Synthetic modular forms

ISBN	978-90-393-7913-4
DOI	10.33540/3141

Cover	Vincent Jansen
Printing	Ridderprint | ridderprint.nl

Synthetic modular forms

Synthetische modulaire vormen

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof. dr. ir. W. Hazeleger,
ingevolge het besluit van het College voor Promoties
in het openbaar te verdedigen op

dinsdag 11 november 2025 des middags te 12.15 uur

door

Sven van Nigtevecht

geboren op 10 oktober 1998 te Veldhoven

Promotor:
Prof. dr. I. Moerdijk

Copromotor:
Dr. F.L.M. Meier

Beoordelingscommissie:
Prof. dr. M. Behrens
Prof. dr. G.R. Cavalcanti
Prof. dr. D. Gepner
Dr. H.J. Kong
Prof. dr. Z. Xu

Dit proefschrift werd mogelijk gemaakt door financiële steun van de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) door de Vidi-subsidie toegewezen aan dr. F.L.M. Meier met dossiernummer `VI.Vidi.193.111`.

Contents

Introduction

At its core, algebraic topology is concerned with the study of topological spaces up to weak homotopy equivalence; that is, where two topological spaces are considered to be ‘the same’ if there is a map between them that induces an isomorphism on all homotopy groups. Not all questions about topological spaces can be answered in this way, but history has shown that many interesting and nontrivial questions are naturally understood this way. Examples range from more basic ones, such as showing that $\mathbf{R}^n$ is not homeomorphic to $\mathbf{R}^k$ if $n \neq k$, to very involved ones, such as the nonexistence of real division algebras of dimensions greater than 8, counting vector fields on spheres, or counting the number of smooth structures on spheres.

As time went on, the field of algebraic topology grew into the field of *homotopy theory*, stemming from the insight that this homotopical way of thinking arises in many other contexts. Taking this seriously leads one to construct and study objects with meaningful homotopical properties, even though their construction might not have much geometric significance. The *Eilenberg–MacLane spaces* are perhaps the easiest example: these represent cellular cohomology, making them inherently interesting homotopically speaking. Their construction, on the other hand, is not very geometrically meaningful, and it is indeed hard to come by a naturally arising topological space that is an Eilenberg–MacLane space.

We can go even further and build not only tailor-made objects, but even entire tailor-made homotopy theories. An important example is *stable homotopy theory*, which is designed to answer questions about spaces that are invariant under suspension. The objects of stable homotopy theory are called *spectra*; for example, the *sphere spectrum* $\mathbf{S}$ is an object whose homotopy groups are the *stable homotopy groups of spheres*, which are those homotopy groups $\pi_k S^n$ for which suspending defines an isomorphism with $\pi_{k+1} S^{n+1}$. Although simpler, computing these remains an open and motivational problem in stable homotopy theory.

Stable homotopy theory turns out to be significantly easier than the homotopy theory of spaces. A big reason for this is the existence of a very good generalisation of the theory of commutative algebra in the setting of stable homotopy theory, using so-called $\mathbf{E}_\infty$-*ring spectra* in the place of commutative rings; this area is known as *higher algebra*. The sphere spectrum $\mathbf{S}$ has a canonical $\mathbf{E}_\infty$-ring structure, and we can exploit this structure to gain more information about its homotopy groups. For example, we can study maps of $\mathbf{E}_\infty$-ring spectra $\mathbf{S} \to A$. If we can find an $\mathbf{E}_\infty$-ring spectrum A that is easier than $\mathbf{S}$, this becomes particularly effective. Of course, there is a trade-off between how computable A is and how much information from the sphere it can detect. An important first example in this fashion is KO, the spectrum arising from real K-theory.

A much more involved example is the $\mathbf{E}_\infty$-ring of *topological modular forms*, denoted by tmf, introduced by Hopkins and his coauthors. It arises through a combination of higher algebra and the algebraic geometry of elliptic curves (we will come back to this later in this introduction). It gets its name from its homotopy groups: there is a map from $\pi_*(\mathrm{tmf})$ to the ring of (meromorphic) modular forms, but it is neither injective nor surjective. Although the geometric meaning of topological modular forms remains a mystery, this has not stopped homotopy theorists from using it in many different ways. One of the reasons tmf is so important is that it sits in a sweet spot between computability and the nontriviality of the image of $\mathbf{S} \to \mathrm{tmf}$. Although involved, one can compute all of the homotopy groups of tmf. Amazingly, the map $\mathbf{S} \to \mathrm{tmf}$ detects a lot of information, which can be used both in explicit computations of low-dimensional groups of $\mathbf{S}$, and also in understanding the large-scale structure of $\mathbf{S}$. Many important leaps in our understanding of the sphere spectrum rely crucially on tmf: see, for instance, [IWX23]. It has thus cemented itself as a cornerstone of computational stable homotopy theory.

Shockingly, even many years after its definition, a complete proof of the computation of $\pi_*(\mathrm{tmf})$ has never appeared in the literature: all available proofs take certain information about tmf as input. As a result, there is currently no non-circular proof that can be assembled from the literature. This thesis solves this problem by giving a new and independent computation of $\pi_*(\mathrm{tmf})$, thereby also putting the rest of the computational literature on solid ground. This computation is part of joint work with Carrick and Davies [CDvN25; CDvN24], reproduced in this thesis in modified form.

Our approach is made possible by the use of *synthetic spectra*[1] as introduced by Pstrągowksi [Pst22]. This is a recent example of a homotopy theory that is tailor-made to be a tool for understanding stable homotopy theory. In this thesis, we introduce *synthetic modular forms*, which is the analogue of topological modular forms in the setting of synthetic spectra. Its *raison d'être* is not for geometric reasons, but rather to transport tmf to this more exotic homotopy theory. Once it exists in that world, all of the synthetic toolkit becomes available to this computation, which we carry out in detail in this thesis.

[1]The use of the word 'synthetic' in this context has nothing to do with the usual meaning of *synthetic*

Effectively using synthetic modular forms required further development of existing synthetic tools. The main example of such a result is the *Omnibus Theorem*, which functions as a Rosetta stone for passing between spectral sequences and synthetic spectra. The other contribution of this thesis is to develop these tools, and more specifically to give proofs of such results that very straightforwardly generalise to other homotopy theories similar to synthetic spectra. We do this by proving them in yet another homotopy theory, namely that of *filtered spectra*. Not only is this a natural context for these proofs, but it is a homotopy theory that can be used even more generally than synthetic spectra. We hope that our exposition of filtered spectra will make them more accessible to a general homotopical audience.

This thesis is divided into two parts: the first about filtered and synthetic spectra, and the second about the definition and computation of synthetic modular forms. While the second part heavily relies on the first, the first applies much more generally, and is written so that it can be read independently of the second. Each part is accompanied by its own, more detailed introduction; we give a more leisurely overview of both parts here.

Part I: Filtered and synthetic spectra

In both algebraic topology and homotopy theory, far and away the most used tool for computations is that of *spectral sequences*. These arise from choices of resolution, which we often refer to as a *filtrations*, of homotopical objects. For example, if X is a CW complex, then it comes with an increasing sequence[2] consisting of its skeleta:

$$X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots \subseteq X.$$

If h is a homology theory, then applying h to this sequence of inclusions leads to a sequence

$$h_*(X_0) \longrightarrow h_*(X_1) \longrightarrow h_*(X_2) \longrightarrow \cdots \longrightarrow h_*(X).$$

Because X_0 is a discrete set of points, the homology $h_*(X_0)$ is easy to compute. For every n, the difference between $h_*(X_n)$ and $h_*(X_{n+1})$ is 'measured' by $h_*(X_{n+1}/X_n)$, meaning that these are related by a long exact sequence. The quotient X_{n+1}/X_n is a wedge of spheres, so its homology is also easy to compute. The spectral sequence arising from this filtration (known as the *Atiyah–Hirzebruch spectral sequence*) is a device that describes how to reconstruct $h_*(X)$ out of the data $\{\, h_*(X_{n+1}/X_n) \,\}_n$. One summarises this situation by writing

$$h_*(X_{n+1}/X_n) \implies h_*(X).$$

mathematics. Instead, it is used to evoke the sense that these objects are tailor-made for a specific purpose.

The precise meaning behind this is beyond the scope of this introduction, and is deferred to the main text.

In general, a spectral sequence usually starts with a more computable, algebraic approximation, and tries to extract from this approximation a more complicated invariant. We say *tries*, because this is where the true difficulty of spectral sequences lies: there is no algorithm for performing this step, and often requires (significant) further input of some geometric kind. Providing all of the necessary information (the *differentials*) is referred to as *computing* the spectral sequence, which is more or less a synonym for saying that one can successfully compute the object of interest with it.

In stable homotopy theory, the most famous example is the *Adams spectral sequence*. Ever since its introduction by Adams in [Ada58], it has been the instrument of all advances in computations of the stable homotopy groups of spheres. If X is a spectrum (or space), then the starting point of this spectral sequence is computable in terms of the $\mathbf{F}_2$-homology of X, and it tries to compute the stable homotopy groups of X.

More recently, people have realised that it is very useful to work directly with the filtrations rather than the spectral sequences they give rise to. The main reason is that filtrations constitute a good homotopy theory, so that one can use homotopical constructions to form new filtrations. Important examples include the works by Bhatt–Morrow–Scholze [BMS19] and Hahn–Raksit–Wilson [HRW24], where they define filtrations on trace theories by taking homotopy limits of filtrations, leading to new advances in computations of algebraic K-theory. Luckily, setting up a good homotopy theory of filtrations does not require much technology; the insight is realising the value of this approach. To make this methodology more approachable, one of the goals of Part I is to give an introduction to *filtered spectra*, the homotopy theory describing stable filtrations.

A more involved idea in this direction is to form new homotopy theories where the algebraic approximation (in the Adams example, $\mathrm{H}_*(X;\mathbf{F}_2)$) and the object one wants to compute something about (in the Adams example, X) live in the same category. Such a category should essentially consist of resolutions of a special type, namely those resolutions that give rise to the spectral sequence one is trying to study. For Adams spectral sequences, this is what *synthetic spectra* are. Because these are geared towards more specific spectral sequences, they allow for even more powerful computational techniques. This has led to significant advances in recent years in computing stable homotopy groups of spheres [Isa19; IWX20; IWX23; LWX25]. The main player in synthetic spectra is a map called τ, which governs the

[2]The use of the word 'sequence' in *spectral sequence* is most likely not related to this sequence of spaces (or spectra) giving rise to it. The originally intended meaning of the phrase *spectral sequence* is in all likelihood lost to history.

relation between synthetic spectra and Adams spectral sequences. The main result formalising this idea is known as the *Omnibus Theorem* of Burklund–Hahn–Senger [BHS23, Theorem 9.19]. The other goal of Part I is to give an introduction to synthetic spectra.

What distinguishes our exposition from the existing literature is that we focus on the relationship between filtered and synthetic spectra, and introduce τ in filtered spectra first. This has both didactic and practical benefits: the filtered version of τ has a more concrete meaning, but filtered spectra also provide for a natural place to prove many synthetic results. The main showcase for this is that we prove a generalised version of the synthetic Omnibus Theorem by first proving it in the filtered setting, and deducing the generalised synthetic version directly from this. In addition, many (although not all) synthetic computational techniques turn out to be filtered methods. In fact, many aspects of this story are not specific to synthetic spectra: we show how these filtered results export to any so-called *deformation*.

Part II: Synthetic modular forms

Part II *of this thesis is based on the joint works* [CDvN25; CDvN24] *with Christian Carrick and Jack Davies.*

To explain the use of synthetic spectra for topological modular forms, we briefly discuss the aforementioned gap in the literature; we give a more detailed explanation in Chapter 6. To compute $\pi_*(\mathrm{tmf})$, one needs to find a spectral sequence for it. It turns out that the most approachable candidate is the Adams–Novikov spectral sequence, which requires the MU-homology of tmf as its input (where MU is the homology theory of complex bordism). From the definition of tmf however, it is not clear how to compute $\mathrm{MU}_*(\mathrm{tmf})$. The spectrum tmf is defined as the connective cover of a variant labelled Tmf, and the latter is defined as the global sections of a sheaf of $\mathbf{E}_\infty$-ring spectra on the (compactified) moduli stack of elliptic curves. This spectral algebro-geometric definition of Tmf gives one enough control to compute $\mathrm{MU}_*(\mathrm{Tmf})$. However, in general, there is no clear way in which homology theories interact with connective covers. In this case, it turns out we need a piece knowledge about $\pi_*(\mathrm{Tmf})$ known as the *Gap Theorem*, which says that $\pi_n(\mathrm{Tmf})$ vanishes for $-21 < n < 0$. The problem in the literature arises because the only available computation of $\pi_*(\mathrm{Tmf})$ uses the full computation of the Adams–Novikov spectral sequence for tmf.

One way to solve this problem is to compute $\pi_*(\mathrm{Tmf})$ directly. This is more difficult than it sounds, as all potential spectral sequences one could use for it have their own specific problems. The Adams–Novikov spectral sequence for Tmf is riddled with computational problems; see Section 6.3 for more information. A better-behaved one is the *descent spectral sequence*, which arises from the algebro-geometric definition of Tmf. Its defect however is that it has too little structure: it is a very complicated spectral sequence, and in order to compute it, one needs to input information from different spectral sequences. Without this, it is essentially impossible to compute the descent spectral sequence.

Our solution is to define a synthetic spectrum that implements the descent spectral sequence for Tmf: we call this object *synthetic modular forms*, and denote it by Smf. The mere fact that the descent spectral sequence now exists as a synthetic spectrum is enough to solve the problems outlined above, for two reasons.

- Within synthetic spectra, there is an object $\nu\mathbf{S}$ implementing the Adams–Novikov spectral sequence for the sphere spectrum. It has a universal property, being the initial $\mathbf{E}_\infty$-algebra in synthetic spectra. Because Smf is an $\mathbf{E}_\infty$-algebra, this immediately yields a map $\nu\mathbf{S} \to \mathrm{Smf}$, which on underlying spectral sequences is a map along which we can import the information we need as a starting point.
- The fact that Smf implements the descent spectral sequence makes the synthetic toolkit available for its computation, significantly upgrading ordinary spectral sequence techniques.

In fact, it turns out to be crucial to not only have a map of spectral sequences, but rather to have the map $\nu\mathbf{S} \to \mathrm{Smf}$ of $\mathbf{E}_\infty$-algebras in synthetic spectra. We stress then that our usage of synthetic spectra is not merely a matter of taste: we do not know how to get around the aforementioned issues without using these technologies.

We give a very detailed account of the computation of the descent spectral sequence for Tmf. This both proves the Gap Theorem and computes the homotopy groups of tmf; see Chapter 10 for the proofs. An overview of the literature of topological modular forms, and of how our results tie into it, is provided in Chapter 11.

Theorem A. *The homotopy groups $\pi_n \,\mathrm{Tmf}$ vanish for $-21 < n < 0$.*

Theorem B. *The descent spectral sequence for* Tmf *takes the form depicted in Figures* A.2 *to* A.6 *at the prime* 2, *depicted in Figure* A.1 *at the prime* 3, *and collapses otherwise as detailed in Theorem* 9.66.

We hope that the detailed account of this computation will be a helpful example for learning how to apply the filtered and synthetic techniques from Part I.

Remark. Gheorge–Isaksen–Krause–Ricka [GIKR21] define what they call *motivic modular forms*, which through a comparison of synthetic and motivic spectra can also be thought of as a synthetic version of tmf. This is a different object than the synthetic modular forms Smf introduced in this thesis, and in particular cannot be used in the place of Smf to solve the problems outlined above. We give a more detailed comparison in Example 5.43, Remark 8.3, and Chapter 11.

Background and conventions

This thesis is written for a homotopical audience. However, the different goals of the two parts make their prerequisites differ too.

- For Part I, 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.
- For Part II, we have in mind a reader who is familiar with the basic setup of topological modular forms, although we do include a brief review of the definitions and background. We moreover assume that readers of Part II are familiar with Part I.

Throughout the thesis, 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 $\mathscr{S}$. The ∞-category of spectra is denoted by Sp, and we write $\otimes$ for the smash product of spectra. By the term $\mathbf{E}_\infty$-*ring*, we mean an $\mathbf{E}_\infty$-ring spectrum.

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

Note on collaboration

As mentioned above, large parts of this thesis are based on the two papers [CDvN25] and [CDvN24], both of which are joint work with Christian Carrick and Jack Davies. The three of us are all responsible for the entire contents of these papers.

Part I

Filtered and synthetic spectra

Chapter 1

Introduction to Part I

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 the first part to this thesis, 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. The second part of this thesis crucially uses this in the case of synthetic modular forms.

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 the rest of Part I.

1.1 Filtered spectra

A *filtered spectrum* is a functor $X\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$, 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 this thesis 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 $C^s_\bullet$ 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 a 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

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

The (attempted) passage from the homotopy groups $\pi_n \mathrm{Gr}^s X$ 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 C. 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 $\pi_n X^{s+1} \to \pi_n X^s$; 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 $\pi_n X^s \to \pi_n X^{s-r}$ (for some $r \geqslant 1$); 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 $X^s \to X^{s-1}$. 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 C 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 $\mathrm{Gr}^s X \to \Sigma X^{s+1}$ encode information about differentials of *all* lengths. We refer to these are *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 $\mathbf{E}_\infty$-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

$$\mathrm{E}_2^{n,s} \cong \mathrm{Ext}_{E_*E}^{s,\,n+s}(E_*,\, E_*X) \implies \pi_n X,$$

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 $\mathbf{F}_p$-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 $E = \mathrm{MU}$, 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 $E = \mathbf{F}_2$, the state of the art in terms of computing the Ext groups is for $n + s \leqslant 200$ 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 this thesis, 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 $\mathrm{Syn}_E(\mathrm{Sp})$ of *E-synthetic spectra*, along with functors

$$\nu_E\colon \mathrm{Sp} \longrightarrow \mathrm{Syn}_E(\mathrm{Sp}) \qquad \text{and} \qquad \sigma\colon \mathrm{Syn}_E(\mathrm{Sp}) \longrightarrow \mathrm{FilSp},$$

called the *E-synthetic analogue* and *signature* functor, respectively. A series of computations in $\mathrm{Syn}_E(\mathrm{Sp})$ 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

$$\mathrm{FilSp} \underset{\sigma}{\overset{\rho}{\rightleftarrows}} \mathrm{Syn}_E(\mathrm{Sp}).$$

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 $\mathrm{Syn}_E(\mathrm{Sp})$. In this way, all results in the τ-formalism directly apply to $\mathrm{Syn}_E(\mathrm{Sp})$ 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 $\mathrm{Syn}_E(\mathrm{Sp})$ is simpler than FilSp are the following.

- There is a t-structure on $\mathrm{Syn}_E(\mathrm{Sp})$, 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 $\mathrm{Syn}_E(\mathrm{Sp})$ 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 $\mathrm{Syn}_E(\mathrm{Sp})$ is closely related to E-Adams spectral sequences: for instance, the Ext groups on the E_2-page 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 $\mathrm{Syn}_E(\mathrm{Sp})$ is the natural home for Adams spectral sequences.

As another example, we have a commutative diagram

$$\begin{tikzcd} \mathrm{Syn}_E(\mathrm{Sp}) \arrow[r, "\sigma"] \arrow[d, "C\tau\otimes-"'] & \mathrm{FilSp} \arrow[d, "C\tau\otimes-"] \\ \mathrm{Mod}_{C\tau}(\mathrm{Syn}_E(\mathrm{Sp})) \arrow[r, "\sigma"] & \mathrm{Mod}_{C\tau}(\mathrm{FilSp}). \end{tikzcd}$$

As explained before, the vertical functor on the right can be identified with the associated graded functor. The diagram now tells us that this factors through $C\tau$-modules in $\mathrm{Syn}_E(\mathrm{Sp})$. 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]. This specific application will not be used to deduce differentials in this thesis, but 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 $\mathrm{Syn}_E(\mathrm{Sp})$ (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, one that we will use heavily in this thesis, 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 Appendix B.

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 $\mathrm{Sp}_{\mathbf{C}} \to \mathrm{Sp}$ from **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

[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.

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 uses 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 movitic 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 *$C\tau$-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 defined E-synthetic spectra, and proved that if $E = \mathrm{MU}$, 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 $E = \mathrm{MU}$. 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 this thesis, 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

We continue to follow the conventions mentioned in the introduction to the thesis. In addition, we adhere to the following conventions.

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

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

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

$$\mathrm{E}_2^{n,s} \cong \mathrm{Ext}_{E_*E}^{s,\, n+s}(E_*,\, E_*X).$$

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, 2.28, 2.38 and 2.51, Construction 2.29, and Definition 2.31.
- First vs. second-page indexing: Remarks 2.37 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 C. 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 $\operatorname{Gr} F$ given by

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

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

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 $f\colon A \to B$ 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 $\lim^1_s F^s A$ vanishes;*

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

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

Then f is an isomorphism of abelian groups, and moreover restricts to an isomorphism $F^s A \xrightarrow{\cong} F^s B$ 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

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

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 $\lim^1_s F^s$ 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 this thesis, 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 $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Ab}$, where we view $\mathbf{Z}$ as a poset under the usual ordering. We write

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

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

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

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

(4) If A is a filtered abelian group, then its **associated graded** is the graded abelian group $\mathrm{Gr}\,A$ given by

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

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

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

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 this thesis, 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 2.61.

Remark 2.9 (Homological vs. cohomological grading). Fitting with Remark 2.3, we regard the usage of $\mathbf{Z}^{\mathrm{op}}$ 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 $\mathbf{Z}^{\mathrm{op}}$ 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, together with a map $A^{-\infty} \to B$.

  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

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

  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 spectrum need not be injective.

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

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

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

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 2.26 below. More specifically, by post-composing with the inclusion, a filtered abelian group A determines a functor $\mathbf{Z}^{\mathrm{op}} \to \mathcal{D}(\mathrm{Ab})$ to the derived ∞-category of abelian groups. Then A is derived complete if and only if the (derived) limit of $\mathbf{Z}^{\mathrm{op}} \to \mathcal{D}(\mathrm{Ab})$ is zero. Because the forgetful functor $\mathcal{D}(\mathrm{Ab}) \to \mathrm{Sp}$ preserves limits, this is equivalent to viewing A as a functor $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$ landing in discrete (a.k.a. Eilenberg–MacLane) spectra, and asking for the (homotopy) limit of this functor to vanish. See Remark 2.48 for an elaboration on this point.

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

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

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

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

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

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

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

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

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

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 $A^s \to F^s$ 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 $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$, where we view $\mathbf{Z}$ as a poset under the usual ordering. We write

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

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

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

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

(4) A **graded spectrum** is a functor $\mathbf{Z}^{\mathrm{discr}} \to \mathrm{Sp}$, where $\mathbf{Z}^{\mathrm{discr}}$ is the discrete category with objects $\mathbf{Z}$. We write

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

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

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

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

We often depict a filtered spectrum $X\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$ as a diagram

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

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

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

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 $\pi_*\colon \mathrm{Sp} \to \mathrm{grAb}$ induces a functor

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

We can find corepresenting objects for the individual abelian groups of this functor. By Yoneda, the transition maps 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 $i\colon \mathbf{Z} \to \mathrm{FilSp}$ for the functor given by

$$s \longmapsto \Sigma^\infty_+ \mathrm{Hom}_{\mathbf{Z}}(-, s).$$

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 $\Sigma^{n,s}\colon \mathrm{FilSp} \to \mathrm{FilSp}$ for the functor given by tensoring with $\mathbf{S}^{n,s}$ on the left.

(3) We write $\pi_{n,s}\colon \mathrm{FilSp} \to \mathrm{Ab}$ 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

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

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 2.37 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 $\mathrm{Map}(\mathbf{S}^{0,s}, X) \cong \Omega^\infty X^s$, 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

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

We write $\tau_{\geqslant n}^{\mathrm{diag}}$ and $\tau_{\leqslant n}^{\mathrm{diag}}$ 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*

$$\pi_{n,s} X = 0 \qquad \textit{whenever } n < s.$$

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

$$\pi_{n,s} X = 0 \qquad \textit{whenever } n > s.$$

(c) *The connective cover $\tau_{\geqslant 0}^{\mathrm{diag}} X \to X$ induces an isomorphism*

$$\pi_{n,s}(\tau_{\geqslant 0}^{\mathrm{diag}} X) \xrightarrow{\cong} \pi_{n,s}(X) \qquad \textit{whenever } n \geqslant s.$$

Likewise, the 0*-truncation $X \to \tau_{\leqslant 0}^{\mathrm{diag}} X$ induces an isomorphism*

$$\pi_{n,s}(X) \xrightarrow{\cong} \pi_{n,s}(\tau_{\leqslant 0}^{\mathrm{diag}} X) \qquad \textit{whenever } n \leqslant s.$$

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

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

(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 $\mathrm{Const}\colon \mathrm{Sp} \to \mathrm{FilSp}$ given by precomposition with $\mathbf{Z}^{\mathrm{op}} \to \Delta^0$.

