG] J. Lurie. ‘Spectral Algebraic Geometry’. 3rd Feb. 2018. URL: <https://www.math.ias.edu/~lurie/papers/SAG-rootfile.pdf>.
- [Ker] J. Lurie. *Kerodon*. URL: <https://kerodon.net>.
- [Mah81] M. Mahowald. ‘*bo*-resolutions’. *Pacific Journal of Mathematics* 92.2 (1st Feb. 1981), pp. 365–383. DOI: 10.2140/pjm.1981.92.365.
- [Mah82] M. Mahowald. ‘The Image of J in the EHP Sequence’. *The Annals of Mathematics* 116.1 (July 1982), p. 65. DOI: 10.2307/2007048.

- [Man21] L. Mantovani. *Localizations and completions of stable ∞ -categories*. 5th May 2021. arXiv: 2105.02285v1 [math]. Preprint.
- [Mar24] P. Marek. *H \mathbb{F}_2 -synthetic homotopy groups of topological modular forms*. 22nd Feb. 2024. arXiv: 2202.11305v2 [math]. Preprint.
- [MNN17] A. Mathew, N. Naumann and J. Noel. ‘Nilpotence and Descent in Equivariant Stable Homotopy Theory’. *Advances in Mathematics* 305 (Jan. 2017), pp. 994–1084. DOI: 10.1016/j.aim.2016.09.027.
- [McC00] J. McCleary. *A User’s Guide to Spectral Sequences*. 2nd ed. Cambridge University Press, 27th Nov. 2000. DOI: 10.1017/CB09780511626289.
- [Mil81] H. R. Miller. ‘On Relations between Adams Spectral Sequences, with an Application to the Stable Homotopy of a Moore Space’. *Journal of Pure and Applied Algebra* 20.3 (Mar. 1981), pp. 287–312. DOI: 10.1016/0022-4049(81)90064-5.
- [Mil12] H. R. Miller. ‘The Adams Spectral Sequence: Course Notes’. Lecture Notes. 13th Sept. 2012. URL: <https://math.mit.edu/~hrm/papers/notes-adams-spectral-sequence.pdf>.
- [Mou21] T. Moulinos. ‘The Geometry of Filtrations’. *Bulletin of the London Mathematical Society* 53.5 (Oct. 2021), pp. 1486–1499. DOI: 10.1112/blms.12512.
- [NPR24] N. Naumann, L. Pol and M. Ramzi. *A Symmetric Monoidal Fracture Square*. 8th Nov. 2024. arXiv: 2411.05467v1 [math]. Preprint.
- [Pal05] J. H. Palmieri. ‘A user’s guide to the Bockstein spectral sequence’. 2005. URL: <https://sites.math.washington.edu/~palmieri/Papers/bss.pdf>. Preprint.
- [PP23] I. Patchkoria and P. Pstragowski. *Adams Spectral Sequences and Franke’s Algebraicity Conjecture*. 9th July 2023. arXiv: 2110.03669v2 [math]. Preprint.
- [Pst22] P. Pstragowski. ‘Synthetic Spectra and the Cellular Motivic Category’. *Inventiones mathematicae* (20th Dec. 2022). DOI: 10.1007/s00222-022-01173-2.
- [Pst25] P. Pstragowski. ‘Perfect Even Modules and the Even Filtration’. *Journal of the European Mathematical Society* (11th July 2025). DOI: 10.4171/jems/1669.
- [PV22] P. Pstragowski and P. VanKoughnett. ‘Abstract Goerss-Hopkins Theory’. *Advances in Mathematics* 395 (Feb. 2022), p. 108098. DOI: 10.1016/j.aim.2021.108098.
- [Rav78] D. C. Ravenel. ‘A Novice’s Guide to the Adams-Novikov Spectral Sequence’. In: *Geometric Applications of Homotopy Theory II*. Ed. by M. G. Barratt and

- M. E. Mahowald. Vol. 658. Springer Berlin Heidelberg, 1978, pp. 404–475. DOI: 10.1007/BFb0068728.
- [Rav04] D. C. Ravenel. *Complex Cobordism and Stable Homotopy Groups of Spheres*. 2nd ed. AMS Chelsea Publishing v. 347. AMS Chelsea Pub, 2004. 395 pp. DOI: 10.1090/chel/347.
- [Rog12] J. Rognes. ‘MAT9580 - Algebraic Topology III: The Adams Spectral Sequence’. Lecture Notes. 5th June 2012. URL: <https://www.mn.uio.no/math/personer/vit/rognes/kurs/mat9580v12/notes.050612.pdf>.
- [Rog21] J. Rognes. ‘MAT9580 - Algebraisk Topologi III: Spectral Sequences’. Lecture Notes. 2021. URL: <https://www.uio.no/studier/emner/matnat/math/MAT9580/v21/dokumenter/spseq.pdf>.
- [Sch20] D. Schäppi. *Flat Replacements of Homology Theories*. 24th Nov. 2020. arXiv: 2011.12106v1 [math]. Preprint.
- [Stacks] Stacks Project Authors. *The Stacks Project*. URL: <https://stacks.math.columbia.edu>.
- [Sta21] S.-T. Stahn. *The motivic Adams-Novikov spectral sequence at odd primes over \mathbb{C} and \mathbb{R}* . 22nd Jan. 2021. arXiv: 1606.06085 [math]. Preprint.
- [Voe03] V. Voevodsky. ‘Motivic cohomology with $\mathbf{Z}/2$ -coefficients’. *Publications mathématiques de l’IHÉS* 98.1 (Dec. 2003), pp. 59–104. DOI: 10.1007/s10240-003-0010-6.
- [Voe10] V. Voevodsky. ‘Motivic Eilenberg-MacLane Spaces’. *Publications mathématiques de l’IHÉS* 112.1 (Nov. 2010), pp. 1–99. DOI: 10.1007/s10240-010-0024-9.
- [Yan22] L. Yanovski. ‘The monadic tower for ∞ -categories’. *Journal of Pure and Applied Algebra* 226.6 (June 2022), p. 106975. DOI: 10.1016/j.jpaa.2021.106975.