(2) The **Whitehead filtration** is the functor $\mathrm{Wh}\colon \mathrm{Sp} \to \mathrm{FilSp}$ given by $\tau_{\geqslant 0}^{\mathrm{diag}} \circ \mathrm{Const}$.

(3) The **Postnikov tower** is the functor $\mathrm{Post}\colon \mathrm{Sp} \to \mathrm{FilSp}$ given by $\tau_{\leqslant 0}^{\mathrm{diag}} \circ \mathrm{Const}$.

It is a straightforward exercise to see that these definitions result in the Whitehead filtration and Postnikov tower as we know them: if X is a spectrum, then $\mathrm{Wh}\, X$ is

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

and $\mathrm{Post}\, X$ is

$$\cdots \longrightarrow \tau_{\leqslant 1} X \longrightarrow \tau_{\leqslant 0} X \longrightarrow \tau_{\leqslant -1} X \longrightarrow \cdots .$$

Note that, from the definition, the functors Wh and Post come with natural maps

$$\mathrm{Wh}\, X \longrightarrow \mathrm{Const}\, X \qquad \text{and} \qquad \mathrm{Const}\, X \longrightarrow \mathrm{Post}\, X.$$

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

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.

Definition 2.26. Let X be a filtered spectrum. We say that X is **complete** if the limit X^∞ of X vanishes. We write $\widehat{\mathrm{FilSp}}$ for the full subcategory of FilSp on the complete filtered spectra.

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

$$\mathrm{FilSp} \longrightarrow \widehat{\mathrm{FilSp}}, \quad X \longmapsto \widehat{X}.$$

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

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

Note that, because $\pi_* \colon \mathrm{Sp} \to \mathrm{grAb}$ 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 this thesis, 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 C.

Remark 2.28. 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 C. 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 this thesis).

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

Construction 2.29. Let X be a filtered spectrum.

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

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

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.

$$\begin{tikzcd} \pi_n X^s \arrow[rr, "{(0,-1)}"] & & \pi_n X^s \arrow[dl, "{(0,0)}"] \\ & \pi_n \operatorname{Gr}^s X \arrow[ul, "{(-1,1)}"] & \end{tikzcd}$$

(2) The **underlying spectral sequence** of X, denoted by $\{\, \mathrm{E}_r^{n,s}(X),\ d_r \,\}_{r\geqslant 1}$, 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

$$\mathrm{E}_1^{n,s}(X) := \pi_n \operatorname{Gr}^s X \implies \pi_n X^{-\infty}$$

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

The term *abutment* and the above notation are not 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 2.30. 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

$$\mathrm{E}_1^{n,s} \cong \begin{cases} \pi_n \mathbf{S} & \text{if } s = 0, \\ 0 & \text{if } s \neq 0. \end{cases}$$

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. ▲

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

Definition 2.31. Let X be a filtered spectrum. Let $r \geqslant 0$, and let $n, s \in \mathbf{Z}$.

(1) Write $\mathrm{Z}_r^{n,s} \subseteq \mathrm{E}_1^{n,s}$ for the subset on the ***r*-cycles**, i.e., those x such that the differentials $d_1(x), \ldots, d_r(x)$ vanish. This leads to a sequence of inclusions

$$\cdots \subseteq \mathrm{Z}_2^{n,s} \subseteq \mathrm{Z}_1^{n,s} \subseteq \mathrm{Z}_0^{n,s} = \mathrm{E}_1^{n,s}.$$

We define

$$\mathrm{Z}_\infty^{n,s} := \lim_r \mathrm{Z}_r^{n,s} = \bigcap_r \mathrm{Z}_r^{n,s}.$$

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

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

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

We define

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

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

(3) If $r \geqslant 1$, 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

$$E_\infty^{n,s} := Z_\infty^{n,s} / B_\infty^{n,s},$$
$$RE_\infty^{n,s} := \lim_r{}^1 Z_r^{n,s}.$$

Warning 2.32. 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 2.33. 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 $r \geqslant 1$. 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 $X^s \to X^{-\infty}$.

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

Remark 2.34. The underlying spectral sequence from Construction 2.29 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

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

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 SSeq(bigrAb). 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 2.35 (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 C. There are a number of alternative approaches in the literature, including at least the following.

- Cartan–Eilenberg systems [CE16, 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. This was introduced by Antieau, 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 $\widehat{\mathrm{FilSp}}$ 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] shows that the spectral sequence arising from this Cartan–Eilenberg system is isomorphic to the one arising from iterating the décalage functor. More generally, he shows this when working with filtered objects of a stable ∞-category with sequential limits and colimits that is equipped with a t-structure.

[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 2.61 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 free summand of $\pi_* X$.

- 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 2.29. Creemers [Cre25, Theorem 5.1] shows that the spectral sequence arising from this Cartan–Eilenberg system is isomorphic to the one arising from this exact couple.

Remark 2.36. 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 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 2.37 (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

$$\widetilde{\mathrm{E}}_{r+1}^{n,s} := \mathrm{E}_r^{n,\,s+n}.$$

It is straightforward to check that this turns d_r-differentials into $\widetilde{d}_{r+1}$-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)

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

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 2.38. 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 2.51.

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.

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 this thesis. An alternative introduction is given by Hedenlund in [Hed20, Section 1.2.2].

Recall from Definition 2.27 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 2.39. Let X be a filtered spectrum. Write $\partial_{n,s}\colon \pi_n \operatorname{Gr}^s X \to \pi_{n-1} X^{s+1}$ 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

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

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

$$\operatorname{Gr}^s\,\pi_n X^{-\infty} = \frac{F^s\,\pi_n X^{-\infty}}{F^{s+1}\,\pi_n X^{-\infty}} \hookrightarrow \mathrm{E}_\infty^{n,s}. \tag{2.40}$$

Definition 2.41. 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.,

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

(b) the natural map (2.40) is an isomorphism for all n and s.

Warning 2.42. 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 spectral sequence.

[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.

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 2.43. Because the map (2.40) is always injective, we learn the following, even in the absence of any of the above convergence criteria. If $x \in \mathrm{E}_1^{n,s}$ 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 (2.40), the image of α defines a nonzero map in $\mathrm{Gr}^s \, \pi_n X^{-\infty}$. 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 2.44. 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 2.43).

Remark 2.45. Boardman [Boa99, Definition 5.2] 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 2.41 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 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 2.46. 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 2.47. As with strong convergence, this terminology is abusive: conditional convergence is a condition on the filtration rather than the underlying spectral sequence.

Remark 2.48. 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

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

Warning 2.49. 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:

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

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

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

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 2.50 below. In general, Boardman's *whole-plane obstruction* is the obstruction to this implication; see Remark 2.54 below for a further discussion, and [Boa99, Lemma 8.11] for the precise result alluded to here. Clearly, if both 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 2.54.

Definition 2.50. Let X be a filtered spectrum.

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

(2) We say that X is **left concentrated** if the transition maps $X^s \to X^{s+1}$ 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 2.51. In practice, we usually reindex a right-concentrated filtered spectrum to be of the form

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

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 2.52 (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* 2.31 (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 2.49):

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

■

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: 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 2.53. *Let X be a filtered spectrum. Suppose that for every n, the groups $\pi_n X^s$ 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 2.39, it is clear that the map (2.40) is an isomorphism if $Z_\infty^{n,s} = \ker \partial_{n,s}$ for all n, s. The inclusion $Z_\infty^{n,s} \subseteq \ker \partial_{n,s}$ 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 2.48. ∎

Remark 2.54 (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 2.35 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 $\pi_* \colon \mathrm{Sp} \to \mathrm{grAb}$ 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 this thesis. Let us also point out that this duality is a feature specific to the stable setting.

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

$$\begin{aligned} (-)^{\mathrm{tow}} \colon \mathrm{FilSp} &\longrightarrow \mathrm{FilSp}, \quad X \longmapsto \mathrm{cofib}(\Sigma^{0,-1} X \to \mathrm{Const}\, X^{-\infty}), \\ (-)^{\mathrm{fil}} \colon \mathrm{FilSp} &\longrightarrow \mathrm{FilSp}, \quad X \longmapsto \mathrm{fib}(\mathrm{Const}\, X^{\infty} \to \Sigma^{0,1} X). \end{aligned}$$

Concretely, if X is a filtered spectrum, then

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

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 2.56. Let X be a filtered spectrum. The **fibre-associated graded** of X is the graded spectrum $\mathrm{fibGr}\, X$ given by

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

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

$$\begin{tikzcd} & \mathrm{fibGr}^1 X \arrow[d] & \mathrm{fibGr}^0 X \arrow[d] & \mathrm{fibGr}^{-1} \arrow[d] & \\ \cdots \arrow[r] & X^1 \arrow[r] & X^0 \arrow[r] & X^{-1} \arrow[r] & \cdots . \end{tikzcd}$$

The following is a straightforward diagram chase.

Proposition 2.57. *Let X be a filtered spectrum. Then there are natural isomorphisms of graded spectra*

$$\mathrm{fibGr}(X^{\mathrm{tow}}) \cong \mathrm{Gr}\, X \qquad \textit{and} \qquad \mathrm{Gr}(X^{\mathrm{fil}}) \cong \mathrm{fibGr}\, X.$$

The reflection functors from Definition 2.55 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 2.58. We call a filtered spectrum **cocomplete** if its colimit vanishes. We write $\breve{\mathrm{FilSp}}$ for the full subcategory of FilSp on the cocomplete filtered spectra.

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

$$\mathrm{FilSp} \longrightarrow \breve{\mathrm{FilSp}}, \quad X \longmapsto \check{X} := \mathrm{fib}(X \to X^{-\infty}).$$

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

$$\begin{tikzcd} \breve{\mathrm{FilSp}} \arrow[r, hook, shift left] & \mathrm{FilSp}. \arrow[l, shift left, "\breve{(-)}"] \end{tikzcd}$$

We think of $\widehat{\mathrm{FilSp}}$ as the conditionally convergent filtrations, and of $\breve{\mathrm{FilSp}}$ as the conditionally convergent towers. The reflection functors of Definition 2.55 translate between these in the following way.

Proposition 2.59.

(1) *We have commutative diagrams*

$$\begin{tikzcd} \mathrm{FilSp} \arrow[r, "(-)^{\mathrm{tow}}"] \arrow[d, "\widehat{(-)}"'] & \mathrm{FilSp} \\ \widehat{\mathrm{FilSp}} \arrow[r, "(-)^{\mathrm{tow}}"'] & \breve{\mathrm{FilSp}} \arrow[u, hook] \end{tikzcd} \qquad \textit{and} \qquad \begin{tikzcd} \mathrm{FilSp} \arrow[r, "(-)^{\mathrm{fil}}"] \arrow[d, "\breve{(-)}"'] & \mathrm{FilSp} \\ \breve{\mathrm{FilSp}} \arrow[r, "(-)^{\mathrm{fil}}"'] & \widehat{\mathrm{FilSp}}. \arrow[u, hook] \end{tikzcd}$$

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

$$(X^{\mathrm{tow}})^{\infty} \cong (\hat{X})^{-\infty} \qquad \textit{and} \qquad (X^{\mathrm{fil}})^{-\infty} \cong (\check{X})^{\infty}.$$

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

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

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

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

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

$$\widehat{\mathrm{FilSp}} \underset{(-)^{\mathrm{fil}}}{\overset{(-)^{\mathrm{tow}}}{\rightleftarrows}} \widecheck{\mathrm{FilSp}}.$$

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 $X \to \widehat{X}$ becomes an isomorphism after taking associated towers. Write T for X^{tow} and R for $(\widecheck{X})^{\mathrm{tow}}$. Then for every s, we have a commutative diagram where all rows and columns are cofibre sequences

$$\begin{tikzcd}
X^{\infty} \arrow[r] \arrow[d, equal] & X^{s+1} \arrow[r] \arrow[d] & (\widehat{X})^{s+1} \arrow[d] \\
X^{\infty} \arrow[r] \arrow[d] & X^{-\infty} \arrow[r] \arrow[d] & (\widehat{X})^{-\infty} \arrow[d] \\
0 \arrow[r] & T^{s} \arrow[r] & R^{s}.
\end{tikzcd}$$

It follows that $T^s \to R^s$ 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

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

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 $((-)^{\mathrm{tow}})^{\mathrm{fil}}$ 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 for X^{tow}, and F for T^{fil}. Then using that $X^{\infty} = 0$, we have a commutative diagram where all rows and columns are cofibre sequences

$$\begin{tikzcd}
\Omega X^{s} \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & X^{s} \arrow[d] \\
0 \arrow[r] \arrow[d] & X^{-\infty} \arrow[r, equal] \arrow[d] & X^{-\infty} \arrow[d] \\
F^{s} \arrow[r] & T^{\infty} \arrow[r] & T^{s-1}.
\end{tikzcd}$$

This supplies a natural cofibre sequence $\Omega X \to 0 \to F$, proving the claim. ∎

Example 2.60. 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

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

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

Example 2.61. 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 $X^\wedge_p$. If we take the associated filtration of this tower, then we obtain the p-adic filtration on $X^\wedge_p$:

$$\cdots \xrightarrow{\;p\;} X^\wedge_p \xrightarrow{\;p\;} X^\wedge_p \xrightarrow{\;p\;} X^\wedge_p =\!=\!= \cdots.$$

▲

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 2.73). This limitation is not essential by any means, and can avoided by working with filtered resolutions (Remark 2.85). Later in this thesis, we will use synthetic spectra to set up this improved version; see Section 4.4.

2.5.1 The Tot spectral sequence

Definition 2.62. Let $\mathcal{C}$ be a pointed ∞-category, and let $X^\bullet \colon \Delta \to \mathcal{C}$ be a cosimplicial object of $\mathcal{C}$.

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

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

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

$$\operatorname{Tot}_n X^\bullet = \lim_{\Delta_{\leqslant n}} X^\bullet.$$

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

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

we obtain a tower

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

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

We call this a *tower* in accordance with Remark 2.8; below in Definition 2.68, we will turn this tower into a filtration. Note that the Tot tower is in particular a cocomplete object of $\operatorname{Fil}(\mathcal{C})$ in the sense of Definition 2.58.

Remark 2.63 (Cubes). The n-th partial totalisation can be computed as follows. We write

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

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

$$\Delta^{\mathrm{inj}}_{/[n]} \longrightarrow \Delta_{\leqslant n},$$

and this functor is homotopy initial; see [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 2.64. 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 **semicosimplial object** of $\mathcal{C}$ is a functor $\Delta_{\mathrm{inj}} \to \mathcal{C}$. The definition of the Tot tower can be mimicked for a semicosimplicial object, instead taking the limit over Δ_{inj} or over $(\Delta_{\mathrm{inj}})_{\leqslant n}$. There is a subtle difference between these two.

Remark 2.65. The inclusion $j\colon \Delta_{\mathrm{inj}} \subseteq \Delta$ is homotopy initial; see [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

$$\begin{tikzcd} \mathrm{Sp}^{\Delta} \arrow[rr, "\operatorname{Tot}"] \arrow[dr, "j^*"'] & & \mathrm{Sp}. \\ & \mathrm{Sp}^{\Delta_{\mathrm{inj}}} \arrow[ur, "\operatorname{Tot}"'] & \end{tikzcd}$$

However, the same is not true for the partial totalisations: the inclusion $(\Delta_{\mathrm{inj}})_{\leqslant n} \subseteq \Delta_{\leqslant n}$ is not homotopy initial; see [Dug17, Section 21.6]. Limits over $(\Delta_{\mathrm{inj}})_{\leqslant n}$ are bigger in some sense; see Remark 2.66 and Proposition 2.71.

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

Remark 2.66. Right Kan extension along $j\colon \Delta_{\mathrm{inj}} \subseteq \Delta$ is a right adjoint to the forgetful functor $\mathrm{Sp}^{\Delta} \to \mathrm{Sp}^{\Delta_{\mathrm{inj}}}$, and likewise for Ab in the 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

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

with the product ranging over all surjective maps. 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 $\Delta_{\leqslant n}$ is right Kan extended from the restriction of $X^\bullet$ to $\Delta_{\leqslant n}$. 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 simplicial objects and certain towers. In the following, let us write $\mathrm{Fil}^{\geqslant 0}(\mathcal{C})$ 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 2.62 lands in $\mathrm{Fil}^{\geqslant 0}(\mathcal{C})$.

Theorem 2.67 (Stable Dold–Kan correspondence, Lurie). *Let $\mathcal{C}$ be a stable ∞-category. Then the Tot tower functor restricts to an equivalence*

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

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

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

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

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 2.68. Let $X^\bullet$ be a cosimplicial spectrum. By reflecting the Tot tower of $X^\bullet$ via the functor $(-)^{\mathrm{fil}}$ from Definition 2.55, we obtain a filtered spectrum that we call the **totalisation filtration** (or *Tot filtration*) of $X^\bullet$. We denote its terms by

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

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

$$\cdots \longrightarrow \mathrm{Tot}^2 X^\bullet \longrightarrow \mathrm{Tot}^1 X^\bullet \longrightarrow \mathrm{Tot}\, X^\bullet =\!=\!= \mathrm{Tot}\, X^\bullet =\!=\!= \cdots$$

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 Tot $X^\bullet$. In general however, the conditional convergence of the Tot spectral sequence need not be strong.

Remark 2.69. From Theorem 2.67 and Proposition 2.59, it follows that the Tot filtration functor induces an equivalence from $\mathcal{C}^\Delta$ to the full subcategory of $\widehat{\mathrm{Fil}}(\mathcal{C})$ 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 2.70.

(1) Let $A^\bullet$ be a semicosimplicial abelian group. The **unnnormalised cochain complex** $\mathrm{C}(A^\bullet)$ of $A^\bullet$ is the cochain complex

$$\cdots \longrightarrow 0 \longrightarrow \mathrm{C}^0(A^\bullet) \longrightarrow \mathrm{C}^1(A^\bullet) \longrightarrow \mathrm{C}^2(A^\bullet) \longrightarrow \cdots$$

given by, for $m \geqslant 0$,

$$\mathrm{C}^m(A^\bullet) = A^m$$

and with differential $\mathrm{C}^m(A^\bullet) \to \mathrm{C}^{m+1}(A^\bullet)$ 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** $\mathrm{N}(A^\bullet)$ of $A^\bullet$ is the cochain complex

$$\cdots \longrightarrow 0 \longrightarrow \mathrm{N}^0(A^\bullet) \longrightarrow \mathrm{N}^1(A^\bullet) \longrightarrow \mathrm{N}^2(A^\bullet) \longrightarrow \cdots$$

given by, for $m \geqslant 0$,

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

and with differential $\mathrm{N}^m(A^\bullet) \to \mathrm{N}^{m+1}(A^\bullet)$ 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

$$\mathrm{N}(A^\bullet) \longrightarrow \mathrm{C}(A^\bullet).$$

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].

Proposition 2.71.

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

$$\mathrm{E}_1^{n,s}(X^\bullet) \cong \mathrm{N}^s(\pi_{n+s}\, X^\bullet)$$

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

(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 2.66). For all integers n and s, there is a natural isomorphism*

$$\mathrm{E}_1^{n,s}(j_* X^\bullet) \cong \mathrm{C}^s(\pi_{n+s}\, X^\bullet)$$

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

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

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

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 $\mathcal{C} = \mathrm{Sp}^{\mathrm{op}}$) the discussion of [HA, Remark 1.2.4.4, Variant 1.2.4.9]. ∎

Remark 2.72. One can phrase item (3) at the level of filtered spectra using the décalage functor. Let $\mathrm{Déc}\colon \mathrm{FilSp} \to \mathrm{FilSp}$ denote Antieau's décalage functor; see [Ant24] or [Hed20, Part II]. Then if $X^\bullet$ is a cosimplicial spectrum, the above implies that the unit $X^\bullet \to j_* j^* X^\bullet$ induces an isomorphism of filtered spectra

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

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

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

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 2.35), the cosimplicial version has a simpler expression. These two definitions of décalage actually coincide: see Proposition 2.77.

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 2.73. Let $X^\bullet \colon \Delta \to \mathrm{Sp}$ be a semicosimplicial spectrum. We define the **décalage** of $X^\bullet$ as the filtered spectrum given by

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

If $X^\bullet$ is a cosimplicial spectrum, then we let $\mathrm{Déc}^{\Delta} X^\bullet$ denote the décalage of the underlying semicosimplicial object $j^* X^\bullet$; cf. Remark 2.65.

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

Remark 2.74. The functor $\mathrm{Déc}^{\Delta} \colon \mathrm{Sp}^{\Delta_{\mathrm{inj}}} \to \mathrm{FilSp}$ is naturally lax symmetric monoidal. Indeed, it is the composite

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

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 $\mathrm{FilSp} \to \mathrm{Fil}(\mathrm{Sp})^{\Delta_{\mathrm{inj}}}$ 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 2.37.

Theorem 2.75 (Levine). *Let $X^\bullet$ be a (semi)cosimplicial spectrum, and let $\{ \mathrm{E}_r^{*,*}(X^\bullet) \}_{r\geqslant 1}$ denote the Tot spectral sequence associated to $X^\bullet$. Then there is an isomorphism of spectral sequences (where $r \geqslant 1$)*

$$\mathrm{E}_r^{n,s}(\mathrm{Déc}^{\Delta} X^\bullet) \cong \mathrm{E}_{r+1}^{n,\,s-n}(X^\bullet).$$

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 $\mathrm{Déc}^{\Delta} X^{\bullet}$. Phrased like this, the above theorem says that this second-page indexed $\mathrm{Déc}^{\Delta} X^{\bullet}$ is isomorphic to the second page onwards of $\mathrm{E}_r^{*,*}(X^{\bullet})$.

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

Proposition 2.76. *Let $X^{\bullet}$ be a (semi)cosimplicial spectrum. Then the filtered spectrum $\mathrm{Déc}^{\Delta} X^{\bullet}$ is naturally a complete filtration of $\mathrm{Tot}\, X^{\bullet}$, meaning that its limit vanishes and its colimit is naturally isomorphic to $\mathrm{Tot}\, X^{\bullet}$.*

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

$$\tau_{\geqslant s+1} X^{\bullet} \longrightarrow X^{\bullet} \longrightarrow \tau_{\leqslant s} X^{\bullet}.$$

Taking totalisations, one therefore has a natural fibre sequence

$$\mathrm{Tot}(\tau_{\geqslant s+1} X^{\bullet}) \longrightarrow \mathrm{Tot}\, X^{\bullet} \longrightarrow \mathrm{Tot}(\tau_{\leqslant s} X^{\bullet}).$$

Taking colimits over s, one has a natural cofibre sequence

$$\operatorname*{colim}_s \mathrm{Tot}(\tau_{\geqslant s+1} X^{\bullet}) \longrightarrow \mathrm{Tot}\, X^{\bullet} \longrightarrow \operatorname*{colim}_s \mathrm{Tot}(\tau_{\leqslant s} X^{\bullet}).$$

By definition, the left-hand term gives the colimit of $\mathrm{Déc}^{\Delta}(X^{\bullet})$, 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 $\mathrm{Tot}\, \tau_{\leqslant s} X^{\bullet}$ 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 $\mathrm{Déc}\colon \mathrm{FilSp} \to \mathrm{FilSp}$ for Antieau's décalage functor.

Proposition 2.77. *Let $X^{\bullet}$ be a (semi)cosimplicial spectrum. Then there is an isomorphism, natural in $X^{\bullet}$, of filtered spectra*

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

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 Beĭlinson t-structure, as explained in detail in [Ant24] or [Hed20, Part II].

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

$$\cdots \longrightarrow \mathrm{filTot}(\tau_{\geqslant 1} X^\bullet) \longrightarrow \mathrm{filTot}(\tau_{\geqslant 0} X^\bullet) \longrightarrow \cdots \longrightarrow \mathrm{filTot}(X^\bullet). \tag{2.78}$$

We claim that this is the Whitehead filtration of $\mathrm{filTot}(X^\bullet)$ in the Beĭlinson t-structure on FilSp. First, we show that for every integer n, the filtered spectrum $\mathrm{filTot}(\tau_{\geqslant n} X^\bullet)$ is Beĭlinson n-connective. For this, we ought to show that its associated graded in filtration s is an $(n-s)$-connective spectrum. Using Proposition 2.71, we compute

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

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 $\mathrm{filTot}(\tau_{\geqslant n} X^\bullet) \to \mathrm{filTot}(X^\bullet)$ factors through a map

$$\mathrm{filTot}(\tau_{\geqslant n} X^\bullet) \longrightarrow \tau^{\mathrm{Bei}}_{\geqslant n} \, \mathrm{filTot}(X^\bullet).$$

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 2.71 and [Hed20, Proposition II.2.4].

We conclude that (2.78) is indeed the Beĭlinson Whitehead filtration of $\mathrm{filTot}\, X^\bullet$. From the definition of décalage, it follows that

$$\mathrm{D\acute{e}c}(\mathrm{filTot}\, X^\bullet) = \operatorname*{colim}_s \mathrm{Tot}^s(\tau_{\geqslant *} X^\bullet) \cong \mathrm{Tot}(\tau_{\geqslant *} X^\bullet) = \mathrm{D\acute{e}c}^\Delta \, X^\bullet. \qquad \blacksquare$$

Using this comparison result, the properties of décalage from Theorem 2.75 and Proposition 2.76 also follow from general properties of Antieau's décalage functor from [Ant24] and [Hed20, Part II].

Remark 2.79. The cosimplicial décalage functor $\mathrm{D\acute{e}c}^\Delta$ still has an advantage over the composite functor $\mathrm{D\acute{e}c} \circ \mathrm{filTot}$: the former is naturally lax symmetric monoidal (see Remark 2.74). Although the functor $\mathrm{D\acute{e}c}\colon \mathrm{FilSp} \to \mathrm{FilSp}$ is also lax symmetric monoidal, the functor filTot is not,[4] so the lax monoidal structure only arises through the version of Definition 2.73.

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

Remark 2.80 (Beĭlinson 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 Beĭlinson t-structure on a subcategory of FilSp. This is remarked by Lawson in [Law24a, Remark 3.16]. In more detail: recall from Remark 2.69 that the functor $\mathrm{filTot}\colon \mathrm{Sp}^{\Delta} \to \widehat{\mathrm{FilSp}}$ 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 Beĭlinson 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 2.81. Let E be a spectrum with a map $\mathbf{S} \to E$, and let X be a spectrum. The map $\mathbf{S} \to E$ gives rise to an augmented semicosimplicial spectrum $\Delta_{\mathrm{inj},+} \to \mathrm{Sp}$ of the form

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

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

$$\mathrm{ASS}_E^{\Delta}(X) := E^{[\bullet]} \otimes X = \Big(E \otimes X \rightrightarrows E \otimes E \otimes X \rightrightarrows\!\!\!\!\to \cdots \Big).$$

Define the filtered spectrum

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

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

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

and the filtered spectrum

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

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 $\mathbf{E}_1$-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 $\mathrm{ASS}_E^{\Delta}(Y, X)$ to a cosimplicial spectrum, which we will denote by the same notation. We would then define $\mathrm{ASS}_E(Y, X)$ using the cosimplicial Tot tower, which has a more efficient first page compared to the semicosimplicial approach; see Proposition 2.71.

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.

Variant 2.82. Alternatively, we could have defined the Adams spectral sequence as arising from the filtered spectrum

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

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 2.37.

Remark 2.83. 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 $\mathbf{E}_\infty$-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 2.84. 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 2.85 (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 $\mathbf{S} \to E$, let $\overline{E}$ denote the fibre of this map. The *filtered E-Adams resolution*[5] of the sphere is the filtered spectrum given by

$$\cdots \longrightarrow \overline{E} \otimes \overline{E} \longrightarrow \overline{E} \longrightarrow \mathbf{S} =\!=\!= \mathbf{S} =\!=\!= \cdots$$

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 2.67, the Tot tower of the cosimplicial spectrum $E^{[\bullet]} \otimes X$ is matched up with the associated tower (in the sense of Definition 2.55) of the filtered E-Adams resolution for X. Because the associated tower functor factors through completion (Proposition 2.59), 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 2.81, 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 2.86 ([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

$$E^* E_\alpha \longrightarrow \mathrm{Hom}_{E_*}(E_* E_\alpha,\, E_*) \tag{2.87}$$

is an isomorphism.

Remark 2.88. If E is an $\mathbf{E}_1$-ring, then the condition on a finite-projective E_α that the map (2.87) 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 $E_* E_\alpha$ is projective implies that the

[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.

[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.

resulting Künneth spectral sequence computing $E^* E_\alpha$ is concentrated in filtration 0, implying that (2.87) 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 2.89. 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 2.90.

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

(2) If E is $\mathbf{F}_p$, or more generally if $\pi_* E$ 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 $\pi_* E$ 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 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 $E = \mathrm{MU}$ and $E = \mathrm{BP}$ 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 2.89 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 2.91. Let (A, Γ) be a graded Hopf algebroid. We write $\mathrm{grComod}_{(A,\Gamma)}$ 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 $\mathrm{grComod}_{(A,\Gamma)}$ 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 2.92. *Let (A, Γ) be a graded Hopf algebroid. Suppose that (A, Γ) is flat. Then the category* $\mathrm{grComod}_{(A,\Gamma)}$ *is a Grothendieck abelian category, and the forgetful functor to* grAb *preserves small colimits and finite limits* (*in particular, it is exact*). *In particular,* $\mathrm{grComod}_{(A,\Gamma)}$ *has enough injectives.*

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

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

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 $\mathrm{grComod}_{(A,\Gamma)}$ is Grothendieck abelian also. ∎

The forgetful functor $\mathrm{grComod}_{(A,\Gamma)} \to \mathrm{grMod}_A$ 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 $\mathrm{grComod}_{(A,\Gamma)}$ is merely an additive category.

Notation 2.93. 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

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

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

Construction 2.94. 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 2.89. This gives us a good symmetric monoidal abelian category of comodules over (E_*, E_*E). We will abbreviate this category by

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

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 2.95. Let E be a homotopy-associative ring spectrum. Then E being of Adams type implies that E_*E is flat; see Remark 2.89. 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 2.81. In some cases, such as when Y is a sphere, these agree; see Proposition 2.99 below. We use the following (nonstandard) names to distinguish between them.

Definition 2.96. Let X, Y and E be spectra.

(1) The **algebraic *E*-Adams filtration** on $[Y, X]$ is the strict filtration where, for every $s \geqslant 1$, 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 $s \geqslant 1$, 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 2.97. *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 $\mathrm{ASS}_E(Y, X)$ is the topological E-Adams filtration.*

Proof. We freely use the notation introduced in Remark 2.85. By loc. cit., we may compute the induced strict filtration on $[Y, X]_*$ using the Adams resolution

$$\cdots \longrightarrow \overline{E} \otimes \overline{E} \otimes X \longrightarrow \overline{E} \otimes X \longrightarrow X. \tag{2.98}$$

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

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

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 (2.98) becomes null after tensoring with E. As a result, a map $\Sigma^n Y \to X$ that lifts to $\overline{E}^{\otimes s} \otimes X$ for $s \geqslant 1$ is of topological Adams filtration $\geqslant s$. For the converse, it is enough to consider the case $s = 1$ and $n = 0$, i.e., we have a map $f\colon Y \to X$ that becomes nullhomotopic after tensoring with E. Then also the composite $Y \to E \otimes Y \to E \otimes X$ is nullhomotopic, so it follows that there exists a dashed factorisation

$$\begin{tikzcd} & Y \arrow[dl, dashed] \arrow[d, "f"] & \\ \overline{E} \otimes X \arrow[r] & X \arrow[r] & E \otimes X, \end{tikzcd}$$

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

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 2.99. *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 $\pi_* E$ is a graded field, or if $Y = \mathbf{S}^n$). Then the following hold.*

(1) *There is an isomorphism*

$$\mathrm{E}_2^{n,s}(Y, X) \cong \mathrm{Ext}^{s,\, n+s}_{E_* E}(E_* Y,\, E_* X).$$

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]. ■

Remark 2.100. 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–Pstrągowski [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 $\mathrm{ASS}_E(X)$ is a complete filtration by Remark 2.69, so it might appear as if there are no convergence problems. The issue, however, is whether the underlying spectrum of $\mathrm{ASS}_E(X)$ 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 2.101. Let X be a spectrum, and let E be spectrum with a map $\mathbf{S} \to E$. The ***E*-nilpotent completion** of X is the spectrum

$$X^\wedge_E := \mathrm{Tot}(E^{[\bullet]} \otimes X).$$

The map $\mathbf{S} \to E^{[\bullet]}$ induces a natural map $X \to X^\wedge_E$. 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 $\mathrm{map}(Y, X^\wedge_E)$. If Y is the sphere, then this is simply $X^\wedge_E$. In both cases then, the question of convergence is whether the map $X \to X^\wedge_E$ 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 $X \to X^\wedge_E$ factors through a map

$$L_E X \longrightarrow X^\wedge_E. \tag{2.102}$$

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 2.103 (Bousfield). *Let E and X be as above. Then the map* (2.102) *is an isomorphism if and only if the natural map $X^\wedge_E \to (X^\wedge_E)^\wedge_E$ 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 2.104 (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 $L_E X \to X^\wedge_E$ 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 $L_E X \to X^\wedge_E$ is an isomorphism, and moreover*
$$L_E X \cong X^\wedge_n.$$

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

Remark 2.105. 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 $L_E \mathbf{Z} \to \mathbf{Z}^\wedge_E$ is not an isomorphism.

Example 2.106. The ring MU satisfies the first condition for $J = \varnothing$, so that for all spectra X that are bounded below, we have

$$L_{\mathrm{MU}} X \cong X^\wedge_{\mathrm{MU}} \cong X.$$

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

$$L_{\mathrm{BP}} X \cong X^\wedge_{\mathrm{BP}} \cong X_{(p)}.$$

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

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

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. There, we discuss strengthened versions of the ordinary Leibniz rule, which will play a major role in Part II of this thesis.

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

Proposition 3.1. *The functor*

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

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

Mathematically, there is nothing deep about this statement. The value is in the human aspect: it can be less mentally taxing to think in terms of algebraic equations involving τ, than it is to picture the diagram that is a filtered abelian group. Even for 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** $\mathrm{Const}\colon \mathrm{Ab} \to \mathrm{FilAb}$ is given by sending an abelian group A to the constant filtration on A, given by

$$\cdots = A = A = A = \cdots .$$

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 $\mathrm{FilAb}[\tau^{-1}]$ for the full subcategory on these objects.

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

$$\begin{tikzcd} \mathrm{FilAb} \arrow[r, bend left, "\tau^{-1}"] \arrow[r, bend right] & \mathrm{FilAb}[\tau^{-1}] \arrow[l, hook'] \end{tikzcd} \qquad \textit{become} \qquad \begin{tikzcd} \mathrm{FilAb} \arrow[r, bend left, "A \mapsto A^{-\infty}"] \arrow[r, bend right, "A \mapsto A^{\infty}"'] & \mathrm{Ab}. \arrow[l, hook', "\mathrm{Const}"'] \end{tikzcd}$$

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 $\mathrm{Ab} \simeq \mathrm{FilAb}[\tau^{-1}]$. ∎

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 $\tau^{-1}\colon \mathrm{FilAb} \to \mathrm{FilAb}[\tau^{-1}]$ for the functor sending A to

$$A[\tau^{-1}] := \operatorname{colim}\big(A \xrightarrow{\tau} A \xrightarrow{\tau} \cdots \big).$$

We write $(-)^{\tau=1}\colon \mathrm{FilAb} \to \mathrm{Ab}$ for the composite

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

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

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

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

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

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

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

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*

$$M \longrightarrow M^{\wedge}_{x} := \lim_{k} M/x^k$$

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

$$\begin{tikzcd}
\vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] \\
M \arrow[r, "x^3"] \arrow[d, "x"'] & M \arrow[r] \arrow[d, equal] & M/x^3 \arrow[d] \\
M \arrow[r, "x^2"] \arrow[d, "x"'] & M \arrow[r] \arrow[d, equal] & M/x^2 \arrow[d] \\
M \arrow[r, "x"] & M \arrow[r] & M/x.
\end{tikzcd}$$

[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.

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

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

These are precisely the limit and first-derived limit of the x-adic filtration on M. In other words, we see that $M \to M_x^\wedge$ 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. ■

Variant 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 $\mathbf{Z}^{\mathrm{op}} \to \mathrm{grAb}$. Note that this by definition means that the transition maps preserve degrees. We write $\mathrm{FilgrAb} := \mathrm{Fun}(\mathbf{Z}^{\mathrm{op}}, \mathrm{grAb})$ 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

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

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

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

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

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

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

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 $\widetilde{\eta}$ or $\widetilde{\nu}$, but instead use the symbols η and ν to directly record the filtration of the elements in the ring A. We 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 $\pi_{n,s} \to \pi_{n,s-1}$ is induced by a map $\mathbf{S}^{n,s-1} \to \mathbf{S}^{n,s}$.

Definition 3.13.

(1) The map $\tau\colon \mathbf{S}^{0,-1} \to \mathbf{S}$ is the image of the morphism $-1 \to 0$ in $\mathbf{Z}$ under the functor $i\colon \mathbf{Z} \to \mathrm{FilSp}$ from Definition 2.17.

(2) If X is a filtered spectrum, then tensoring $\tau\colon \mathbf{S}^{0,-1} \to \mathbf{S}$ with X results in a map $\Sigma^{0,-1}X \to X$. 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

$$\begin{tikzcd} \cdots \arrow[r] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & \mathbf{S} \arrow[r, equal] \arrow[d, equal] & \cdots \\ \cdots \arrow[r] & 0 \arrow[r] & \mathbf{S} \arrow[r, equal] & \mathbf{S} \arrow[r, equal] & \cdots . \end{tikzcd}$$

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

$$\begin{tikzcd} \cdots \arrow[r] & X^2 \arrow[r, "f_1"] \arrow[d, "f_1"] & X^1 \arrow[r, "f_0"] \arrow[d, "f_0"] & X^0 \arrow[r] \arrow[d, "f_{-1}"] & \cdots \\ \cdots \arrow[r] & X^1 \arrow[r, "f_0"] & X^0 \arrow[r, "f_{-1}"] & X^{-1} \arrow[r] & \cdots . \end{tikzcd}$$

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

Remark 3.14. Using the above notation, the functor $i\colon \mathbf{Z} \to \mathrm{FilSp}$ 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

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

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 $\tau\colon\ \Sigma^{0,-1}\, X \to X$ is an isomorphism. We write $\mathrm{FilSp}[\tau^{-1}]$ for the full subcategory of FilSp on the τ-invertible filtered spectra. Write $\tau^{-1}\colon \mathrm{FilSp} \to \mathrm{FilSp}$ for the functor sending X to the colimit

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

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

$$\mathrm{FilSp} \underset{}{\overset{\tau^{-1}}{\rightleftarrows}} \mathrm{FilSp}[\tau^{-1}].$$

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 structure.

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 $\mathrm{FilSp}[\tau^{-1}] \subseteq \mathrm{FilSp}$ preserves colimits and has a further right adjoint.*

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

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

■

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* Const: Sp $\to$ FilSp *restricts to an equivalence onto the τ-invertible filtered spectra:*

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

Under this equivalence, the adjunctions

$$\begin{tikzcd} \mathrm{FilSp} \arrow[r, bend left, "\tau^{-1}"] \arrow[r, bend right] & \mathrm{FilSp}[\tau^{-1}] \arrow[l, hook'] \end{tikzcd} \qquad \textit{become} \qquad \begin{tikzcd} \mathrm{FilSp} \arrow[r, bend left, "X \mapsto X^{-\infty}"] \arrow[r, bend right, "X \mapsto X^{\infty}"'] & \mathrm{Sp}. \arrow[l, hook', "\mathrm{Const}"'] \end{tikzcd}$$

Paralleling Notation 3.4, we will sometimes use the following notation to distinguish between the two equivalent ∞-categories Sp and 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

$$\mathrm{FilSp} \xrightarrow{\tau^{-1}} \mathrm{FilSp}[\tau^{-1}] \xrightarrow{\simeq} \mathrm{Sp}.$$

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 2.33) 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 $\mathrm{E}_1^{n,s}$ not being zero, which is implied by the definition of detection from Definition 2.33.

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 $i(-)[\tau^{-1}] \cong \operatorname{Const}\mathbf{S}$. As τ-inversion and Const are exact functors, they preserve suspensions, so we find that

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

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 $k \geqslant 1$.

- We write $C\tau^k$ for the cofibre of the map $\tau^k\colon \mathbf{S}^{0,-k} \to \mathbf{S}$.
- 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

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

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

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 $k \geqslant 0$, the filtered spectrum $C\tau^k$ admits (uniquely up to contractible choice) the structure of a filtered $\mathbf{E}_\infty$-ring such that its unit map $\mathbf{S} \to C\tau^k$ is an isomorphism in filtrations $0, -1, \ldots, -k+1$.*

Proof. See [Lur15, Proposition 3.2.5], bearing in mind that Lurie writes $\operatorname{Rep}(\mathbf{Z})$ 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

$$d\colon \mathrm{grSp} \longrightarrow \mathrm{FilSp}$$

given by left Kan extension along the functor $\mathbf{Z}^{\mathrm{discr}} \to \mathbf{Z}^{\mathrm{op}}$, with $\mathbf{Z}^{\mathrm{discr}}$ 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

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

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*

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

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

$$\begin{tikzcd} \mathrm{FilSp} \arrow[r, "\mathrm{Gr}"] \arrow[dr, "C\tau\otimes -"'] & \mathrm{grSp} \arrow[d, "\simeq"] \\ & \mathrm{Mod}_{C\tau}(\mathrm{FilSp}). \end{tikzcd}$$

Proof. See [Lur15, Proposition 3.2.7], bearing in mind that Lurie writes $\mathrm{Rep}(\mathbf{Z})$ for FilSp, writes $\mathrm{Rep}(\mathbf{Z}^{\mathrm{ds}})$ 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 $\mathrm{Gr}\colon \mathrm{FilSp} \to \mathrm{grSp}$, because the functor $C\tau \otimes - \colon \mathrm{FilSp} \to \mathrm{Mod}_{C\tau}(\mathrm{FilSp})$ 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 $\mathrm{FilSp}^{\wedge}_{\tau}$ for the full subcategory of FilSp on the τ-adically complete filtered spectra. This results in an adjunction

$$\mathrm{FilSp} \underset{\hookleftarrow}{\xrightarrow{(-)^{\wedge}_{\tau}}} \mathrm{FilSp}^{\wedge}_{\tau}.$$

This notion is, in fact, nothing but the notion of completeness previously introduced in Definition 2.26.

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 $X \to \mathrm{cofib}(X^{\infty} \to X)$ is the τ-completion of X.*

Proof. By definition, a filtered spectrum X is τ-complete if and only if $\mathrm{Map}(Y, X)$ 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 $\mathrm{cofib}(X^{\infty} \to X)$ 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*

$$\begin{tikzcd} X \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & X^{\wedge}_{\tau} \arrow[d] \\ X[\tau^{-1}] \arrow[r] & (X^{\wedge}_{\tau})[\tau^{-1}]. \end{tikzcd}$$

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

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

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 limit and on the associated graded. As such, Proposition 3.31 can be thought of as the stable analogue of Proposition 2.2.

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 2.41 is asking $\pi_{*,*}X$ to be τ-complete. As explained in Warning 2.49, the τ-completeness of X (i.e., the vanishing of X^{∞}) need not imply this. The conditional convergence criteria of Boardman from Theorem 2.52 and Remark 2.54 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 $n \geqslant 1$. 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 $N \geqslant n$, 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 $d_n, \ldots, d_{N-1}$-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 $d_1, \ldots, d_{n-1}$ 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

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

where the first map is the transition map. Recall from Appendix C that the boundary map $\mathrm{Gr}^s \to \Sigma X^{s+1}$ 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 $n \geqslant k \geqslant 1$, and let $\infty \geqslant N \geqslant n$. Then we have commutative diagrams*

$$\begin{tikzcd} X/\tau^n \arrow[r, "\partial_n^\infty"] \arrow[dr, "\partial_n^N"'] & \Sigma^{1,-n} X \arrow[d] \\ & \Sigma^{1,-n} X/\tau^{N-n} \end{tikzcd} \quad \textit{and} \quad \begin{tikzcd} X/\tau^n \arrow[r, "\partial_n^\infty"] \arrow[d] & \Sigma^{1,-n} X \arrow[d, "\tau^{n-k}"] \\ X/\tau^k \arrow[r, "\partial_k^\infty"] & \Sigma^{1,-k} X, \end{tikzcd}$$

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

$$\begin{tikzcd} X \arrow[r, equal] \arrow[d, "\tau^{N-n}"'] & X \arrow[d, "\tau^N"] \\ X \arrow[r, "\tau^n"] & X. \end{tikzcd}$$

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

$$\begin{tikzcd} X \arrow[r, "\tau^n"] \arrow[d] & X \arrow[r] \arrow[d] & X/\tau^n \arrow[r, "\partial_n^\infty"] \arrow[d, equal] & X \arrow[d] \\ X/\tau^{N-n} \arrow[r, "\tau^n"] & X/\tau^N \arrow[r] & X/\tau^n \arrow[r, "\partial_n^N"] & X/\tau^{N-n}. \end{tikzcd}$$

The right-most square is the first claimed diagram.

The second diagram comes from the commutative diagram

$$\begin{tikzcd} X \arrow[r, "\tau^n"] \arrow[d, "\tau^{n-k}"'] & X \arrow[r] \arrow[d, equal] & X/\tau^n \arrow[r, "\partial_n^\infty"] \arrow[d] & X \arrow[d, "\tau^{n-k}"] \\ X \arrow[r, "\tau^k"] & X \arrow[r] & X/\tau^k \arrow[r, "\partial_k^\infty"] & X \end{tikzcd}$$

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 $k \geqslant m \geqslant 0$ 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

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

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

$$\begin{tikzcd} \Sigma^{0,-m}\, X/\tau^k \arrow[r, "\tau^m"] \arrow[d] & X/\tau^k \\ \Sigma^{0,-m}\, X/\tau^{k-m} \arrow[ur, "\tau^m"'] & \end{tikzcd}$$

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 θ mod τ^{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 $x \in \mathrm{E}_1^{n,s}$. View x as an element of $\pi_{n,s}\, X/\tau$. Let $r \geqslant 0$.*

(1a) *If $\partial_1^\infty(x)$ is τ^r-divisible, then $d_1(x), \ldots, d_r(x)$ 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 $d_1(x), \ldots, d_r(x)$ 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 $\mathrm{E}_1^{n-1,s+r+1}$.*

(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

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

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 $n \geqslant 1$, and let $n \leqslant N \leqslant \infty$. 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 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 $n \geqslant 1$, the map*

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

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.

We will see both of these results in action in a myriad of cases in Chapter 9.

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 $x, y \in \mathrm{E}_1^{*,*}$ 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 $t \in \mathrm{E}_1^{*,*}$ for its mod τ reduction, and suppose that we have a multiplicative relation in $\mathrm{E}_1^{*,*}$ 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

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

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 the proof of Corollary C in Chapter 10. In Tmf, there is a 2-extension in the 110-stem:

$$2 \cdot \kappa_4 = \eta_1^2 \bar{\kappa}^3. \tag{3.43}$$

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

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

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 this thesis. 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 $\mathrm{BF}_\tau\, X$, is the bifiltered spectrum $\mathbf{Z}^{\mathrm{op}} \to \mathrm{FilSp}$ given by

$$\cdots \xrightarrow{\ \tau\ } \Sigma^{0,-2}\,X \xrightarrow{\ \tau\ } \Sigma^{0,-1}\,X \xrightarrow{\ \tau\ } X =\!=\!= \cdots$$

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 filtrations 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 2.28. 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 $X\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{FilSp}$ be a bifiltered spectrum. We define a trigraded exact couple

$$\mathrm{A}^{n,w,s} = \pi_{n,w}(X^s) \qquad \text{and} \qquad \mathrm{E}^{n,w,s} = \pi_{n,w}(\mathrm{Gr}^s X),$$

fitting into the following diagram, where each map is annotated by its tridegree.

$$\begin{tikzcd} \mathrm{A}^{n,w,s} \arrow[rr, "{(0,0,-1)}"] & & \mathrm{A}^{n,w,s} \arrow[dl, "{(0,0,0)}"] \\ & \mathrm{E}^{n,w,s} \arrow[ul, "{(-1,0,1)}"] & \end{tikzcd}$$

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

$$\mathrm{E}_1^{n,w,s} = \pi_{n,w}(\mathrm{Gr}^s X) \implies \pi_{n,w} X^{-\infty}.$$

The induced strict filtration on $\pi_{n,w} X^{-\infty}$ is

$$F^s \pi_{n,w} X^{-\infty} = \mathrm{im}(\pi_{n,w} X^s \longrightarrow \pi_{n,w} X^{-\infty}).$$

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

$$\pi_{n,w}(\mathrm{BF}_\tau^s X) = \pi_{n,w}(\Sigma^{0,-s} X) = \pi_{n,\,w+s} X \qquad \text{and} \qquad \mathrm{Gr}^s(\mathrm{BF}_\tau X) = \Sigma^{0,-s} X/\tau.$$

In particular, the first page is of the form

$$\mathrm{E}_1^{n,w,s} = \begin{cases} \pi_{n,\,w+s}(X/\tau) & \text{if } s \geqslant 0, \\ 0 & \text{else.} \end{cases} \tag{3.49}$$

The induced strict filtration on $\pi_{*,*} X$ is

$$F^s \pi_{n,w} X = \mathrm{im}(\tau^s\colon\ \pi_{n,\,w+s} X \longrightarrow \pi_{n,w} X). \qquad \blacktriangle$$

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 $\overline{\tau}$ to have tridegree $(0, -1, 1)$. Let X be a filtered spectrum, and $\{ \mathrm{E}_r^{*,*,*} \}$ denote its τ-BSS. We put a $\mathbf{Z}[\overline{\tau}]$-module structure on $\mathrm{E}_1^{*,*,*}$ in such a way that the isomorphism (3.49) becomes an isomorphism trigraded $\mathbf{Z}[\overline{\tau}]$-modules

$$\mathrm{E}_1^{*,*,*} \cong \mathbf{Z}[\overline{\tau}] \otimes \pi_{*,*}(X/\tau),$$

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 differentials 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

$$d_r^\tau(x) = \overline{\tau}^r \cdot y.$$

The insertion of this $\overline{\tau}$-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 $\overline{\tau}$-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 $d_{\leqslant r-1}$-cycles, implicitly considering these elements as defining classes in E_r.

Theorem 3.52. *Let X be a filtered spectrum. Let $\{ \mathrm{E}_r^{*,*}, d_r \}$ denote the spectral sequence underlying X, and let $\{ \mathrm{E}_r^{*,*,*}, d_r^\tau \}$ denote the τ-Bockstein spectral sequence of X.*

(1) *There is a natural isomorphism of trigraded $\mathbf{Z}[\overline{\tau}]$-modules*

$$\mathrm{E}_1^{*,*,*} \cong \mathbf{Z}[\overline{\tau}] \otimes \pi_{*,*}(X/\tau) = \mathbf{Z}[\overline{\tau}] \otimes \mathrm{E}_1^{*,*},$$

where $\pi_{n,w}(X/\tau)$ is placed in tridegree $(n,w,0)$ and $\overline{\tau}$ has tridegree $(0,-1,1)$.

(2) *The differentials are $\overline{\tau}$-linear and $\overline{\tau}$-divisible: for $x \in \mathrm{E}_1^{n,w,s}$ and $y \in \mathrm{E}_1^{n-1,w,s+r}$, there is a differential*

$$d_r^{\tau}(x) = y,$$

if and only if for any (hence all) $m \geqslant 0$, there is a differential

$$d_r^{\tau}(\overline{\tau}^m x) = \overline{\tau}^m y.$$

In particular, the $\mathbf{Z}[\overline{\tau}]$-module structure on the first page induces a $\mathbf{Z}[\overline{\tau}]$-module structure on later pages.

(3) *The target of a d_r^{τ}-differential is a $\overline{\tau}^r$-multiple: for every $x \in \mathrm{E}_r^{n,w,s}$, there is an element $y \in \mathrm{E}_r^{n-1,w+r,s}$ such that*

$$d_r^{\tau}(x) = \overline{\tau}^r \cdot y,$$

where the multiplication denotes the $\mathbf{Z}[\overline{\tau}]$-module structure on $\mathrm{E}_r^{*,*,*}$ from (2).

(4) *If $x \in \mathrm{E}_1^{n,w}$ and $y \in \mathrm{E}_r^{n-1,w+r}$, then there is a differential*

$$d_r(x) = y$$

if and only if there is a differential

$$d_r^{\tau}(x) = \overline{\tau}^r \cdot y.$$

(5) *Suppose that $x \in \mathrm{E}_1^{n,w,s}$ detects an element $\theta \in \pi_{n,w} X$. Then θ is τ^s-divisible. Moreover, for every $m \geqslant 0$, the class $\overline{\tau}^m \cdot x$ 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 $\mathrm{RE}_{\infty}^{n,s}$ (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 $\mathbf{Z}[\overline{\tau}]$-module structure on the first page, so that the statement of $\overline{\tau}$-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 $\mathbf{Z}[\overline{\tau}]$-module structure, so that we can talk about $\overline{\tau}$-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

$$\mathrm{E}_2^{n,w,s} \cong \mathrm{E}_2^{n,w+s}$$

  for all $s \geqslant 1$. Since the d_2^{τ}-differential only hits filtrations $s \geqslant 2$, 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 $\overline{\tau}$ 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 $\overline{\tau} \cdot x$ 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 $\overline{\tau}$-power multiples of differentials originating from filtration 0. As a result, both the differential $d_r^\tau(x) = y$ and $d_r^\tau(\overline{\tau} x) = \overline{\tau} y$ 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 $\overline{\tau}$ induces a surjection

$$\overline{\tau} \colon \mathrm{E}_r^{n,w,s} \twoheadrightarrow \mathrm{E}_r^{n,\, w-1,\, s+1}.$$

This claim, in turn, we prove by induction on r. For $r = 1$ it is clear from item (1). Assume that for some $r \geqslant 1$, multiplication by $\overline{\tau}$ induces a surjection $\mathrm{E}_r^{n,w,s} \twoheadrightarrow \mathrm{E}_r^{n,\, w-1,\, s+1}$. Write $(\ker d_r)^{n,w,s}$ for the kernel of the d_r-differential out of $\mathrm{E}_r^{n,w,s}$, and write $(\mathrm{im}\, d_r)^{n,w,s}$ for the image of d_r into $\mathrm{E}_r^{n,w,s}$. Then we have

$$\mathrm{E}_{r+1}^{n,w,s} = \frac{(\ker d_r)^{n,w,s}}{(\mathrm{im}\, d_r)^{n,w,s}} \qquad \text{and} \qquad \mathrm{E}_{r+1}^{n,\, w-1,\, s+1} = \frac{(\ker d_r)^{n,\, w-1,\, s+1}}{(\mathrm{im}\, d_r)^{n,\, w-1,\, s+1}}.$$

The $\overline{\tau}$-divisibility of the differentials from (2) implies that multiplication by $\overline{\tau}$ restricts to a surjection $(\ker d_r)^{n,w,s} \to (\ker d_r)^{n,\, w-1,\, s+1}$, 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:

$$\cdots \longrightarrow X^{w_0+2} \longrightarrow X^{w_0+1} \longrightarrow X^{w_0} = X^{w_0} = \cdots,$$

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 $\mathrm{BF}_\tau X\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{FilSp}$ 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

$$\mathrm{A}^{n,w_0,s}(\mathrm{BF}_\tau X) \xrightarrow{\cong} \mathrm{A}^{n,w_0+s}(Y) \qquad \mathrm{E}^{n,w_0,s}(\mathrm{BF}_\tau X) \xrightarrow{\cong} \mathrm{E}^{n,w_0+s}(Y)$$

where $s \geqslant 0$. 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 2.33 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 $\pi_{n,w} \mathrm{BF}_\tau^s X \to \pi_{n,w} \mathrm{BF}_\tau^0 X$. 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 $\mathbf{E}_\infty$-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 $\overline{\tau} \cdot x \in \mathrm{E}_1^{n,w-1,s+1}$. 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 $\tau^{s+1}\colon \pi_{n,w+s} X \to \pi_{n,w-1} X$, so indeed $\overline{\tau} \cdot x$ 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 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 2.52, 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 $s \geqslant 0$)

$$\mathrm{RE}_\infty^{n,w,s} = \lim_r{}^1 \mathrm{Z}_r^{n,w,s} \cong \lim_r{}^1 \mathrm{Z}_r^{n,w+s} = \mathrm{RE}_\infty^{n,w+s}. \qquad \blacksquare$$

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 $s \geqslant 0$)

$$\mathrm{Z}_r^{n,w,s} \cong \mathrm{Z}_r^{n,w+s} \qquad \mathrm{B}_r^{n,w,s} \cong \mathrm{B}_{\min(r,s)}^{n,w+s} \qquad \mathrm{E}_{r+1}^{n,w,s} \cong \mathrm{Z}_r^{n,w+s}/\mathrm{B}_{\min(r,s)}^{n,w+s}.$$

In particular, we have isomorphisms

$$\mathrm{Z}_\infty^{n,w,s} \cong \mathrm{Z}_\infty^{n,w+s} \qquad \mathrm{B}_\infty^{n,w,s} \cong \mathrm{B}_s^{n,w+s} \qquad \mathrm{E}_\infty^{n,w,s} \cong \mathrm{Z}_\infty^{n,w+s}/\mathrm{B}_s^{n,w+s}.$$

[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 trigraded abelian groups, where these conditions are certainly met, and Boardman's arguments go through without any change (merely tagging on an additional grading).

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 2.54 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 2.49. 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^τ-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

$$\widetilde{\mathrm{E}}_{r+1}^{n,w,s} := \mathrm{E}_r^{n,\,w,\,s+w}.$$

In this indexing, $\overline{\tau}$ has tridegree $(0,-1,0)$, and the non-$\overline{\tau}$-divisible groups are located in tridegrees of the form (n,w,w).

Warning 3.56. The $\overline{\tau}$-divisions appearing in items (2) and (3) are not necessarily unique, due to the $\overline{\tau}$-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 $k \geqslant 1$. 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 $k \geqslant 1$, we have a morphism of bifiltered spectra $\mathrm{BF}_\tau X \to \mathrm{tr}_k \mathrm{BF}_\tau X$ of the form

$$\begin{tikzcd} \cdots \arrow[r, "\tau"] & \Sigma^{0,-k} X \arrow[r, "\tau"] \arrow[d] & \Sigma^{0,-k+1} X \arrow[r, "\tau"] \arrow[d] & \Sigma^{0,-k+2} X \arrow[r, "\tau"] \arrow[d] & \cdots \arrow[r, "\tau"] & X \arrow[d] \\ \cdots \arrow[r] & 0 \arrow[r] & \Sigma^{0,-k+1} X/\tau \arrow[r, "\tau"] & \Sigma^{0,-k+2} X/\tau^2 \arrow[r, "\tau"] & \cdots \arrow[r, "\tau"] & X/\tau^k \end{tikzcd}$$

constructed by letting each nontrivial square be a pushout. Formally, it is left Kan extended from the subdiagram

$$\begin{tikzcd} \cdots \arrow[r, "\tau"] & \Sigma^{0,-k-1} X \arrow[r, "\tau"] \arrow[d] & \Sigma^{0,-k} X \arrow[r, "\tau"] \arrow[d] & \cdots \arrow[r, "\tau"] & X. \\ \cdots \arrow[r] & 0 \arrow[r] & 0 \end{tikzcd}$$

Likewise, for $m \geqslant k \geqslant 1$, we have a morphism $\mathrm{tr}_m \mathrm{BF}_\tau X \to \mathrm{tr}_k \mathrm{BF}_\tau X$, 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 $k \geqslant 1$. Let $\{ \mathrm{E}_r^{*,*}, d_r \}$ denote the spectral sequence underlying X, and let $\{ \mathrm{E}_r^{*,*,*}, d_r^\tau \}$ denote the k-truncated τ-Bockstein spectral sequence of X.*

(1) *The map on spectral sequences induced by $\mathrm{BF}_\tau X \to \mathrm{tr}_k \mathrm{BF}_\tau X$ from Construction 3.59 is, on first pages, given by the quotient*

$$\mathbf{Z}[\overline{\tau}] \otimes \pi_{*,*}(X/\tau) \twoheadrightarrow \mathbf{Z}[\overline{\tau}]/\overline{\tau}^k \otimes \pi_{*,*}(X/\tau).$$

For $m \leqslant k$, the map on spectral sequences induced by $\mathrm{tr}_k \mathrm{BF}_\tau X \to \mathrm{tr}_m \mathrm{BF}_\tau X$ from Construction 3.59 is, on first pages, given by the quotient

$$\mathbf{Z}[\overline{\tau}]/\tau^k \otimes \pi_{*,*}(X/\tau) \twoheadrightarrow \mathbf{Z}[\overline{\tau}]/\overline{\tau}^m \otimes \pi_{*,*}(X/\tau).$$

(2) *The differentials are $\overline{\tau}$-linear: for $x \in \mathrm{E}_1^{n,w,s}$ and $y \in \mathrm{E}_1^{n-1,w,s+r}$, if there is a differential*

$$d_r^\tau(x) = y,$$

then for all $m \geqslant 0$, there is a differential

$$d_r^\tau(\overline{\tau}^m x) = \overline{\tau}^m y.$$

In particular, the $\mathbf{Z}[\overline{\tau}]$-module structure on the first page induces a $\mathbf{Z}[\overline{\tau}]$-module structure on later pages.

(3) *The target of a d_r^τ-differential is a $\overline{\tau}^r$-multiple: for every $x \in \mathrm{E}_r^{n,w,s}$, there is an element $y \in \mathrm{E}_r^{n-1,w+r,s}$ such that*

$$d_r^\tau(x) = \overline{\tau}^r \cdot y,$$

where the multiplication denotes the $\mathbf{Z}[\overline{\tau}]$-module structure on $\mathrm{E}_r^{*,*,*}$ from (2).

(4) *Let $r \leqslant k-1$. If $x \in \mathrm{E}_1^{n,w}$ and $y \in \mathrm{E}_r^{n-1,w+r}$, then there is a differential*

$$d_r(x) = y$$

if and only if there is a differential

$$d_r^\tau(x) = \overline{\tau}^r \cdot y.$$

(5) *Suppose that $x \in \mathrm{E}_1^{n,w,s}$ detects an element $\theta \in \pi_{n,w} X/\tau^k$. Then θ is in the image of the map*

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

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 $s+m \leqslant k-1$, the class $\overline{\tau}^m \cdot x$ detects $\tau^m \cdot \theta$. For every m such that $s+m \geqslant k$, 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 $\overline{\tau}$-divisible in general, and that item (4) only applies to differentials of length $r \leqslant k-1$ (longer d_r^τ-differentials vanish for degree reasons).

Proof. The definition of the map $\mathrm{BF}_\tau X \to \mathrm{tr}_k \mathrm{BF}_\tau X$ from Construction 3.59 as a left Kan extension shows that in filtrations $0 \leqslant s \leqslant k-1$, 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 $s \geqslant k$, so that (the trigraded analogue of) Proposition 2.53 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

$$\operatorname{tr}_k \mathrm{BF}^s_\tau X = \Sigma^{0,-s} X/\tau^{k-s} \longrightarrow X/\tau^k = \operatorname{tr}_k \mathrm{BF}^0_\tau X,$$

in other words, that $\tau^s(\alpha) = \theta$. If $s+m \leqslant k-1$, then

$$\pi_{n,w-m}(\mathrm{Gr}^{s+m} \operatorname{tr}_k \mathrm{BF}_\tau X) = \pi_{n,w}\, X/\tau,$$

and $\overline{\tau}^m \cdot x$ 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 $\overline{\tau}^m \cdot x$ to $\pi_{n,w-m}(\operatorname{tr}_k \mathrm{BF}^{s+m}_\tau X) = \pi_{n,w+s}\, X/\tau^{k-s-m}$. Evidently, the image of this element in $\operatorname{tr}_k \mathrm{BF}^0_\tau X$ 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 $s+m \geqslant k$, then this means $m \geqslant k-s$. 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. ∎

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

$$\mathrm{E}_1^{n,s} = \pi_{n,s}(C\tau \otimes X) \implies \pi_n(X^{\tau=1}).$$

In Appendix C, 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 $x \in \mathrm{E}_1^{n,s} = \pi_{n,s}(X/\tau)$ 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 $\pi_{n,s} X \to \pi_n X^{\tau=1}$.*

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 $\mathrm{BF}_\tau X \to \mathrm{tr}_1 \mathrm{BF}_\tau X$ from Construction 3.59 is, on underlying objects, equal to the reduction map $X \to X/\tau$. Meanwhile the induced map on the first page of the resulting spectral sequences is given by quotienting by $\overline{\tau}$:

$$\mathbf{Z}[\overline{\tau}] \otimes \pi_{*,*}(X/\tau) \twoheadrightarrow \pi_{*,*}(X/\tau),$$

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 $\overline{\tau}^{r-1} \cdot x$ detects $\tau^{r-1} \cdot \alpha$. To show that $\tau^{r-1} \cdot \alpha$ is nonzero, by (the trigraded analogue of) Remark 2.43 we need to show that $\overline{\tau}^{r-1} \cdot x$ survives to page ∞. Again using Theorem 3.52 (4), we see that $\overline{\tau}^{r-1} \cdot x$ 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 $\overline{\tau}^{r-1} \cdot x$ also survives to page ∞.

Properties (2b) and (3b) are a restatement of the definition of detection from Definition 2.33, with (2b) specifically following from Remark 2.43.

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 $d_r^\tau(y) = \overline{\tau}^r \cdot x$. 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.

$$\begin{tikzcd}
\Sigma^{0,-r} X \arrow[r, "\tau^{r-1}"] \arrow[d] & \Sigma^{0,-1} X \arrow[r, "\tau"] & X \arrow[d] \\
\Sigma^{0,-r} X/\tau & & X/\tau \arrow[ul, dashed, "\partial_1^\infty"]
\end{tikzcd}$$

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 $M_\tau^\wedge$ 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. ■

Warning 3.63. Suppose that $x \in \mathrm{E}_1^{n,s}$ 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 $d \geqslant 0$, 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 2.33).

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 $k \geqslant 1$. Let $x \in \mathrm{E}_1^{n,s} = \pi_{n,s}(X/\tau)$ be a class. Then the following are equivalent.*

(1a) *The differentials $d_1(x), \ldots, d_{k-1}(x)$ 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 $r \leqslant k$, then $\tau^{r-1} \cdot \alpha$ is nonzero.*

(2b) *The image of α under $\partial_k^{k+1}\colon\ \pi_{n,s}(X/\tau^k) \to \pi_{n-1,s+r}(X/\tau)$ 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 $\operatorname{tr}_k \mathrm{BF}_\tau X \to \operatorname{tr}_1 \mathrm{BF}_\tau X$ from Construction 3.59 is, on underlying objects, the reduction map $X/\tau^k \to X/\tau$. On first pages of the underlying spectral sequences on the other hand, the map takes the form of quotienting $\overline{\tau}$:

$$\mathbf{Z}[\overline{\tau}]/\overline{\tau}^k \otimes \pi_{*,*}\, X/\tau \longrightarrow\!\!\!\!\!\rightarrow \pi_{*,*}\, X/\tau.$$

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 $\overline{\tau}^{r-1} \cdot x$. The assumption that x survives to page r implies that $\overline{\tau}^{r-1} \cdot x$ is a permanent cycle that survives to page r, and for degree reasons it then also survives to page ∞. Using Remark 2.43, 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 $d_k^\tau(x)$ 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.

$$\begin{tikzcd}
\Sigma^{0,-k} X \arrow[r, "\tau^{k-1}"] \arrow[d] & \Sigma^{0,-1} X \arrow[r, "\tau"] & X \arrow[d] \\
\Sigma^{0,-k} X/\tau & & X/\tau \arrow[ul, dashed, "\partial_1^\infty"]
\end{tikzcd}$$

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 $\partial_1^\infty(a)$, so that its projection to $\pi_{n,s+k}\, X/\tau$ is a representative for $d_k^\tau(x)$. 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 $\overline{\tau}^{r-1} \cdot x$ 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^τ-differential

$$d_r^\tau(y) = \overline{\tau}^r x.$$

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.

$$\begin{tikzcd}
\Sigma^{0,-r} X/\tau^{k-r} \arrow[r, "\tau^{r-1}"] \arrow[d] & \Sigma^{0,-1} X/\tau^{k-1} \arrow[r, "\tau"] & X/\tau^k \arrow[d] \\
\Sigma^{0,-r} X/\tau & & X/\tau \arrow[ul, dashed, "\partial_1^k"]
\end{tikzcd}$$

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 $d_{\leqslant r-1}$-cycle. Choose a lift $\widetilde{y}$. It then follows that $\partial_r^k(\widetilde{y})$ is a valid choice for β as above. To show that this β lifts to X/τ^k, we need to show that $\partial^k_{k-r}(\beta) = 0$. Note that $\partial^k_{k-r} \circ \partial^k_r$ can be written as (omitting shifts for readability)

$$X/\tau^{k-1} \xrightarrow{\partial^\infty_k} X \longrightarrow X/\tau^{k-r} \xrightarrow{\partial^\infty_{k-r}} X \longrightarrow X/\tau^r.$$

The middle two maps are part of a cofibre sequence, so in particular their composition is zero. This means that indeed $\partial^k_{k-r}(\beta) = 0$, 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 $\pi_{n,s}\, X \to \pi_{n,s}\, X/\tau = \mathrm{E}_1^{n,s}$ has image given by the subgroup $\mathrm{Z}_\infty^{n,s}$ of permanent cycles. Postcomposing with the quotients, we in particular obtain, for every $1 \leqslant r \leqslant \infty$, a map

$$\pi_{n,s}\, X \longrightarrow \mathrm{E}_r^{n,s}$$

which is surjective in the case $r = \infty$.

(2) Let $r \geqslant 1$. By Theorem 3.67 (1), the reduction map $\pi_{n,s}\, X/\tau^r \to \pi_{n,s}\, X/\tau = \mathrm{E}_1^{n,s}$ has image given by the subgroup $\mathrm{Z}_{r-1}^{n,s}$ of $(r-1)$-cycles. Postcomposing with the quotient, we in particular obtain a surjective map

$$\pi_{n,s}\, X/\tau^r \twoheadrightarrow \mathrm{E}_r^{n,s}.$$

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 $r \geqslant 1$, and let n and s be integers. Then we have a commutative diagram*

$$\begin{tikzcd}
\pi_{n,s}\, X/\tau^r \arrow[r, "\partial^\infty_r"] \arrow[d] & \pi_{n-1,\,s+r}\, X \arrow[d] \\
\mathrm{E}_r^{n,s} \arrow[r, "d_r"] & \mathrm{E}_r^{n-1,\,s+r}.
\end{tikzcd}$$

More generally, for $R \geqslant r$, we have a commutative diagram

$$\begin{tikzcd}
\pi_{n,s}\, X/\tau^r \arrow[r, "\partial^R_r"] \arrow[d] & \pi_{n-1,\,s+r}\, X/\tau^{R-r} \arrow[d] \\
\mathrm{E}_r^{n,s} \arrow[r, "d_r"] & \mathrm{Z}_m^{n-1,\,s+r}/\mathrm{B}_{r-1}^{n-1,\,s+r}
\end{tikzcd}$$

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 $i \leqslant k$ 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 $X/\tau^k \to X/\tau^m$. 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 $r \geqslant 1$.*

(1) *Let $k \geqslant 1$ be fixed. Suppose that for every $1 \leqslant i \leqslant k$, 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

$$\mathrm{Z}_{i-1}^{n,\, s+k-i} / \mathrm{B}_{k-i}^{n,\, s+k-i}. \tag{3.71}$$

(*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 $1 \leqslant m \leqslant k$ be fixed. Suppose that for every $1 \leqslant i \leqslant k-m$, 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 $k \geqslant m \geqslant 1$.*

(1) *If $\pi_{n,s+d}\, X/\tau$ vanishes for $0 \leqslant d \leqslant k-1$, then $\pi_{n,s}\, X/\tau^k$ vanishes also.*

(2) *If $\pi_{n,s+d}\, X/\tau$ vanishes for $m \leqslant d \leqslant k-1$, then the reduction map $\pi_{n,s}\, X/\tau^k \to \pi_{n,s}\, X/\tau^m$ is injective.*

Proof. In the notation of Theorem 3.70, we have $\mathrm{E}_1^{n,s} = \pi_{n,s}\, X/\tau$, and $\mathrm{Z}_r^{n,s}$ 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 this thesis 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, which will also appear throughout Part II.

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 $\mathrm{Pr}^{\mathrm{L}}_{\mathrm{st}}$.

A module over FilSp in $\mathrm{Pr}^{\mathrm{L}}_{\mathrm{st}}$ 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 $\tau\colon \Sigma^{0,-1}X \to X$ the map given by tensoring the map $\tau\colon \mathbf{S}^{0,-1} \to \mathbf{S}$ 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

$$\mathcal{C}[\tau^{-1}] := \mathrm{Mod}_{\mathbf{S}[\tau^{-1}]}(\mathcal{C}),$$

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

$$\mathrm{Mod}_{C\tau}(\mathcal{C}),$$

and for $X \in \mathcal{C}$, we write

$$X/\tau := C\tau \otimes X.$$

Finally, we write $\mathcal{C}^{\wedge}_{\tau}$ for the full subcategory of $\mathcal{C}$ on the $C\tau$-local objects, which we call **τ-complete**. This results in a localisation

$$\begin{tikzcd} \mathcal{C} \arrow[r, shift left, "(-)^{\wedge}_{\tau}"] & \mathcal{C}^{\wedge}_{\tau} \arrow[l, shift left, hook'] \end{tikzcd}.$$

Remark 3.75. Using that the tensoring over FilSp preserves colimits in each variable, we find that

$$X[\tau^{-1}] = \mathrm{colim}\big(\ X \xrightarrow{\ \tau\ } \Sigma^{0,1}X \xrightarrow{\ \tau\ } \cdots\ \big)$$

and

$$X/\tau = \mathrm{cofib}(\tau\colon\ \Sigma^{0,-1}\, X \longrightarrow X).$$

Moreover, since $\mathbf{S}[\tau^{-1}]$ is an idempotent in FilSp, it follows that $\mathcal{C} \to \mathcal{C}[\tau^{-1}]$ is a smashing localisation, and the forgetful functor $\mathcal{C}[\tau^{-1}] \to \mathcal{C}$ is fully faithful with essential image consisting of those X on which τ is an isomorphism. In most examples, the forgetful functor $\mathrm{Mod}_{C\tau}(\mathcal{C}) \to \mathcal{C}$ 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,

$$\mathrm{Mod}_{R^{\tau=1}}(\mathrm{Sp}) \qquad \text{and} \qquad \mathrm{Mod}_{R/\tau}(\mathrm{grSp}).$$

If R is a filtered $\mathbf{E}_\infty$-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. ▲

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*

$$\begin{tikzcd} X \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & X^\wedge_\tau \arrow[d] \\ X[\tau^{-1}] \arrow[r] & (X^\wedge_\tau)[\tau^{-1}]. \end{tikzcd}$$

In particular, a map $X \to Y$ in $\mathcal{C}$ is an isomorphism if and only if the maps

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

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. ∎

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

$$\begin{tikzcd} \mathcal{C}[\tau^{-1}] \arrow[r, hook] & \mathcal{C} \arrow[l, bend right, "\tau^{-1}"'] \arrow[l, bend left] \arrow[r, "(-)^\wedge_\tau"] & \mathcal{C}^\wedge_\tau \arrow[l, bend right, hook'] \arrow[l, bend left, hook] \end{tikzcd}$$

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 $i\colon \mathbf{Z} \to \mathrm{FilSp}$ from Definition 2.17. Define the **filtered mapping spectrum** $\mathrm{filmap}_{\mathcal{C}}(X, Y)$ from X to Y as the filtered spectrum

$$\mathrm{map}_{\mathcal{C}}(i(-) \otimes X,\, Y)\colon \mathbf{Z}^{\mathrm{op}} \longrightarrow \mathrm{Sp},$$

where $\mathrm{map}_{\mathcal{C}}(-,-)$ denotes the mapping spectrum of $\mathcal{C}$. This is naturally functorial in X and Y, leading to a functor $\mathrm{filmap}\colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathrm{FilSp}$. Concretely, the $\mathrm{filmap}_{\mathcal{C}}(X, Y)$is given by

$$\cdots \longrightarrow \mathrm{map}_{\mathcal{C}}(\Sigma^{0,1} X,\, Y) \longrightarrow \mathrm{map}_{\mathcal{C}}(X, Y) \longrightarrow \mathrm{map}_{\mathcal{C}}(\Sigma^{0,-1} X,\, Y) \longrightarrow \cdots$$

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 $i\colon \mathbf{Z} \to \mathrm{FilSp}$ 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*

$$\mathrm{Fun}(\mathbf{Z}, \mathcal{C}) \simeq \mathrm{LFun}(\mathrm{FilSp},\, \mathcal{C})$$

such that, if $f\colon \mathbf{Z} \to \mathcal{C}$ corresponds to $F\colon \mathrm{FilSp} \to \mathcal{C}$, 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*
$$\mathrm{Fun}^{\otimes}(\mathbf{Z}, \mathcal{C}) \simeq \mathrm{LFun}^{\otimes}(\mathrm{FilSp}, \mathcal{C}),$$
where $\mathrm{Fun}^{\otimes}$ *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 $\mathrm{FilSp} \simeq \mathrm{Sp}(\mathrm{PSh}(\mathbf{Z}))$, where $\mathrm{PSh}(\mathbf{Z})$ 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 $f\colon \mathbf{Z} \to \mathcal{C}$ be a monoidal functor.

- We typically reserve the letter ρ for for the monoidal functor $\mathrm{FilSp} \to \mathcal{C}$ induced by f. This functor in particular turns $\mathcal{C}$ into a (left) FilSp-module, which informally is given by (where $A \in \mathrm{FilSp}$ 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 $\mathcal{C} \to \mathrm{FilSp}$ 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 $f(-1 \to 0)$. Thus, we can think of f as a 'prescription' for what the map τ in $\mathcal{C}$ ought to be.

[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 $\mathrm{Pr}^{\mathrm{L}}_{\mathrm{st}}$. A monoidal functor $f\colon \mathbf{Z} \to \mathcal{C}$ gives rise to an $\mathbf{E}_1$-monoidal functor $\mathrm{FilSp} \to \mathcal{C}$, but this does not equip $\mathcal{C}$ with the structure of an $\mathbf{E}_1$-algebra in $\mathrm{Pr}^{\mathrm{L}}_{\mathrm{st}}$. (This is the usual difference between an $\mathbf{E}_1$-algebra in $\mathrm{Mod}_A(\mathcal{C})$ and an $\mathbf{E}_1$-map $A \to R$ in $\mathcal{C}$.) However, in this thesis we are mostly concerned with the symmetric monoidal (i.e., $\mathbf{E}_\infty$) case where this distinction goes away, so this abuse of terminology does not have many ramifications for the rest of this thesis.

Example 3.84. Let R be a filtered $\mathbf{E}_2$-ring. Then $\mathrm{Mod}_R(\mathrm{FilSp})$ is naturally an $\mathbf{E}_1$-monoidal ∞-category, and the functor $R \otimes -\colon \mathrm{FilSp} \to \mathrm{Mod}_R(\mathrm{FilSp})$ is a monoidal functor, turning $\mathrm{Mod}_R(\mathrm{FilSp})$ into a monoidal deformation. If R is a filtered $\mathbf{E}_\infty$-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. ▲

Using that ρ is monoidal, we can import structure from FilSp into $\mathcal{C}$. For example, the $\mathbf{E}_\infty$-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 $\mathbf{E}_\infty$-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

$$\cdots \longrightarrow \mathrm{map}_{\mathcal{C}}(f(1), X) \longrightarrow \mathrm{map}_{\mathcal{C}}(f(0), X) \longrightarrow \mathrm{map}(f(-1), X) \longrightarrow \cdots,$$

where $\mathrm{map}_{\mathcal{C}}(-,-)$ 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 $\mathrm{filmap}(\mathbf{1}_{\mathcal{C}}, -)$ 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 $A \in \mathrm{FilSp}$ 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 $A \in \mathrm{FilSp}$ 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 $f\colon \mathbf{Z} \to \mathcal{C}$.*

(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*

$$\begin{tikzcd} \mathcal{C} \arrow[r, "\sigma"] \arrow[d, "(-)/\tau"'] & \mathrm{FilSp} \arrow[d, "(-)/\tau"] \\ \mathrm{Mod}_{C\tau}(\mathcal{C}) \arrow[r, "\sigma"] & \mathrm{Mod}_{C\tau}(\mathrm{FilSp}). \end{tikzcd} \quad \text{and} \quad \begin{tikzcd} \mathcal{C} \arrow[r, "\sigma"] \arrow[d, "(-)^\wedge_\tau"'] & \mathrm{FilSp} \arrow[d, "(-)^\wedge_\tau"] \\ \mathcal{C}^\wedge_\tau \arrow[r, "\sigma"] & \mathrm{FilSp}^\wedge_\tau. \end{tikzcd}$$

In particular, for $X \in \mathcal{C}$, we have natural isomorphisms

$$\sigma(C\tau \otimes X) \cong C\tau \otimes \sigma(X) \qquad \text{and} \qquad \sigma(X^\wedge_\tau) \cong \sigma(X)^\wedge_\tau.$$

(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

$$\begin{tikzcd} \mathcal{C} \arrow[r, "\sigma"] \arrow[d, "\tau^{-1}"'] & \mathrm{FilSp} \arrow[d, "\tau^{-1}"] \\ \mathcal{C}[\tau^{-1}] \arrow[r, "\sigma"] & \mathrm{FilSp}[\tau^{-1}]. \end{tikzcd}$$

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

$$\operatorname{map}_{\mathrm{FilSp}}(\mathbf{S}^{0,s}, \sigma X) \cong \operatorname{map}_{\mathcal{C}}(\rho(\mathbf{S}^{0,s}), X) \cong \operatorname{map}_{\mathcal{C}}(f(s), X)$$

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 $\pi_{n,s} \colon \mathcal{C} \to \mathrm{Ab}$ 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

$$\mathrm{E}_1^{n,s} = \pi_{n,s}(X/\tau) \implies [\Sigma^n \mathbf{1}_{\mathcal{C}}, X[\tau^{-1}]] \cong \pi_{n,*}(X)[\tau^{-1}].$$

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 $\mathrm{Mod}_{C\tau}(\mathcal{C})$ is simpler than the original case of filtered spectra, where $\mathrm{Mod}_{C\tau}(\mathrm{FilSp}) \simeq \mathrm{grSp}$ 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 $\mathbf{Z}^{\mathrm{op}} \to \mathcal{C}$ given by

$$\cdots \xrightarrow{\ \tau\ } \Sigma^{0,-2} X \xrightarrow{\ \tau\ } \Sigma^{0,-1} X \xrightarrow{\ \tau\ } X = \!=\!= \cdots .$$

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

$$\mathrm{E}_1^{n,w,s} \cong \pi_{n,\,w+s}(X/\tau) \implies \pi_{n,w}\,X.$$

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 $S^0 \to S^\sigma$, which stably results in a map called the *Euler class*

$$a_\sigma \colon \mathbf{S}^{-\sigma} \longrightarrow \mathbf{S}.$$

We expect, but do not check in detail, that this assembles to a monoidal functor

$$\cdots \xrightarrow{a_\sigma} \mathbf{S}^{-\sigma} \xrightarrow{a_\sigma} \mathbf{S} \xrightarrow{a_\sigma} \mathbf{S}^{\sigma} \xrightarrow{a_\sigma} \cdots .$$

(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 $\mathrm{FilSp} \to \mathrm{Sp}_{C_2}$, 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*

$$\begin{tikzcd} X^{C_2} \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & X^{hC_2} \arrow[d] \\ X^{\Phi C_2} \arrow[r] & X^{tC_2}. \end{tikzcd}$$

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. ▲

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 $f\colon \mathbf{Z} \to \mathrm{Sp}$ given by

$$\cdots \xrightarrow{p} \mathbf{S} \xrightarrow{p} \mathbf{S} \xrightarrow{p} \mathbf{S} \xrightarrow{p} \cdots .$$

This functor induces an adjunction

$$\mathrm{FilSp} \underset{\sigma}{\overset{\rho}{\rightleftarrows}} \mathrm{Sp}.$$

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 $\mathbf{Z} \to \mathcal{D}(\mathrm{Ab})$ 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 $\mathcal{D}(\mathrm{Ab})$, so equivalently is given by a functor $\mathbf{Z} \to \mathrm{Ab}$. 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 $\mathcal{D}(\mathrm{Ab})$ 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

$$\sigma\colon \mathrm{Syn}_E \longrightarrow \mathrm{FilSp}$$

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

$$\nu\colon \mathrm{Sp} \longrightarrow \mathrm{Syn}_E$$

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 to be used later in Part II of this thesis.

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 discuss these obstruction theories in this thesis.

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 Pstrągowski, 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], Pstrągowski constructs a symmetric monoidal ∞-category Syn_E of ***E*-based synthetic spectra**, together with a unital lax symmetric monoidal functor $\nu\colon \mathrm{Sp} \to \mathrm{Syn}_E$. 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 $\mathbf{E}_\infty$-ring** is an $\mathbf{E}_\infty$-algebra object in Syn_E. If E is clear from the context, then we may also refer to such an object as a *synthetic* $\mathbf{E}_\infty$*-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 Pstrągowski'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;

[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.

see Remark 2.84. 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.

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 2.90).

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 2.86.

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 νP is a compact and dualisable object in Syn_E, with dual $\nu(P^\vee)$. 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 $\Sigma^k \nu P$, 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 comapct dualisable generators; see, e.g., (the footnote to) [NPR24, Terminology 4.8]. ■

Next, we turn to properties of the functor ν.

Proposition 4.9.

(1) *The functor* $\nu\colon \mathrm{Sp} \to \mathrm{Syn}_E$ *is fully faithful, additive, and preserves filtered colimits. In particular,* ν *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 $Z \to \Sigma X$ *is zero on* E_**-homology.*

(3) *The comparison map* $\nu X \otimes \nu Y \to \nu(X \otimes Y)$ *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* $\nu X \otimes \nu Y \to \nu(X \otimes Y)$ *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 $\nu E \otimes X \to \nu(E \otimes X)$ 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 π_*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 π_*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 $\mathrm{Wh}\colon \mathrm{Sp} \to \mathrm{FilSp}$. 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 $\Sigma^{n,s} \colon \mathrm{Syn}_E \to \mathrm{Syn}_E$ for the functor given by tensoring with $\mathbf{S}^{n,s}$ on the left.

(3) We write $\pi_{n,s} \colon \mathrm{Syn}_E \to \mathrm{Ab}$ for the functor

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

(4) The map $\tau \colon \mathbf{S}^{0,-1} \to \mathbf{S}^{0,0}$ is the colimit-comparison map

$$\tau \colon \mathbf{S}^{0,-1} = \Sigma(\nu \mathbf{S}^{-1}) \longrightarrow \nu\mathbf{S} = \mathbf{S}^{0,0}.$$

If X is a synthetic spectrum, then tensoring it with the map $\tau \colon \mathbf{S}^{0,-1} \to \mathbf{S}$ results in a map $\Sigma^{0,-1} X \to X$, which we denote by τ_X, or by τ when there is no risk of confusion.

Note that $\pi_{*,*}$ naturally lifts to a functor $\mathrm{Syn}_E \to \mathrm{Mod}_{\mathbf{Z}[\tau]}(\mathrm{bigrAb})$.

Remark 4.14. If X is a spectrum, then we also have the natural colimit-comparison map $\Sigma\nu(\Sigma^{-1}X) \to \nu X$. 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 this thesis is 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 this thesis.

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

$$(n,s) \longmapsto (n,\, n+s) \qquad \text{and} \qquad (t,w) \longmapsto (t,\, w-t).$$

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 *νE-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

$$\nu E_{n,s}(X) = 0 \qquad \text{whenever } s > 0.$$

We will write $\tau_{\geqslant n}$ and $\tau_{\leqslant n}$ 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*

$$\nu E_{n,s}(X) = 0 \qquad \text{whenever } s > 0.$$

(b) *A synthetic spectrum X is 0-truncated if and only if X is νE-local and*

$$\nu E_{n,s}(X) = 0 \qquad \text{whenever } s < 0.$$

(c) *Let X be a synthetic spectrum. The connective cover $\tau_{\geqslant 0} X \to X$ induces an isomorphism*

$$\nu E_{n,s}(\tau_{\geqslant 0} X) \xrightarrow{\cong} \nu E_{n,s}(X) \qquad \text{whenever } s \leqslant 0.$$

Likewise, the 0-truncation $X \to \tau_{\leqslant 0} X$ induces an isomorphism

$$\nu E_{n,s}(X) \xrightarrow{\cong} \nu E_{n,s}(\tau_{\leqslant 0} X) \qquad \text{whenever } s \geqslant 0.$$

(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

$$\begin{tikzcd} \mathrm{Sp} \arrow[r, "\tau_{\leqslant 0}\nu"] \arrow[dr, "E_*(-)"'] & \mathrm{Syn}_E^{\heartsuit} \arrow[d, "\simeq"] \\ & \mathrm{grComod}_{E_*E}. \end{tikzcd}$$

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 $(\mathrm{Syn}_E)_{\geqslant 0}$ 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 $E = \mathrm{K}(n)$ and where Sp is replaced by ($\mathrm{K}(n)$-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 $\tau^{\leqslant n}$ for what we normally denote by $\tau_{\geqslant -n}$, and $\tau^{\geqslant n}$ for $\tau_{\leqslant -n}$. 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 $\pi_n^{\heartsuit}$ for the functor $\Sigma^{-n}\,\tau_{\leqslant n}\tau_{\geqslant n}$ 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 $\pi_n^{\heartsuit}$ is *not* given by bigraded homotopy groups. Instead, for $X \in \mathrm{Syn}_E$, by (the comment following) [Pst22, Theorem 4.18], we have an isomorphism of graded E_*E-comodules

$$\pi_n^{\heartsuit}(X) \cong \nu E_{*+n,-n}(X) = \nu E_{*,-n}(X)[-n].$$

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 $\pi_n^{\heartsuit}$.

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* $\mathbf{Z} \to \mathrm{Syn}_E$ *given by the multiplication-by-*τ *tower on the unit:*

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

Proof. We follow the argument given in the proof of [Law24b, Corollary 6.1]. In Pstrągowski's model [Pst22], the synthetic sphere $\mathbf{S}^{0,s}$ is defined as the sheafification of the presheaf $\tau_{\geqslant -s} \operatorname{map}(-, \mathbf{S})$, with τ induced by the suspension-comparison map. The Whitehead filtration functor is lax symmetric monoidal (Remark 2.25), as is sheafification, so the $\mathbf{E}_\infty$-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 $\mathbf{Z} \to \mathrm{Syn}_E$ from Lemma 4.28 induces an adjunction

$$\mathrm{FilSp} \underset{\sigma}{\overset{\rho}{\rightleftarrows}} \mathrm{Syn}_E$$

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 $k \geqslant 1$, the synthetic spectrum $C\tau^k$ inherits an $\mathbf{E}_\infty$-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.

[2] Alternatively, one can use Proposition 4.47 and the monoidality of the homological t-structure to give $C\tau$ an $\mathbf{E}_\infty$-structure; this is how it is done in [Pst22, Corollary 4.30].

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

$$\cdots \xrightarrow{\tau} \mathrm{map}(\mathbf{S}, \Sigma^{0,-1}X) \xrightarrow{\tau} \mathrm{map}(\mathbf{S}, X) \xrightarrow{\tau} \mathrm{map}(\mathbf{S}, \Sigma^{0,1}X) \xrightarrow{\tau} \cdots.$$

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.

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

$$\mathbf{S}^{n,s}_{\mathrm{fil}} \quad \text{and} \quad \mathbf{S}^{n,s}_{\mathrm{syn}}$$

for the filtered and synthetic spheres, respectively, and similarly $\pi^{\mathrm{fil}}_{*,*}$ and $\pi^{\mathrm{syn}}_{*,*}$ for the homotopy groups.

Proposition 4.34.

(1) *For all n and s, we have an isomorphism*

$$\rho(\mathbf{S}^{n,s}_{\mathrm{fil}}) \cong \mathbf{S}^{n,s-n}_{\mathrm{syn}}.$$

(2) *For all n, we have a natural isomorphism of graded $\mathbf{Z}[\tau]$-modules (where $X \in \mathrm{Syn}_E$)*

$$\pi^{\mathrm{syn}}_{n,*}(X) \cong \pi^{\mathrm{fil}}_{n,*+n}(\sigma X).$$

(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 $\mathbf{S}^{0,s}_{\mathrm{fil}}$ to $\mathbf{S}^{0,s}_{\mathrm{syn}}$ for all s. In particular, ρ is exact, so it preserves arbitrary suspensions. Using the identifications

$$\mathbf{S}^{n,s}_{\mathrm{fil}} \cong \Sigma^n\, \mathbf{S}^{0,s}_{\mathrm{fil}} \qquad \text{and} \qquad \Sigma_{\mathrm{syn}}(-) \cong \mathbf{S}^{1,-1}_{\mathrm{syn}} \otimes -,$$

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 $X \in \mathrm{Syn}_E$)

$$\pi^{\mathrm{syn}}_{n,s}(X) \cong \pi^{\mathrm{fil}}_{n,s+n}(\sigma X).$$

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 $(\mathrm{FilSp})_{\geqslant 0}$ is the smallest subcategory generated under colimits by the objects $\mathbf{S}^{n,s}_{\mathrm{fil}}$ for $n - s \geqslant 0$. Since synthetic analogues are connective, it follows that $\mathbf{S}^{k,u}_{\mathrm{syn}} = \Sigma^{-u}\, \nu(\mathbf{S}^{k+u})$ is $(-u)$-connective in the homological t-structure. It follows that ρ sends $\mathbf{S}^{n,s}_{\mathrm{fil}}$ for $n - s \geqslant 0$ to a connective synthetic spectrum. Because ρ preserves colimits, it follows that it restricts to a functor $(\mathrm{FilSp})_{\geqslant 0} \to (\mathrm{Syn}_E)_{\geqslant 0}$, proving the claim. ∎

Warning 4.35. Even though $\rho(\mathbf{S}^{n,s}_{\mathrm{fil}})$ is a synthetic sphere, the filtered spectrum $\sigma(\mathbf{S}^{n,s}_{\mathrm{syn}})$ is *very* different from a filtered sphere. Indeed, the spectral sequence associated to a filtered sphere is uninteresting (see Example 2.30), while the spectral sequence underlying $\sigma(\mathbf{S}^{n,s}_{\mathrm{syn}})$ 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

$$\mathrm{E}_2^{n,s} = \pi_{n,s}(C\tau \otimes X) \implies \pi_n(X[\tau^{-1}]).$$

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 substracts 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 2.37.

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.

Variant 4.39. If X is a synthetic spectrum, then its **τ-adic filtration** is the filtered synthetic spectrum $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Syn}_E$ given by

$$\cdots \xrightarrow{\tau} \Sigma^{0,-2}\,X \xrightarrow{\tau} \Sigma^{0,-1}\,X \xrightarrow{\tau} X =\!=\!= \cdots .$$

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

$$\mathrm{E}_1^{n,w,s} \cong \begin{cases} \pi_{n,\,w+s}(X/\tau) & \text{if } s \geqslant 0 \\ 0 & \text{else} \end{cases} \implies \pi_{n,w}\,X.$$

Its differential d_r^τ is of tridegree $(-1,1,r)$ for $r \geqslant 1$. 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

$$\mathrm{E}_1^{n,w,s}(X) \cong \mathrm{E}_1^{n,\,w+n,\,s}(\sigma X).$$

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 $\nu\colon \mathrm{Sp} \to \mathrm{Syn}_E$ 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 $\mathrm{Syn}_E[\tau^{-1}]$ 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 $\text{よ}\colon \mathrm{Sp} \to \mathrm{Syn}_E$ for the functor $\nu(-)[\tau^{-1}]$.

The functor よ is also referred to as the *spectral Yoneda embedding*. By definition, よ lands in τ-invertible synthetic spectra.

Theorem 4.41 ([Pst22], Theorem 4.37). *The functor よ is fully faithful, exact, and symmetric monoidal, and restricts to a symmetric monoidal equivalence*

$$\text{よ}\colon \mathrm{Sp} \xrightarrow{\simeq} \mathrm{Syn}_E[\tau^{-1}].$$

Notation 4.42. We write $(-)^{\tau=1}$ for the composite

$$\mathrm{Syn}_E \xrightarrow{\tau^{-1}} \mathrm{Syn}_E[\tau^{-1}] \simeq \mathrm{Sp}.$$

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

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

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. ▲

This identification is compatible with our earlier identification of $\mathrm{FilSp}[\tau^{-1}]$ 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*

$$\begin{tikzcd} & \mathrm{Sp} \arrow[dl, "\mathrm{Const}"'] \arrow[dr, "\text{よ}"] & \\ \mathrm{FilSp}[\tau^{-1}] \arrow[rr, shift left, "\rho"] & & \mathrm{Syn}_E[\tau^{-1}]. \arrow[ll, shift left, "\sigma"] \end{tikzcd}$$

In particular, for every spectrum X, the colimit of the filtered spectrum $\sigma(\nu X)$ is naturally isomorphic to X.

Proof. The functors $\rho\,\mathrm{Const}$ and よ are symmetric monoidal colimit-preserving functors $\mathrm{Sp} \to \mathrm{Syn}_E$. By the universal property of 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 $\sigma(\nu(X)[\tau^{-1}]) \cong \mathrm{Const}\, X$ and the fact that σ preserves τ-inversion by Proposition 4.31. ■

So far, we have focussed on ν and defined よ in terms of it. It is also possible to go in the other direction and characterise ν in terms of よ, using the homological t-structure.

Proposition 4.45. *Let X be a spectrum.*

(1) *The τ-inversion map*

$$\nu X \longrightarrow \text{よ}(X)$$

is a connective cover with respect to the homological t-structure.

(2) *There is a natural isomorphism of functors $\mathbf{Z}^{\mathrm{op}} \to \mathrm{Syn}_E$ between the Whitehead filtration of $\text{よ}(X)$,*

$$\cdots \longrightarrow \tau_{\geqslant 1}\,\text{よ}(X) \longrightarrow \tau_{\geqslant 0}\,\text{よ}(X) \longrightarrow \tau_{\geqslant -1}\,\text{よ}(X) \longrightarrow \cdots$$

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. ■

Remark 4.46. The functor $\nu\colon \mathrm{Sp} \to \mathrm{Syn}_E$ 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 $\nu \cong \tau_{\geqslant 0} \circ \text{よ}$ and by pasting adjunctions: the horizontal composites in

$$\mathrm{Sp} \underset{\text{よ}}{\overset{(-)^{\tau=1}}{\rightleftarrows}} \mathrm{Syn}_E[\tau^{-1}] \overset{\tau^{-1}}{\rightleftarrows} \mathrm{Syn}_E \underset{\tau_{\geqslant 0}}{\rightleftarrows} (\mathrm{Syn}_E)_{\geqslant 0}$$

form the adjunction

$$\mathrm{Sp} \underset{\nu}{\overset{(-)^{\tau=1}}{\rightleftarrows}} (\mathrm{Syn}_E)_{\geqslant 0}.$$

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 $\nu X \to C\tau \otimes \nu X$ 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 $\mathrm{grComod}_{E_*E} \simeq \mathrm{Syn}_E^{\heartsuit}$.

Proof. The first statement is [Pst22, Lemma 4.29]. The final claim follows by combining this with Theorem 4.23 (e). ■

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 $\mathcal{D}(\mathrm{grComod}_{E_*E})$ 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]

[3]Note that this is a phenomenon that does not occur in $\mathcal{D}(\mathrm{Mod}_A)$ or $\mathcal{D}(\mathrm{grMod}_A)$ for a (graded) commutative ring A, and is specific to the comodule setting.

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 $\mathcal{D}(\mathrm{grComod}_{(A,\Gamma)})$ 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

$$\mathrm{Stable}_{(A,\Gamma)} = \mathrm{Ind}(\mathrm{Perf}_\Gamma).$$

The inclusion functor $\mathrm{Perf}_\Gamma \to \mathcal{D}(\mathrm{grComod}_{(A,\Gamma)})$ induces a functor $\mathrm{Stable}_{(A,\Gamma)} \to \mathcal{D}(\mathrm{grComod}_{(A,\Gamma)})$, and this turns out to have a fully faithful right adjoint:

$$\mathrm{Stable}_{(A,\Gamma)} \rightleftarrows \mathcal{D}(\mathrm{grComod}_{(A,\Gamma)}). \tag{4.49}$$

As Perf_Γ is closed under tensor products, the ∞-category $\mathrm{Stable}_{(A,\Gamma)}$ is naturally a symmetric monoidal functor, and the localisation functor to the derived is symmetric monoidal. By definition, the unit of $\mathrm{Stable}_{(A,\Gamma)}$ 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*

$$\mathrm{Mod}_{C\tau}(\mathrm{Syn}_E) \hookrightarrow \mathrm{Stable}_{E_*E}$$

with the following properties.

(1) *This functor sits in a commutative diagram of lax monoidal functors*

$$\begin{tikzcd}
\mathrm{Sp} \arrow[r, "\nu(-)/\tau"] \arrow[d, "E_*(-)"'] & \mathrm{Mod}_{C\tau}(\mathrm{Syn}_E) \arrow[d, hook] \\
\mathrm{grComod}_{E_*E} \arrow[r, hook] & \mathrm{Stable}_{E_*E}.
\end{tikzcd}$$

(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]. ∎

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*

$$L_{\nu E}\,\mathrm{Mod}_{C\tau}(\mathrm{Syn}_E) \simeq \mathcal{D}(\mathrm{grComod}_{E_*E})$$

which is naturally symmetric monoidal if E is homotopy commutative.

Notation 4.54. Following [Pst22], we will also write $\widehat{\mathrm{Syn}}_E$ for $L_{\nu E}\,\mathrm{Syn}_E$. In op. cit., objects of $\widehat{\mathrm{Syn}}_E$ 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 this thesis, we will always work with a fixed E, so this confusion should not arise.

Warning 4.55. For general E, the unit in $\widehat{\mathrm{Syn}}_E$ is not compact for general E, because the unit of $\mathcal{D}(\mathrm{grComod}_{E_*E})$ is usually not compact. As a result, although $\widehat{\mathrm{Syn}}_E$ is a symmetric monoidal deformation in its own right (obtained by νE-localising the functor ρ), the resulting right adjoint $\widehat{\mathrm{Syn}}_E \to \mathrm{FilSp}$ 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:

$$[\nu Y/\tau,\, \nu X/\tau]_{C\tau} \cong \mathrm{Hom}_{E_*E}(E_*Y,\, E_*X).$$

For an integer k, let us denote the k-fold grading-shift functor on $\mathcal{D}(\mathrm{grComod}_{E_*E})$ 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

$$[\Sigma^{n,s}\,\nu Y/\tau,\, \nu X/\tau]_{C\tau} \cong [(E_*Y)[n+s],\, \Sigma^s\, E_*X]_{\mathcal{D}(\mathrm{grComod}_{E_*E})} = \mathrm{Ext}^{s,\,n+s}_{E_*E}(E_*Y,\, E_*X).$$

In particular, we have

$$\pi_{n,s}(\nu X/\tau) = [\mathbf{S}^{n,s},\, \nu X/\tau] \cong [C\tau \otimes \mathbf{S}^{n,s},\, \nu X/\tau]_{C\tau} \cong \mathrm{Ext}^{s,\,n+s}_{E_*E}(E_*,\, E_*X). \qquad \blacktriangle$$

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. $\blacktriangle$

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. We will see examples of this in practice in Part II. 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 $(S^\wedge_\tau)[\tau^{-1}]$, 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 $X\colon I \to \mathrm{Sp}$ be a diagram of spectra. Then we have an isomorphism*

$$\nu(\lim X) \cong \tau_{\geqslant 0}(\lim \nu(X)),$$

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 $\tau_{\geqslant n} S \to S$ 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

$$S[\tau^{-1}] = \operatorname{colim}(\ S \xrightarrow{\ \tau\ } \Sigma^{0,1} S \xrightarrow{\ \tau\ } \Sigma^{0,2} S \xrightarrow{\ \tau\ } \cdots\ ),$$

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 $\tau_{\geqslant n} S \to S$ is $\tau_{\leqslant n-1} S$, 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 $(\mathrm{Syn}_E)_{\geqslant 0}$. By Remark 4.46, the functor $\nu\colon \mathrm{Sp} \to (\mathrm{Syn}_E)_{\geqslant 0}$ is right adjoint to τ-inversion, implying that ν sends limits of spectra to limits in $(\mathrm{Syn}_E)_{\geqslant 0}$. A limit in the connective subcategory is computed as the connective cover of the limit in Syn_E, so we find that

$$\nu(\lim X) = \tau_{\geqslant 0}(\lim \nu(X)).$$

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 $\mathrm{Sp} \to \mathrm{FilSp}$. The better and more modern definition is the following.

Definition 4.61. The ***E*-based Adams filtration** is the functor $\sigma \circ \nu_E \colon \mathrm{Sp} \to \mathrm{FilSp}$.

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 2.81. (For an explanation why the décalage appears, see Variant 2.82 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 $\mathbf{E}_\infty$-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 $Y \to X$ 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 $s \leqslant 0$ 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

$$\pi_{*,\leqslant 0}(\nu X) \cong \pi_*(X)[\tau],$$

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

$$\mathrm{Ext}^{s-2,\,n+s-1}_{E_*E}(E_*,\, E_*X) \longrightarrow \pi_{n,s}(\nu X) \xrightarrow{\tau} \pi_{n,\,s-1}(\nu X) \longrightarrow \mathrm{Ext}^{s-1,\,n+s-1}_{E_*E}(E_*, E_*X).$$

Note that $\mathrm{Ext}^{s,t}_{E_*E}(E_*,\, E_*X) = 0$ whenever $s < 0$. Therefore if $s \leqslant 0$, 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 Pstrągowski's proof.

Proof. Using the previous result, we only have to show that $\pi_{n,s}(\nu M)$ vanishes when $s \geqslant 1$. We first show this for $s = 1$. Since M is a homotopy E-module, the Hurewicz homomorphism

$$E_*(-)\colon \pi_n M \longrightarrow \mathrm{Hom}_{E_*E}(E_*[n],\, E_*M)$$

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

$$\mathrm{Ext}^{-1,n}_{E_*E}(E_*,\, E_*M) \longrightarrow \pi_{n,1}(\nu M) \xrightarrow{\ \tau\ } \pi_{n,0}(\nu M) \longrightarrow \mathrm{Ext}^{0,n}_{E_*E}(E_*, E_*M).$$

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

$$\mathrm{Ext}^{s,t}_{E_*E}(E_*,\, E_*M) \cong \mathrm{Ext}^{s,t}_{E_*}(E_*,\, M_*).$$

In particular, we see that these Ext groups vanish whenever $s \geqslant 1$. By the long exact sequence, this means that multiplication by τ induces an isomorphism

$$\tau\colon\ \pi_{n,\,s+1}(\nu X) \xrightarrow{\cong} \pi_{n,s}(\nu X)$$

for all $s \geqslant 1$. 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

$$\nu E_{*,*}(C\tau \otimes X) \cong (\nu E_{*,*}(\nu X))/\tau \cong E_*(X),$$

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

$$\nu E_{*,*}(\text{よ} X) = \nu E_{*,*}(\nu X[\tau^{-1}]) \cong E_*(X)[\tau^{\pm}].$$

This gives an indication of why $\nu X \to \text{よ} X$ is a connective cover, and more generally, why the Whitehead tower of $\text{よ} X$

$$\cdots \longrightarrow \tau_{\geqslant 1}(\text{よ} X) \longrightarrow \tau_{\geqslant 0}(\text{よ} X) \longrightarrow \tau_{\geqslant -1}(\text{よ} X) \longrightarrow \cdots$$

can be identified with the multiplication-by-τ tower on νX

$$\cdots \xrightarrow{\tau} \Sigma^{0,-1}\nu X \xrightarrow{\tau} \nu X \xrightarrow{\tau} \Sigma^{0,1}\nu X \xrightarrow{\tau} \cdots. \qquad \blacktriangle$$

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*

$$\sigma(\nu M) \cong \operatorname{Wh} M$$

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 $\nu X \to \nu X[\tau^{-1}]$ results in a natural (symmetric monoidal) transformation

$$\sigma(\nu X) \longrightarrow \sigma(\nu X)[\tau^{-1}] \cong \operatorname{Const} X, \tag{4.68}$$

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

$$\sigma(\nu M) \longrightarrow \tau_{\geqslant 0}^{\text{diag}}(\operatorname{Const} M) = \operatorname{Wh} M.$$

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 $\sigma(\nu M)^s \to \tau_{\geqslant s} M$ 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 $s \to \infty$) of

$$\cdots \xrightarrow{\tau} \Sigma^{0,-s}\,\nu M \xrightarrow{\tau} \cdots \xrightarrow{\tau} \Sigma^{0,-1}\,\nu M \xrightarrow{\tau} \nu M$$

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

$$\operatorname{map}(\nu P, \lim_s \Sigma^{0,-s}\,\nu M) \cong \lim_s \operatorname{map}(\mathbf{S}^{0,s},\, \nu(P^\vee) \otimes \nu M) \cong \lim_s \operatorname{map}(\mathbf{S}^{0,s},\, \nu(P^\vee \otimes M)),$$

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 2.81.

Construction 4.70. Since E is a homotopy ring spectrum, its unit map $\mathbf{S} \to E$ gives rise to a semicosimplicial spectrum $E^{[\bullet]} \colon \Delta_{\mathrm{inj}} \to \mathrm{Sp}$ 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 $n \geqslant 1$, the spectrum $E^{\otimes n} \otimes X$ has the structure of a homotopy E-module. By Corollary 4.67, it follows that $\sigma(\nu(E^{\otimes n} \otimes X)) \cong \mathrm{Wh}(E^{\otimes n} \otimes X)$. As such, the above diagram induces a map

$$\sigma(\nu X) \longrightarrow \mathrm{Tot}(\mathrm{Wh}(E^{[\bullet]} \otimes X)) = \mathrm{Déc}^{\Delta}(E^{[\bullet]} \otimes X).$$

The target of this map is $\mathrm{Déc}(\mathrm{ASS}_E(X))$, as follows by combining Proposition 2.77 and Definition 2.81.

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

$$\sigma(\nu X) \longrightarrow \mathrm{Déc}(\mathrm{ASS}_E(X))$$

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

$$\nu X \longrightarrow \mathrm{Tot}(\nu(E^{[\bullet]} \otimes X)). \tag{4.72}$$

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 $\mathcal{D}(\mathrm{grComod}_{E_*E}) \subseteq \mathrm{Mod}_{C\tau}(\mathrm{Syn}_E)$ of the form

$$E_*X \longrightarrow \mathrm{Tot}((E_*E)^{\otimes[\bullet]} \otimes E_*X),$$

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 $\mathcal{D}(\mathrm{grComod}_{E_*E})$, because for any $M \in \mathrm{grComod}_{E_*E}$, 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

$$X \longrightarrow \mathrm{Tot}(E^{[\bullet]} \otimes X).$$

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 $\mathrm{ASS}_E(X)$ 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 (décalage of the) latter is completion. We expect that $\sigma(\nu X)$ is isomorphic, as a filtered spectrum, to the décalage of a filtered Adams resolution of X (see Remark 2.85), so that Remark 2.85 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 2.79.)

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 2.74 that $\mathrm{Déc}^{\Delta}$ 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

$$\sigma(\nu X) \longrightarrow \mathrm{Déc}^{\Delta}(E^{[\bullet]} \otimes X)$$

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 $\Delta \to \mathrm{Sp}$; see [MNN17, Construction 2.7]. If E carries an $\mathbf{E}_n$-structure for $1 \leqslant n \leqslant \infty$, 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 $1 \leqslant n \leqslant \infty$, then the functor

$$\mathrm{Sp} \longrightarrow \mathrm{Sp}^{\Delta}, \quad X \longmapsto E^{[\bullet]} \otimes X$$

  is naturally a lax $\mathbf{E}_{n-1}$-monoidal functor. Postcomposing with the lax symmetric monoidal functor $\mathrm{Déc}^{\Delta}$, we obtain a lax $\mathbf{E}_{n-1}$-monoidal functor

$$\mathrm{Sp} \longrightarrow \mathrm{FilSp}, \quad X \longmapsto \mathrm{Déc}^{\Delta}(E^{[\bullet]} \otimes X).$$

  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 structure on $\mathrm{Déc}(\mathrm{ASS}_E(-))$, even if E is merely homotopy-associative. Only when E is $\mathbf{E}_\infty$ 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–Pstrągowski [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 $\mathrm{RE}_\infty^{*,*} = 0$ (for instance, this happens if the spectral sequence converges strongly). Let $x \in \mathrm{E}_2^{n,s} = \pi_{n,s}(X/\tau)$ 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 $\pi_{n,s} X \to \pi_n X^{\tau=1}$.*

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 2.50, this follows from Theorem 2.52.) 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 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 $\pi_{n,*}\,\nu X$ coincides with the E-Adams filtration, recovering [BHS23, Corollary 9.21]. See Definition 2.96 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 $f\colon \mathbf{S}^n \to X$ be a map, and let $s \geqslant 0$. Then f has (algebraic) E-Adams filtration (see Definition 2.96) at least s if and only if there exists a factorisation*

$$\begin{tikzcd} \nu\mathbf{S}^n = \mathbf{S}^{n,0} \arrow[r, "\nu f"] \arrow[d, "\tau^s"'] & \nu X. \\ \Sigma^{0,s}\,\nu\mathbf{S}^n = \mathbf{S}^{n,s} \arrow[ur, dashed] & \end{tikzcd}$$

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

$$F^s\,\pi_n X = \operatorname{im}\left(\ [\mathbf{S}^{n,s}, \nu X] \xrightarrow{\tau=1} [\mathbf{S}^n, X]\ \right).$$

Proof. Combine Propositions 2.97 and 2.99 with Theorem 4.71. ∎

Phrased differently, the functor ν sends a map $f\colon \mathbf{S}^n \to X$ 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

[4]Note that they use second-page indexing, which can be obtained from ours using Remark 3.55. See also Variant 4.39.

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 $\widetilde{\alpha}$ 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 $k \geqslant s$ that are hit by a differential of length $\leqslant k-1$.

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; this is the notation we will use in Part II. ▲

As an another example, in Example 3.12, we studied essentially $\pi_{n,s}\, \mathbf{S}$ of the (2-local) MU-synthetic sphere for $n \leqslant 3$, 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) **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 in this thesis.

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 $\mathcal{C}^{\mathrm{cell}} \subseteq \mathcal{C}$ is closed under colimits and tensor products.*

Consequently, $\mathcal{C}^{\mathrm{cell}}$ inherits the structure of a presentably monoidal ∞-category from $\mathcal{C}$, and we have a colocalisation adjunction

$$\mathcal{C}^{\mathrm{cell}} \underset{(-)^{\mathrm{cell}}}{\overset{\hookrightarrow}{\leftrightarrows}} \mathcal{C}$$

with a monoidal left adjoint. If $\mathcal{C}$ is a symmetric monoidal deformation, then $\mathcal{C}^{\mathrm{cell}}$ 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 $\mathcal{C} \to \mathcal{C}^{\mathrm{cell}}$ 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 $X^{\mathrm{cell}} \to X$ induces an isomorphism

$$\pi_{*,*}\, X \cong \pi_{*,*}(X^{\mathrm{cell}}).$$

In this sense, one may think of the passage from $\mathcal{C}$ to $\mathcal{C}^{\mathrm{cell}}$ 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*

$$\begin{tikzcd} & \mathcal{C}^{\mathrm{cell}} \arrow[dr, hook] & \\ \mathrm{FilSp} \arrow[ur, "\rho"] \arrow[rr, "\rho"'] & & \mathcal{C} \end{tikzcd} \quad \textit{and} \quad \begin{tikzcd} \mathcal{C} \arrow[rr, "\sigma"] \arrow[dr, "(-)^{\mathrm{cell}}"'] & & \mathrm{FilSp}. \\ & \mathcal{C}^{\mathrm{cell}} \arrow[ur, "\sigma"'] & \end{tikzcd}$$

(3) *The adjunction $\rho \dashv \sigma$ restricts to an adjunction between* FilSp *and $\mathcal{C}^{\mathrm{cell}}$, which we denote by*

$$\mathrm{FilSp} \underset{\sigma^{\mathrm{cell}}}{\overset{\rho^{\mathrm{cell}}}{\rightleftarrows}} \mathcal{C}^{\mathrm{cell}}.$$

Proof. The first claim is a restatement of Theorem 3.88 (4). The functor ρ lands in $\mathcal{C}^{\mathrm{cell}}$ 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 $X^{\mathrm{cell}} \to X$ 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 $X^{\mathrm{cell}} \to X$ 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]. Pstrągowski 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 highlights 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

$$\mathrm{FilSp} \underset{\sigma}{\overset{\rho}{\rightleftarrows}} \mathcal{C}$$

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 $\mathbf{E}_\infty$-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 Pstrągowski, 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*

$$\mathrm{FilSp} \underset{\sigma^{\mathrm{cell}}}{\overset{\rho^{\mathrm{cell}}}{\rightleftarrows}} \mathrm{Syn}_E^{\mathrm{cell}}$$

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 $\mathbf{E}_\infty$-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 $\mathbf{E}_\infty$-ring, and the sphere is MU-nilpotent complete (see Example 2.106). 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 ν 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 $\sigma(\nu\mathbf{S})$ 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

$$\mathrm{Mod}_{(\nu\mathbf{S})^\wedge_\tau}(\mathrm{Syn}_E^{\mathrm{cell}}) \xrightarrow{\simeq} \mathrm{Mod}_{\mathrm{Tot}(\mathrm{Wh}(E^{[\bullet]}))}(\mathrm{FilSp}),$$

which is symmetric monoidal if E is an $\mathbf{E}_\infty$-ring. This follows by applying Theorem 5.8 to the symmetric monoidal deformation $\mathrm{Mod}_{(\nu\mathbf{S})^\wedge_\tau}(\mathrm{Syn}_E)$ 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 ν 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 $\sigma\colon \mathrm{Syn}_E \to \mathrm{FilSp}$ 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 Pstrągowski 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

$$\tau\colon\ \pi_{n,\,2s}\,X \longrightarrow \pi_{n,\,2s-1}\,X$$

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 $\mathrm{Wh}\, X$ is even if and only if $\pi_* X$ is concentrated in even degrees, and likewise for $\mathrm{Post}\, X$. ▲

Proposition 5.17.

(1) *The ∞-category* $\mathrm{FilSp}^{\mathrm{ev}}$ *is a presentable stable ∞-category.*

(2) *The subcategory* $\mathrm{FilSp}^{\mathrm{ev}} \subseteq \mathrm{FilSp}$ *is closed under limits, colimits, and tensor products.*

Consequently, $\mathrm{FilSp}^{\mathrm{ev}}$ *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 $\mathrm{Gr}\colon \mathrm{FilSp} \to \mathrm{grSp}$ is symmetric monoidal (see Remark 3.26), and that $\mathrm{grSp}^{\mathrm{ev}} \subseteq \mathrm{grSp}$ 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. ■

Remark 5.18. One may check that the left adjoint to the inclusion $\mathrm{FilSp}^{\mathrm{ev}} \subseteq \mathrm{FilSp}$ sends a filtered spectrum X to

$$\cdots \longrightarrow X^1 = X^1 \longrightarrow X^{-1} = X^{-1} \longrightarrow X^{-3} = \cdots$$

with X^{-1} in filtrations 0 and -1, and the right adjoint sends X to

$$\cdots \longrightarrow X^2 = X^2 \longrightarrow X^0 = X^0 \longrightarrow X^{-2} = \cdots$$

with X^0 in filtrations 0 and -1.

Remark 5.19. There is a symmetric monoidal equivalence $p\colon \mathrm{FilSp}^{\mathrm{ev}} \xrightarrow{\simeq} \mathrm{FilSp}$ that ‘pinches’ the transition maps together: it sends an even filtered spectrum X to

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

For instance, this sends $\mathbf{S}^{n,2s}$ to $\mathbf{S}^{n,s}$. Formally, precomposition with $2\colon \mathbf{Z}^{\mathrm{op}} \to \mathbf{Z}^{\mathrm{op}}$ results in a symmetric monoidal functor $\mathrm{FilSp} \to \mathrm{FilSp}$. Restricting the domain to the even subcategory results in the claimed symmetic monoidal equivalence p. Indeed, this restricted functor has a two-sided inverse given by d^*, where $d\colon \mathbf{Z}^{\mathrm{op}} \to \mathbf{Z}^{\mathrm{op}}$ is the map of posets that sends $2s$ and $2s-1$ to s.

Remark 5.20. We can regard $\mathrm{FilSp}^{\mathrm{ev}}$ as a symmetric monoidal deformation in its own right, with deformation parameter τ^2. Formally, we have a symmetric monoidal functor $2\colon \mathbf{Z} \to \mathbf{Z}$. Writing $i\colon \mathbf{Z} \to \mathrm{FilSp}$ for the functor from Definition 2.17, we see that $i \circ 2\colon \mathbf{Z} \to \mathrm{FilSp}$ is a symmetric monoidal functor that lands in the subcategory $\mathrm{FilSp}^{\mathrm{ev}}$. The underlying functor can be depicted by the diagram

$$\cdots \xrightarrow{\tau^2} \mathbf{S}^{0,-2} \xrightarrow{\tau^2} \mathbf{S} \xrightarrow{\tau^2} \mathbf{S}^{0,2} \xrightarrow{\tau^2} \cdots .$$

Via Notation 3.83, this gives $\mathrm{FilSp}^{\mathrm{ev}}$ the structure of a symmetric monoidal deformation. The resulting right adjoint $\mathrm{FilSp}^{\mathrm{ev}} \to \mathrm{FilSp}$ 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** $\mathcal{C}^{\mathrm{ev}}$ 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 $\mathcal{C}^{\mathrm{ev}}$.

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 $\mathcal{C}^{\mathrm{ev}}$ is a presentable stable ∞-category.*

(2) *The subcategory $\mathcal{C}^{\mathrm{ev}} \subseteq \mathcal{C}^{\mathrm{cell}}$ 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 $\mathcal{C}^{\mathrm{ev}}$.*

(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$^{\mathrm{ev}}$ *and* $\mathcal{C}^{\mathrm{ev}}$*, which we denote by*

$$\mathrm{FilSp}^{\mathrm{ev}} \underset{\sigma^{\mathrm{ev}}}{\overset{\rho^{\mathrm{ev}}}{\rightleftarrows}} \mathcal{C}^{\mathrm{ev}}.$$

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 $\sigma\colon \mathcal{C} \to \mathrm{FilSp}$ 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 $\mathcal{C}^{\mathrm{ev}}$, it lands in FilSp$^{\mathrm{ev}}$, so that the adjunction restricts as indicated. ■

In this case, evenness is preserved under limits too, which is not obvious from the definition of $\mathcal{C}^{\mathrm{ev}}$ from Definition 5.21.

Corollary 5.24. *Let $\mathcal{C}$ be a monoidal deformation satisfying the conditions of Proposition* 5.23*. Then the subcategory $\mathcal{C}^{\mathrm{ev}} \subseteq \mathcal{C}$ 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 $\mathbf{E}_\infty$-algebra in FilSp$^{\mathrm{ev}}$. As FilSp$^{\mathrm{ev}}$ is a monoidal subcategory of FilSp, it follows that σ induces an equivalence of ∞-categories

$$\mathcal{C}^{\mathrm{ev}} \xrightarrow{\simeq} \mathrm{Mod}_{\sigma(\mathbf{1}_{\mathcal{C}})}(\mathrm{FilSp}^{\mathrm{ev}}).$$

Postcomposing this with the (symmetric monoidal) equivalence $p\colon \mathrm{FilSp}^{\mathrm{ev}} \xrightarrow{\simeq} \mathrm{FilSp}$ from Remark 5.19, we obtain an equivalence

$$\mathcal{C}^{\mathrm{ev}} \xrightarrow{\simeq} \mathrm{Mod}_{p(\sigma(\mathbf{1}_{\mathcal{C}}))}(\mathrm{FilSp}).$$

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 2.86) 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 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*

$$\mathrm{grComod}^{\mathrm{ev}}_{(A,\Gamma)} \qquad \textit{and} \qquad \mathrm{grComod}^{\mathrm{odd}}_{(A,\Gamma)}$$

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*

$$\mathrm{grComod}_{(A,\Gamma)} \simeq \mathrm{grComod}^{\mathrm{ev}}_{(A,\Gamma)} \times \mathrm{grComod}^{\mathrm{odd}}_{(A,\Gamma)}.$$

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

$$\rho(\mathbf{S}^{n,s}_{\mathrm{fil}}) = \mathbf{S}^{n,s-n}_{\mathrm{syn}},$$

we should think of a synthetic sphere $\mathbf{S}^{n,s}_{\mathrm{syn}}$ 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

$$\pi^{\mathrm{syn}}_{n,s}(-) \cong \pi^{\mathrm{fil}}_{n,s+n}(\sigma(-)).$$

It follows that σX is even if and only if $\pi^{\mathrm{syn}}_{n,s}(X/\tau)$ vanishes whenever $n+s$ is odd. By Example 4.56, we have an isomorphism

$$\pi_{n,s}(\mathbf{S}_{\mathrm{syn}}/\tau) \cong \mathrm{Ext}^{s,n+s}_{E_*E}(E_*, E_*) = \mathrm{Ext}^{s}_{E_*E}(E_*[n+s], E_*).$$

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* $\mathrm{Syn}_E^{\mathrm{cell}}$ *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*

$$\mathrm{FilSp}^{\mathrm{ev}} \underset{\sigma^{\mathrm{ev}}}{\overset{\rho^{\mathrm{ev}}}{\rightleftarrows}} (\mathrm{Syn}_E^{\mathrm{cell}})^{\mathrm{ev}}.$$

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, Pstrągowski proves that the homological t-structure on Syn_E restricts to a t-structure on $\mathrm{Syn}_E^{\mathrm{ev}}$ whose heart is equivalent to $\mathrm{grComod}^{\mathrm{ev}}_{E_*E}$; 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 $s \geqslant 0$, we have an isomorphism $\pi_{0,s}(C\tau) \cong \mathbf{F}_2 \cdot h_0^s$. It follows that $\sigma(\mathbf{S}_{\mathrm{syn}})$ is not an even filtered spectrum. As a result, the subcategory of $\mathrm{Syn}_{\mathbf{F}_2}$ generated under colimits by the spheres $\mathbf{S}^{n,s}$ for $n+s$ even is not related to $\mathrm{FilSp}^{\mathrm{ev}}$ 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 $(\nu X)^{\mathrm{cell}}$ is even if and only if E_*X is concentrated in even degrees.*

Proof. This follows immediately from the isomorphism from Example 4.56:

$$\pi_{n,s}(\nu X/\tau) \cong \mathrm{Ext}^{s,\,n+s}_{E_*E}(E_*,\, E_*X) = \mathrm{Ext}^{s}_{E_*E}(E_*[n+s],\, E_*X),$$

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 $\mathrm{Syn}_E^{\mathrm{ev}}$. However, the map $\tau^2 \colon \mathbf{S}^{0,-2} \to \mathbf{S}$ 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. ▲

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 Pstrągowski's subcategory $\mathrm{Syn}_E^{\mathrm{ev}} \subseteq \mathrm{Syn}_E$ is closed under limits.*

Proof. This follows from Corollary 5.24. ■

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

$$\mathrm{Mod}_{(\nu\mathbf{S})^\wedge_\tau}(\mathrm{Syn}_E^{\mathrm{ev}}) \simeq \mathrm{Mod}_{\mathrm{Tot}(\tau_{\geqslant 2*}(E^{[\bullet]}))}(\mathrm{FilSp}).$$

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

$$\mathrm{grComod}^{\mathrm{ev}}_{\mathrm{MU}_*\mathrm{MU}} \simeq \mathrm{QCoh}(\mathfrak{M}_{\mathrm{fg}}).$$

As a result, the special fibres of $\mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}}$ and $\widehat{\mathrm{Syn}}_{\mathrm{MU}}^{\mathrm{ev}}$ are, respectively,

$$\mathrm{IndCoh}(\mathfrak{M}_{\mathrm{fg}}) \qquad \text{and} \qquad \mathcal{D}(\mathrm{QCoh}(\mathfrak{M}_{\mathrm{fg}})).$$

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. ▲

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 $\operatorname{Spec}\mathbf{C}$. 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 $\mathbf{E}_\infty$-structure. Gheorghe–Wang–Xu [GWX21] show that modules over $C\tau$ 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 arbit-*

rary prime. There exists a symmetric monoidal equivalence

$$(\mathrm{Sp}_{\mathbf{C}}^{\mathrm{cell}})^{\wedge}_{p} \xrightarrow{\simeq} (\mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}})^{\wedge}_{p}$$

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 **C**-motivic spectra and Adams–Novikov spectral sequences now match up with synthetic notions we previously introduced in this thesis. 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 $\mathrm{Sp}_{\mathbf{C}}$ are quite drastic.

- Most varieties are not cellular, so that $\mathrm{Sp}_{\mathbf{C}}^{\mathrm{cell}}$ 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 nontrivial 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 $\Gamma\colon \mathrm{Sp} \to \mathrm{FilSp}$ given by sending X to

$$\mathrm{Tot}(\tau_{\geqslant 2*}(\mathrm{MU}^{[\bullet]} \otimes X)).$$

Observe that this is a cosimplicial décalage (Definition 2.73) 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

$$\mathrm{Mod}_{\mathrm{Tot}(\tau_{\geqslant 2*}(\mathrm{MU}^{[\bullet]}))}(\mathrm{FilSp}) \simeq \mathrm{Syn}_{\mathrm{MU}}^{\mathrm{ev}}$$

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 **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 **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

$$\mathrm{smf} := \nu\,\mathrm{tmf}.$$

This uses that the MU-homology of tmf is even; see Chapter 11 for a discussion of how this is proved. (Note that, as 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 Smf and SMF to be introduced in Chapter 8; for a comparison, see Remark 8.3. ▲

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 $\mathbf{S}_{\mathrm{mot}}^{t,w}$ for the motivic bigraded (t,w)-sphere. As explained in [Pst22, Section 7.1], this amounts to

$$\mathbf{S}_{\mathrm{mot}}^{0,0} = \Sigma_+^\infty \mathbf{A}^0 \qquad \text{and} \qquad \mathbf{S}_{\mathrm{mot}}^{2,1} = \Sigma_+^\infty \mathbf{P}^1.$$

Let us for the moment use the motivic grading on synthetic spectra, writing $\mathbf{S}_{\mathrm{syn}}^{t,w}$ for the synthetic (t,w)-sphere in the sense of Remark 4.20. Then $\mathbf{S}_{\mathrm{syn}}^{t,w}$ 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

$$\mathbf{S}_{\mathrm{mot}}^{t,w} \qquad \text{and} \qquad \mathbf{S}_{\mathrm{syn}}^{t,2w}.$$

If we instead use Adams grading on synthetic spectra, then this correspondence is between

$$\mathbf{S}_{\mathrm{mot}}^{t,w} \qquad \text{and} \qquad \mathbf{S}_{\mathrm{syn}}^{t,\,2w-t}.$$

Part II

Synthetic modular forms

Chapter 6

Introduction to Part II

Part II *of this thesis consists of modified versions of the joint works* [CDvN25; CDvN24] *with Christian Carrick and Jack Davies.*

Elliptic cohomology and topological modular forms (tmf) play an essential role in modern stable homotopy theory. Aside from its connections to physics and number theory, tmf is a vital approximation to the stable homotopy groups of spheres $\pi_* \mathbf{S}$; see [Goe+05] and [WX17]. The work of Isaksen–Wang–Xu [IWX23], the state of the art in computations of stable homotopy groups of spheres, uses tmf as a necessary tool. A thorough understanding of the homotopy groups of tmf is therefore indispensable to our knowledge of the stable homotopy groups of spheres.

The computation of π_*tmf was first announced by Hopkins–Mahowald in [DFHH, Section 15], and details of this computation have appeared in many sources; see [Bau08], [BR], [Isa+24], [Kon12], [Rez07]. Shockingly, a complete proof has never appeared in the literature: all sources take either $\mathrm{MU}_*(\mathrm{tmf})$ or $\mathrm{H}_*(\mathrm{tmf}; \mathbf{F}_2)$ as input. This circularity in the literature was pointed out by Meier [Mei16]; see Section 6.1 below for a more detailed discussion. Mathew [Mat16] has shown that computing $\mathrm{MU}_*(\mathrm{tmf})$ or $\mathrm{H}_*(\mathrm{tmf}; \mathbf{F}_2)$ requires the Gap Theorem. Giving an independent proof of the Gap Theorem would therefore fix the circularity. This is what we do in this part.

Theorem A. *The homotopy groups $\pi_n \operatorname{Tmf}$ vanish for $-21 < n < 0$.*

The spectrum tmf is by definition the connective cover of Tmf, and the Gap Theorem allows one to deduce $\mathrm{MU}_*(\mathrm{tmf})$ from $\mathrm{MU}_*(\mathrm{Tmf})$. The latter of these homologies follows directly from the definition of Tmf as the global sections of a spectral stack.

As a consequence of our techniques, we also obtain an independent computation of the homotopy groups of both Tmf and tmf; see Corollary C.

The difficulty in proving the Gap Theorem lies in running a spectral sequence converging to $\pi_*\mathrm{Tmf}$ that does not depend on the Adams (ASS) or Adams–Novikov (ANSS) spectral sequences for tmf. In work of Konter [Kon12] for example, the descent spectral sequence (DSS) for Tmf is computed, but only by assuming the ANSS for tmf. Running the DSS without this input is difficult for at least two reasons: the DSS is not *a priori* an algebra over the ANSS for $\mathbf{S}$, and the DSS is much less sparse than the ANSS for tmf. The former makes it difficult to deduce differentials using low-degree information from $\pi_*\mathbf{S}$, and the latter means there are many possible differentials that need to be ruled out. On the other hand, directly computing the ANSS for Tmf would be extremely difficult; we detail the problems in Section 6.3.

Our approach is to use the technology of synthetic spectra. We define an MU-synthetic spectrum Smf called *synthetic modular forms*, which implements the DSS for Tmf as an algebra over the ANSS for $\mathbf{S}$ in the strongest possible sense. Exploiting this structure is the key idea of our computation.

Using the $\mathbf{E}_\infty$*-algebra* Smf *in* MU*-synthetic spectra in tandem with modern techniques, we carefully compute the DSS, leading to the Gap Theorem.*

Below, we detail the history of this circularity and our approach to Tmf using the descent spectral sequence. We then describe the construction of Smf, our techniques, and highlight some key steps in the computation.

6.1 History of the literature

The discovery of *topological modular forms* was announced in Hopkins's 1994 ICM address [Hop95]. At the same time, Hopkins also announced a computation of $\pi_*\mathrm{tmf}$ performed together with Mahowald. It took many years for the construction of topological modular forms to be written down by Behrens in [DFHH, Section 12], and in the meantime, there were many computations of $\pi_*\mathrm{tmf}$.

- In [DFHH, Section 15], originally written in 1998, Hopkins–Mahowald outline how $\pi_*\mathrm{tmf}$ can be computed by either an Adams spectral sequence (ASS) or an Adams–Novikov spectral sequence (ANSS). The E_2-pages of these spectral sequences are assumed here.
- In [Rez07], Rezk assumes the MU-homology of tmf and outlines how one can compute the homotopy groups of tmf.
- In [Bau08], Bauer assumes the MU-homology of tmf, and gives (almost) all of the subsequent details in the ANSS for tmf.
- In [Kon12], Konter uses Bauer's work to compute the DSS for Tmf.
- Much later, in [BR], Bruner–Rognes produce the most thorough computation of π_* tmf using the ASS for tmf. They assume the $\mathbf{F}_p$-homology of tmf as their starting point.

- Recently, in [Isa+24], Isaksen–Kong–Li–Ruan–Zhu compute the ANSS for tmf in the context of **C**-motivic homotopy theory, using the E_2-pages of the ASS and ANSS for tmf as input.

In all of these examples, there is a reliance on either the MU- or $\mathbf{F}_p$-homology of tmf, or on the E_2-page of the ASS or ANSS for tmf, and no reason is given why one might know these *a priori*. Indeed, the definition of tmf as the connective cover of Tmf does not make it clear how to compute its homology from first principles.

Mathew [Mat16] suggests a path to close this hole in the literature. He shows that if Tmf satisfies the Gap Theorem (Theorem A), then one can deduce the MU-homology of tmf from $\mathrm{MU}_*\mathrm{Tmf}$, which itself follows from the algebro-geometric definition of Tmf. The $\mathbf{F}_p$-homology of tmf follows from its MU-homology by a careful analysis of the map $\mathrm{MU} \to \mathbf{F}_p$. (Bear in mind however that the $\mathbf{F}_p$-homology of Tmf vanishes for all p.)

In [Mat16], Mathew implies that a proof of the Gap Theorem can be found in [Kon12], but as alluded to above, obtaining the Gap Theorem in this way would be circular, as this work relies on the computation by Bauer [Bau08]. The first time this circularity is explicitly brought up in the literature is in Meier's review of [Mat16]; see [Mei16].

This left the literature in a precarious situation. Many groundbreaking papers in topology, such as [WX17; IWX23], rely on the homotopy groups of tmf as well as its ASS and ANSS. To fix this circularity, we prove the Gap Theorem using the *descent spectral sequence* for Tmf; see Theorem B. In addition, this yields the ANSS for tmf; see Corollary D.

6.2 The descent spectral sequence: main results

The $\mathbf{E}_\infty$-ring Tmf of (*projective*) *topological modular forms* is defined as the global sections of the Goerss–Hopkins–Miller sheaf $\mathcal{O}^{\mathrm{top}}$ on the moduli stack of generalised elliptic curves $\overline{\mathfrak{M}}_{\mathrm{ell}}$; see [DFHH, Section 12] for a construction and [Dav23] for an axiomatic characterisation. This sheaf has the property that $\pi_{2d}\,\mathcal{O}^{\mathrm{top}}$ is isomorphic to the d-fold tensor power of the canonical line bundle ω. Immediately from this property, one obtains a *descent spectral sequence* (DSS) converging to the homotopy groups of Tmf:

$$E_2^{n,s} \cong H^s(\overline{\mathfrak{M}}_{\mathrm{ell}}, \omega^{\otimes(n+s)/2}) \implies \pi_n \mathrm{Tmf}.$$

The E_2-page is computed purely algebraically; see [Del75] for the computation where $s = 0$ and $n \geqslant 0$, and [Kon12, Figures 10 and 25] for the general pictures.

The main computation of this thesis is to determine all of the differentials in this spectral sequence. The proofs of all of the following results may be found in Chapter 10.

Theorem B. *The descent spectral sequence for* Tmf *takes the form depicted in Figures* A.2 *to* A.6 *at the prime* 2*, depicted in Figure* A.1 *at the prime* 3*, and collapses otherwise as detailed in Theorem* 9.66*.*

The Gap Theorem (Theorem A) immediately follows. In fact, we do not know how to prove the Gap Theorem without computing essentially the whole DSS.

Computing this spectral sequence leads us immediately to the homotopy groups of Tmf, and hence also to the homotopy groups of its connective cover tmf, which we call *connective topological modular forms*.

Corollary C. *The homotopy groups of* Tmf*, and hence also those of* $\mathrm{tmf} = \tau_{\geqslant 0}\mathrm{Tmf}$*, are computed; see Theorem* 9.66 *away from* 6*, Figure* A.1 *at the prime* 3*, and Figures* A.3 *to* A.6 *at the prime* 2*.*

Not only do we obtain the homotopy groups of tmf, but also its ANSS.

Corollary D. *There is an inclusion of the ANSS for* tmf *into the DSS for* Tmf *as a retract of spectral sequences. In particular, the ANSS for* tmf *is the region under the blue line of Figures* A.3 *to* A.6 *at the prime* 2 *from the* E_5*-page, the region under the blue line of Figure* A.1 *and* 3*, and the connective part of Theorem* 9.66 *away from* 6*.*

The ANSS and homotopy groups of periodic topological modular forms, denoted TMF, also follow from the DSS for Tmf.

Corollary E. *The ANSS for* TMF *is obtained from the DSS for* Tmf *by inverting* Δ^{24}*. Specifically, at the prime* 2 *it is obtained by inverting* Δ^8 *in Figures* A.2 *to* A.6*, at the prime* 3 *by inverting* Δ^3 *in Figure* A.1*, and away from* 6 *by inverting* Δ *in Theorem* 9.66*.*

The following flowchart indicates the relationship between these results and the literature on topological modular forms. For a more detailed roadmap to the homotopy groups of tmf, we refer to Chapter 11.

$$
\begin{tikzcd}
\text{DSS for Tmf (Th. B)} \arrow[r] \arrow[dd, "\Delta^{-24}"'] \arrow[dddr, bend right, "\text{Cor. D}"] & \text{Gap Theorem (Th. A)} \arrow[d, "\text{[Mat16, Cor. 5.2]}"'] \arrow[dr, bend left, "\text{[Mat16, Th. 5.13]}"] & \\
 & \text{MU-homology of tmf} \arrow[r, "\text{[Rez07, §21]}"] \arrow[d, shift right, "\text{[Bau08, §5\&7]}"'] \arrow[d, shift left, "\text{[Rez07]}"] & \mathbf{F}_p\text{-homology of tmf} \arrow[d, "\text{[BR, Part I]}"] \\
\text{DSS for TMF} \arrow[d, dash, "\cong"', "\text{Th. G}"] & \text{ANSS } E_2\text{-page for tmf} \arrow[d, shift right, "\text{[Bau08, §6\&8]}"'] \arrow[d, shift left, "\text{[Rez07]}"] & \text{ASS } E_2\text{-page for tmf} \arrow[d, "\text{[BR, Part II]}"] \arrow[dl, "\text{[Isa+24]}"] \\
\text{ANSS for TMF} & \text{ANSS for tmf} \arrow[l, "\Delta^{-24}"] & \text{ASS for tmf}
\end{tikzcd}
$$

Our intention is to give a geodesic route to the homotopy groups and ANSS of tmf. As a result, we in this thesis are not concerned with the DSS for Tmf in stems

-21 and below, as these have no bearing on the Gap Theorem. The same methods nevertheless allow one to compute these; we refer to [CDvN24, Section 6.9] for those arguments. Moreover, we do not compute all of the structure of tmf. For instance, we do not compute all hidden extensions in the ANSS for tmf, nor do we give a comprehensive record of Toda brackets in tmf. We only determine the information needed to deduce Theorems A and B; in particular, the tables of Section A.1 only contain the information we need to establish these results. Generalisations of the techniques to deduce hidden extensions are discussed in Remark 10.4.

A more thorough record of the multiplicative structures on tmf can be found in [BR] and [Isa+24]. We find the techniques of [Isa+24], comparing the ANSS and the ASS of tmf, to be the more methodical and efficient path, once the Gap Theorem has been proven. In addition, we should mention that the multiplicative structure of $\pi_*\mathrm{tmf}$ at odd primes is discussed in [BR, Section 13], and that Marek [Mar24] has computed the bigraded homotopy groups of the $\mathbf{F}_2$-synthetic analogue of tmf based on [BR].

Remark 6.1. It may be possible to deduce the Gap Theorem from the homotopy groups of $L_{K(2)}\mathrm{tmf}$ using an open-closed decomposition of $\overline{\mathfrak{M}}_{\mathrm{ell}}$ via the j-invariant. The homotopy groups of $L_{K(2)}\mathrm{tmf}$ are essentially computed by [Dua+24] using techniques from equivariant homotopy such as restrictions, transfers, norms, and vanishing lines coming from Real bordism theory. However, they do not prove the Gap Theorem, and we do not pursue this approach in this thesis. Our approach proceeds more from first principles using the DSS, and our methods apply much more broadly in contexts outside the homotopy fixed-point spectral sequence.

6.3 The DSS versus the ANSS

We do not use the ANSS for Tmf in any of our computations, and this is a very deliberate choice. While its E_2-page is readily computed using the algebro-geometric nature of Tmf (see [Mat16, Proposition 5.1]), the resulting spectral sequence is incredibly hard to compute. Part of the E_2-page is the same as that of the DSS, which we call the *connective region*; this is the area under the line $5s \leqslant n + 12$. Outside this region, the ANSS differs in a seemingly innocuous way: compared to the DSS, elements there are shifted vertically down by one filtration. We refer to this region as the *nonconnective region*, even though it also lives in stems $n \geqslant 0$, but only in high filtration. This filtration shift in the nonconnective region causes many of the techniques of this article to utterly fail for the ANSS of Tmf.

- The differentials in the DSS are propagated from explicit atomic differentials using the meta-arguments of Section 9.1. These meta-arguments fail for the ANSS for Tmf, as all elements outside of the connective region are Δ-power torsion.
- The class c_4 is seen to be a d_{11}-cycle by arguing that the target of the potential

d_{11} is one that supports a d_{11}; see Proposition 9.51. Not only do we not know how to justify this second nonzero d_{11} due to the first problem above, but now c_4 could *a priori* support a d_{10} in the ANSS. The same issue holds for other classes such as $h_1\Delta$.

- One way of stating the essential difference between Smf and $\nu\,$Tmf is that the former is an *even* MU-synthetic spectrum, and the latter is not. Concretely, this means that the DSS for Tmf satisfies a checkerboard phenomenon, and the ANSS for Tmf does not, and this causes the difficulties mentioned above.

We do not know how to work around these problems without appealing to the computation of the DSS.

As mentioned above, the advantage of the ANSS over the DSS for Tmf is that the former naturally receives a map from the ANSS for **S**. Using synthetic spectra, we are able to define a similar map for the DSS (see Section 6.5.1); using this, the ANSS has no advantage over the DSS.

To motivate this synthetic construction, we need to understand why the ANSS and DSS differ. It is easiest to begin by considering not Tmf, but the periodic version TMF. This arises as the global sections of the sheaf $\mathcal{O}^{\mathrm{top}}$ over the moduli stack of elliptic curves $\mathfrak{M}_{\mathrm{ell}}$, while Tmf is the global sections over the compactification $\overline{\mathfrak{M}}_{\mathrm{ell}}$. For TMF, the descent and Adams–Novikov spectral sequences *do* agree; the reason is as follows. The map $\overline{\mathfrak{M}}_{\mathrm{ell}} \to \mathfrak{M}_{\mathrm{fg}}$ sending an elliptic curve to its completion at the identity is an affine map. In particular, when we pull it back along the cover by the Lazard ring $\operatorname{Spec} L \to \mathfrak{M}_{\mathrm{fg}}$, we obtain an affine scheme; this turns out to compute the MU-homology of TMF:

$$\begin{tikzcd} \operatorname{Spec}(\mathrm{MU}_*\mathrm{TMF}) \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & \mathfrak{M}_{\mathrm{ell}} \arrow[d] \\ \operatorname{Spec} L \arrow[r] & \mathfrak{M}_{\mathrm{fg}}. \end{tikzcd}$$

In the compactified case, the map $\overline{\mathfrak{M}}_{\mathrm{ell}} \to \mathfrak{M}_{\mathrm{fg}}$ is not affine, but merely *quasi-affine*. In the pullback

$$\begin{tikzcd} P \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & \overline{\mathfrak{M}}_{\mathrm{ell}} \arrow[d] \\ \operatorname{Spec} L \arrow[r] & \mathfrak{M}_{\mathrm{fg}}, \end{tikzcd}$$

the global sections of the scheme P are isomorphic to $\mathrm{MU}_*\mathrm{Tmf}$, but it has higher cohomology; specifically, P has cohomological dimension 1 (since it admits a cover by two open affine subschemes). This cohomological dimension 1 is the reason that the E_2-page of the DSS has a nonconnective region shifted up by 1 filtration compared to the ANSS.

A hint at a way to work with the DSS, then, is that we need to work in a setting that can distinguish between elements of cohomological degree 0 and 1 in the derived global sections of the above pullback P. This is precisely why MU-synthetic spectra are the natural context for this problem: the derived ∞-category of quasi-coherent sheaves over $\mathfrak{M}_{\mathrm{fg}}$ is built into the ∞-category $\mathrm{Syn}_{\mathrm{MU}}$ (see Example 5.37). One might suspect therefore that by imitating the definition of Tmf in MU-synthetic spectra, we should obtain an object encoding the descent spectral sequence. This is indeed the case, and is precisely the definition of synthetic modular forms.

Remark 6.2. There is a simpler, more direct definition of synthetic modular forms inspired by the above pullbacks. We prove this as a later consequence in Proposition 10.7. However, as we point out in Remark 10.8, making this definition requires some computational input that we do not know how to obtain without computing (a nontrivial part of) the DSS for Tmf. As such, this description cannot be used as an alternative definition for synthetic modular forms to do this computation with.

6.4 Synthetic descent spectral sequences

At its core, the definition of synthetic modular forms is about the difference between limits in spectra and synthetic spectra. As an easier and motivating example, we consider KU with its C_2-action via complex conjugation. Real K-theory KO is the homotopy fixed-points for this action, and it is well known that the homotopy fixed-point spectral sequence (HFPSS) for this action is isomorphic to the Adams–Novikov spectral sequence (ANSS) for KO. Such identifications are a valuable bridge connecting equivariant and chromatic homotopy theory. For example, there is an analogous identification of spectral sequences for so-called higher real K-theories EO_n which plays an essential role in the *Detection Theorem* of Hill–Hopkins–Ravenel [HHR16].

We can use synthetic spectra to prove a more structured version of this identification. Namely, a consequence of the results of Chapter 7 is that the natural limit-comparison map

$$\nu_{\mathrm{MU}}(\mathrm{KO}) \longrightarrow \nu_{\mathrm{MU}}(\mathrm{KU})^{hC_2}$$

is an isomorphism of MU-synthetic $\mathbf{E}_\infty$-rings. This is not a tautology, as ν_{MU} usually does not preserve limits.

As is often the case with topological modular forms, it behaves similarly to K-theory, provided that one can find the correct formulation of the K-theoretic case. We can consider the above isomorphism through the lens of spectral algebraic geometry. If $\mathfrak{X}$ is a Deligne–Mumford stack with a sheaf of spectra on it, this leads to a *descent spectral sequence* (DSS) computing the homotopy groups of the global sections. For example, the DSS for the quotient stack $\mathrm{Spec}(\mathrm{KU})/C_2$ is equivalent to the HFPSS for KU. In general, we will consider *even-periodic refinements*, which are stacks that are locally constructed from even-periodic affine schemes, such as Spec KU; see [MM15, Section 2.2] or Section 7.1.3 below. In particular, such a spectral stack comes with a flat map $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ from the underlying stack to the moduli stack of formal groups.

Let $\mathfrak{X}$ be a Deligne–Mumford stack with a sheaf $\mathcal{F}$ of spectra on it. Recall that the global sections of $\mathcal{F}$ are a particular limit. Inspired by the above story, to obtain the descent spectral sequence, one should instead take this limit in $\mathrm{Syn}_{\mathrm{MU}}$. We carry this out in Chapter 7, of which the following result is a summary.

Theorem F. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}})$ be an even-periodic refinement. Let $\mathcal{O}^{\mathrm{syn}}$ denote the étale sheafification of the composite functor $\nu \circ \mathcal{O}^{\mathrm{top}}$. Then $\mathcal{O}^{\mathrm{syn}}$ is naturally an étale sheaf of* MU*-synthetic $\mathbf{E}_\infty$-rings, and enjoys the following properties.*

(1) *The composite functor $\tau^{-1} \circ \mathcal{O}^{\mathrm{syn}}$ of $\mathcal{O}^{\mathrm{syn}}$ with the functor $\tau^{-1} \colon \mathrm{Syn}_{\mathrm{MU}} \to \mathrm{Sp}$ is a sheaf of $\mathbf{E}_\infty$-rings, and we even have an isomorphism of sheaves of $\mathbf{E}_\infty$-rings*

$$\tau^{-1} \circ \mathcal{O}^{\mathrm{syn}} \cong \mathcal{O}^{\mathrm{top}}.$$

(2) *There is an isomorphism of filtered $\mathbf{E}_\infty$-rings*

$$\sigma(\mathcal{O}^{\mathrm{syn}}(\mathfrak{X})) \cong \mathrm{DSS}(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}).$$

(3) *There is a natural comparison map*

$$\nu(\mathcal{O}^{\mathrm{top}}(\mathfrak{X})) \longrightarrow \mathcal{O}^{\mathrm{syn}}(\mathfrak{X})$$

which is an isomorphism if and only if the sheaf cohomology groups

$$\mathrm{H}^s(\mathfrak{X} \times_{\mathfrak{M}_{\mathrm{fg}}} \operatorname{Spec} L,\, \omega^{\otimes t})$$

vanish for all $t \in \mathbf{Z}$ and all $s > 0$. This condition is satisfied, for example, if $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ is an affine map.

Proof. The three properties follow, respectively, from Proposition 7.36, Corollary 7.33, and Theorem 7.52. ∎

Property (1) says that $\mathcal{O}^{\mathrm{syn}}$ is a levelwise synthetic lift of $\mathcal{O}^{\mathrm{top}}$, so that its signature can be thought of as a modified ANSS for $\mathcal{O}^{\mathrm{top}}$, while property (2) identifies this signature with the DSS. This puts the DSS on good footing in the setting of Adams–Novikov spectral sequences, allowing them to be more easily compared to Adams–Novikov spectral sequences.

Meanwhile, property (3) generalises the discussion from Section 6.3 on why the DSS and ANSS for Tmf differ: it is due to higher sheaf cohomology of the pullback to Spec L. We can in fact quantify this difference: these sheaf cohomology groups determine the connectivity of $\mathcal{O}^{\mathrm{syn}}(\mathfrak{X})$ in the homological t-structure of synthetic spectra. Property (3) is proved by showing that the map exhibits the source as the connective cover of the target. In this way, property (3) is a showcase of the power of the homological t-structure.

Remark 6.3. In the main text, we work in the generality of a quasi-coherent sheaf over $\mathcal{O}^{\mathrm{top}}$. We define a functor

$$\mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}) \longrightarrow \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{syn}}), \quad \mathcal{F} \longmapsto \mathcal{F}^{\mathrm{syn}},$$

and prove that $\mathcal{F}^{\mathrm{syn}}$ has properties analogous to those of $\mathcal{O}^{\mathrm{syn}}$ outlined in Theorem F (in fact, $\mathcal{O}^{\mathrm{syn}}$ is the image of $\mathcal{O}^{\mathrm{top}}$ under this functor).

The above result recovers the discussion of KO and KU, and more generally homotopy fixed-point spectral sequences; see Section 7.6. More important for us is the case of synthetic modular forms, defined as

$$\mathrm{Smf} := \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}}) \qquad \text{and} \qquad \mathrm{SMF} := \mathcal{O}^{\mathrm{syn}}(\mathfrak{M}_{\mathrm{ell}}).$$

By the above, the signatures of these are given by the DSS for Tmf and TMF, respectively. The sketch from the previous section formalises to the following result.

Theorem G. *The natural comparison map*

$$\nu\,\mathrm{TMF} \longrightarrow \mathrm{SMF}$$

is an isomorphism of synthetic $\mathbf{E}_\infty$*-rings. The natural comparison map*

$$\nu\,\mathrm{Tmf} \longrightarrow \mathrm{Smf}$$

is not an isomorphism of synthetic spectra: the target is (-1)*-connective, but not connective, and the map exhibits the source as the* 0*-connective cover of the target.*

Proof. This is proved in Propositions 8.6 and 8.8, respectively. ∎

Going forward, almost all of our attention is focused on Smf, as this is most approachable to computation and most relevant for our goals.

Remark 6.4. For the connective version tmf, it is not clear how to give a purely synthetic definition that would help us carry out this computation. For instance, when phrased in terms of MU-synthetic spectra, the definition of the $\mathbf{E}_\infty$-algebra mmf from [GIKR21] boils down to $\nu\,\mathrm{tmf}$, so that it captures the ANSS for tmf. As such, the synthetic spectrum $\nu\,\mathrm{tmf}$ does not give a new way to compute the ANSS for tmf, or even its E_2-page, or the MU-homology of tmf. We discuss $\nu\,\mathrm{tmf}$ more in Example 5.43, and compare it to our setup for defining synthetic modular forms in Remark 8.3.

6.5 The computation and synthetic methods

Using the computation of its signature, the Omnibus Theorem tells us that the $\mathbf{Z}[\tau]$-module $\pi_{*,*}\,\mathrm{Smf}$ captures the DSS for Tmf. Now that this spectral sequence is implemented as a synthetic spectrum, all of the synthetic methods from Part I may be brought to bear on this spectral sequence. In fact, these tools prove to be crucial in multiple ways. We note that most of these tools are actually filtered tools as presented in Chapter 3, but two are decidedly synthetic: the construction of the map from the sphere, and the use of Moss's Theorem to compute synthetic Toda brackets. We now give an overview of some of the main examples of our techniques.

6.5.1 A map from the sphere

A major issue is that the DSS has a lot of elements in high filtration. As a result, even simple elements like h_1 (detecting η) could, a priori, support differentials. This problem would be solved if we could show that these elements are in the image from a different spectral sequence where they are permanent. This is where we use the sphere. The synthetic $\mathbf{E}_\infty$-ring Smf receives an $\mathbf{E}_\infty$-map from the unit:

$$\nu\mathbf{S} \longrightarrow \mathrm{Smf}.$$

On underlying spectral sequences, this is a map

$$\mathrm{ANSS}(\mathbf{S}) \longrightarrow \mathrm{DSS}(\mathrm{Tmf}),$$

which is very hard to construct by other means. However, we get more than just this map of spectral sequences: the $\mathbf{E}_\infty$-structure on the map allows one to transport Toda brackets as well. We use this structure to import the classes η, ν, ε, κ and $\bar{\kappa}$, as well as a few Toda bracket expressions for these, into Smf. This is crucial also to our use of tools such as the synthetic Leibniz rule and synthetic Toda brackets. Note that we are only using relatively low-dimensional information from the sphere, which does not require previous knowledge about Tmf to compute.

6.5.2 Truncated synthetic spectra

The elements in high filtration in the DSS cause further difficulties at every step of the way. Using synthetic (or filtered) spectra, we may avoid these problems by not working with Smf directly, but with Smf/τ^k for ever-increasing k. Let us recall the structure of these objects: we have a tower of synthetic $\mathbf{E}_\infty$-rings

$$\mathrm{Smf} \longrightarrow \cdots \longrightarrow \mathrm{Smf}/\tau^2 \longrightarrow \mathrm{Smf}/\tau.$$

The bottom object Smf/τ encodes exactly the second page of the DSS, while for $k > 1$, the object Smf/τ^k captures informations about the pages $\mathrm{E}_2, \ldots, \mathrm{E}_{k+1}$, while still being a highly structured object. In particular, Toda brackets in it make sense; these can be thought of as Toda brackets that are only "temporarily" defined, in the sense that they may involve elements that do not survive past a certain page. Similarly, we can speak of hidden extensions that are only "temporarily" defined. In addition to these features, the Omnibus Theorem tells us that the homotopy group $\pi_{n,s}\,\mathrm{Smf}/\tau^k$ captures information about the n-stem, but does not see phenomena that occur in filtrations $s+k$ and above. This last feature allows one to work with the DSS as if it did have a vanishing line.

6.5.3 Total differentials and the synthetic Leibniz rule

A major concrete advantage of working in synthetic (or filtered) spectra is that the differentials in the spectral sequence may be understood via the total differential ∂_1^∞ from Section 3.3. The linearity of this map (Proposition 3.39) or Burklund's filtered Leibniz rule (Theorem 3.40) are incredibly useful for deducing longer differentials from shorter ones, where the usual versions of the Leibniz rule do not give any information. We give some examples of this that appears in our computation.

Example 6.5 (Linearity of the total differential). In Proposition 9.18, we use information from the sphere to deduce a crucial differential $d_5(\Delta) = h_2 g$. This results in a total differential

$$\partial_4^8(\Delta) = \nu\bar{\kappa},$$

where Δ is the unique lift of $\Delta \in \pi_{24,0}\,\mathrm{Smf}/\tau$ to $\pi_{24,0}\,\mathrm{Smf}/\tau^4$; see Proposition 9.28. We import the hidden 2-extension from 2ν to η^3 in the ANSS for $\mathbf{S}$ to deduce the relation

$$4\nu = \tau^2\eta^3 \qquad \text{in } \pi_{3,1}\,\mathrm{Smf}.$$

Using that ∂_4^8 is $\pi_{*,*}\,\mathrm{Smf}$-linear, we learn that

$$\partial_4^8(4\Delta) = 4\nu\bar{\kappa} = \tau^2\eta^3\bar{\kappa},$$

which results in the "stretched" differential

$$d_7(4\Delta) = h_1^3 g.$$

For a full proof, see Proposition 9.29. ▲

The following differential is crucial, but an independent proof of it has previously not appeared in the literature.

Example 6.6 (Synthetic Leibniz rule). Using the synthetic Leibniz rule, we solve a question posed by Isaksen–Kong–Li–Ruan–Zhu [Isa+24, Problem 1.2]. There is a differential

$$d_7(\Delta^4) = h_1^3 g\Delta^3$$

in the DSS that does not follow from the ordinary Leibniz rule. This first appeared in the context of the ANSS for tmf without proof in [DFHH, Section 14]. It appeared again in Bauer's account of ANSS for tmf [Bau08]; however, it only appears in his charts and is not mentioned in the text. The difficulty of this differential was pointed out by [Isa+24], who showed that it follows from information in the ASS for tmf, which requires the Gap Theorem. We show that it follows quite easily from the total differential $\partial_4^8(\Delta)$ along with the synthetic Leibniz rule. Indeed, applying the synthetic Leibniz rule to the total differential $\partial_4^8(\Delta) = \nu\bar{\kappa}$ yields

$$\partial_4^8(\Delta^4) = 4\Delta^3 \cdot \partial_4^8(\Delta) = 4\nu\bar{\kappa}\Delta^3 = \tau^2\eta^3\bar{\kappa}\Delta^3,$$

resulting in the claimed differential. We are not aware of another way to deduce this differential that does not depend on the Gap theorem. For the full proof of this differential, see Proposition 9.29. ▲

6.5.4 Synthetic Toda brackets and Moss's Theorem

We require the extensive use of Toda brackets to deduce hidden extensions and nonzero differentials, as well as to rule out possible differentials. Toda brackets in synthetic spectra often have advantages over their non-synthetic counterparts. For example, it is often easier in this setting to apply versions of Moss's Theorem, which gives conditions for when a Toda bracket is detected by a Massey product formed on some page of a spectral sequence. Moreover, synthetic Toda brackets often have smaller indeterminacy than their non-synthetic counterparts.

In Appendix B, we give a treatment of Toda brackets formed in the Picard-graded homotopy groups of the unit in a monoidal stable ∞-category and various bracket shuffling formulas. This is far from the most general context in which one can speak of Toda brackets, but is sufficient for our applications, in particular in the bigraded homotopy of synthetic $\mathbf{E}_\infty$-rings like Smf/τ^k. Our approach generalises what is already in the literature by allowing for brackets of arbitrary length and the flexibility of working in a stable monoidal ∞-category. This approach applies quite broadly, for example to motivic and equivariant spectra.

In the appendix, we also describe a general approach to determining where synthetic Toda brackets (of arbitrary length) are detected in their associated spectral sequences. Here, we closely follow ideas of Burklund in [Bur22]. In particular, we prove a general form of Moss's theorem that applies to 3-fold synthetic Toda brackets; see Theorem B.15. We apply this extensively to 3-fold brackets, and we apply our approach to some crucial 4-fold brackets as well.

Example 6.7. One such example is the classical bracket $\langle \kappa, 2, \eta, \nu \rangle$ from the sphere. In Smf, one has $\tau^2\bar{\kappa} \in \langle \kappa, 2, \eta, \nu \rangle$; see Section 8.5.2. As the class $\bar{\kappa}$ comes from the synthetic sphere, the Nishida nilpotence theorem tells us that some power of $\bar{\kappa}$ is

τ-power torsion in the synthetic sphere, and hence also in Smf. Using the 4-fold Toda bracket containing $\bar{\kappa}$ and the shuffling formula for 4-fold Toda brackets in Smf/τ^{24}, we find

$$\tau^{22}\bar{\kappa}^6 \in \tau^{16}\bar{\kappa}^4\langle\kappa, 2, \eta, \nu\rangle\tau^4\bar{\kappa} = \langle\tau^{16}\bar{\kappa}^4, \kappa, 2, \eta\rangle\tau^4\nu\bar{\kappa} = 0,$$

using that $\tau^4\nu\bar{\kappa} = 0$, a consequence of the key d_5-differential of Proposition 9.18. The truncated Omnibus Theorem (Theorem 3.67) now yields the key d_{23}-differential

$$d_{23}(h_1\Delta^5) = g^6.$$

For the full proof, see Proposition 9.57. ▲

Remark 6.8. In [Bau08], Bauer produces this d_{23} using various shuffling formulas for 6-fold Toda brackets. The use of 6-fold brackets requires subtle indeterminacy arguments and some version of Moss's Theorem, which Bauer does not discuss. Using our treatment of Toda brackets along with the approach to Moss's Theorem, it seems plausible one could verify the 6-fold shuffling arguments that Bauer gives for this d_{23}. This would involve delicate indeterminacy checks however, and our approach is much more direct.

6.5.5 A synthetic transfer argument

At several points in our computation, we need additional input to verify that a certain class does not support a differential. In many cases, this follows for degree reasons or because the class comes from a permanent cycle in the ANSS for $\mathbf{S}$, and in other cases we can show the class is a cycle because it detects a Toda bracket in Smf/τ^k for a large enough k. In one crucial case, these methods are not sufficient, and we need to apply a transfer map coming from spectral algebraic geometry.

Example 6.9. The map of stacks

$$\overline{\mathfrak{M}}_1(3) \times \operatorname{Spec} \mathbf{Z}_{(2)} \longrightarrow \overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}_{(2)}$$

determines a *synthetic transfer map*

$$\mathrm{Smf}_1(3) \longrightarrow \mathrm{Smf}.$$

Modulo τ, this induces the algebraic transfer map on sheaf cohomology; see Theorem 8.18. The DSS for $\mathrm{Tmf}_1(3)$ has no differentials, so everything in the image of this transfer map is a permanent cycle in the DSS for Tmf. It follows from this that the class

$$2c_6 \in \pi_{12,0}\,\mathrm{Smf}/\tau$$

is a permanent cycle in the DSS. For a full proof, see Corollary 8.23. ▲

Remark 6.10. If one assumes the Gap Theorem, the corresponding claim in the ANSS for tmf follows easily for degree reasons. In the DSS however, there is a potential d_7 on $2c_6$ that we see no other way of ruling out.

6.6 Outline

This part can itself be divided into two parts: first, a general synthetic approach to descent spectral sequences, and the second the study and computation of synthetic modular forms.

Chapter 7 contains the general setup of descent spectral sequences through synthetic spectra. Aside from discussing the properties of the sheaf $\mathcal{O}^{\mathrm{syn}}$, the main result of this chapter is the use of synthetic spectra to compare the ANSS and DSS. Strictly speaking however, these comparison results are not necessary for the remainder of Part II.

From this point on, we specialise to synthetic modular forms. Their definition is the subject of Chapter 8, where we also compare them to topological modular forms, and lay the groundwork for the computation of synthetic modular forms. The core of the computation is Chapter 9, where we compute the descent spectral sequence of Tmf locally at every prime. We recommend the reader who is interested in learning the techniques of Part I to start there, and refer back to the precise statements in previous chapters as needed. Tables and figures summarising and accompanying these computations may be found in Appendix A. We also make use of Appendix B, where we collect some results about Toda brackets that are proved in [CDvN24].

In Chapter 10, we wrap up our discussion of Tmf by proving all of the main computational results stated in Section 6.2 above. Finally, we end with Chapter 11, where we explain how this thesis fits into a linear path to the theory of topological modular forms.

6.7 Conventions

We continue to follow the conventions mentioned in the introduction to the thesis and in Part I. In addition, we adhere to the following conventions.

Throughout this part, the word *synthetic spectrum* is synonymous with MU-*synthetic spectrum*.

To denote elements of synthetic homotopy groups, we follow Notation 4.80. For instance, we write $\eta \in \pi_{1,1}\,\mathbf{S}$ for the element in the MU-synthetic sphere that reduces mod τ to h_1, and which maps to the usual element $\eta \in \pi_1\,\mathbf{S}$ under τ-inversion.

We assume that all classical Deligne–Mumford stacks are quasi-compact and separated.

What Lurie in [SAG] calls a nonconnective qc separated spectral Deligne–Mumford stack whose underlying ∞-topos is 1-localic, we will simply refer to as a *spectral Deligne–Mumford stack*.

The symbol $\pm$ denotes an ambiguous sign, in the sense that $\pm x$ and $-(\pm x)$ have the same meaning. In particular, we do not write $\mp x$.

We use the term *hidden extension* in a simple way, referring to a relation $\alpha\beta = \tau^r\gamma$, where α, β and γ have nonzero projections mod τ and where $r > 0$. In particular, these are relations that hold in the homotopy groups of a synthetic spectrum (or a truncation thereof), without reference to the E_2-page.

Chapter 7

Synthetic descent spectral sequences

In this chapter, we describe a way to define MU-synthetic spectra that implement the descent spectral sequence (DSS) for a sheaf of spectra on a Deligne–Mumford stack. Although our motivating example is the case of topological modular forms (to be studied in the next chapter), this setup also recovers homotopy fixed-point spectral sequences; see Section 7.6.

This construction has two key benefits.

- It opens up the use of all synthetic techniques in the computation of descent spectral sequences.
- It allows for a very precise study of the difference between the DSS and the Adams–Novikov spectral sequence for the global sections, through the use of the homological t-structure of synthetic spectra.

The first of these will be amply demonstrated in later chapters in this thesis in the case of synthetic modular forms. The second is the main result of this chapter and is the subject of Section 7.4; see in particular Section 7.4.2.

We start with a review of derived algebraic geometry in Section 7.1, purely as a refresher and as a convenient reference for the rest of the chapter. We define the descent spectral sequence in Section 7.2. The construction and basic properties of $\mathcal{O}^{\mathrm{syn}}$ are the subject of Section 7.3. After the aforementioned comparison results of Section 7.4, we include a short study of inverting elements on $\mathcal{O}^{\mathrm{syn}}$ in Section 7.5, and end with a discussion of homotopy fixed-point spectral sequences in Section 7.6.

As with the rest of Part II, in this chapter we work in the MU-synthetic setting, so that *synthetic spectrum* is synonymous with MU-*synthetic spectrum*. The reason for working in this setting is explained in Section 7.3.

7.1 Derived algebraic geometry

We begin by describing the precise version of derived algebraic geometry we will be working in. We will mainly be concerned with three topics: ∞-categorical sheaves, derived stacks, and even-periodic refinements. We review results of Meier from [Mei21] to compare these derived stacks to Lurie's work [SAG].

Notation 7.1. By the phrase *Deligne–Mumford stack*, we will always mean a quasi-compact and separated Deligne–Mumford stack.

Remark 7.2. The assumption that a Deligne–Mumford stack $\mathfrak{X}$ be separated implies the following: if $\operatorname{Spec} A \to \mathfrak{X}$ is any map, then the pullback along itself is an affine scheme over $\operatorname{Spec} A$. As a result, for any map $\operatorname{Spec} A \to \mathfrak{X}$, its Čech nerve is consists levelwise of affine schemes. The quasi-compactness assumption implies that there exists an étale cover $\operatorname{Spec} A \to \mathfrak{X}$ out of an affine scheme, so that we can use the resulting Čech nerve to compute the value of sheaves. We will use these facts repeatedly throughout this chapter. Alternatively, one could assume that our stacks are merely quasi-separated, provided that all that all sheaves are assumed to be hypercomplete sheaves.

Notation 7.3.

- We write DM for the $(2,1)$-category of Deligne–Mumford stacks, which we will freely regard as an ∞-category.
- If $\mathfrak{X}$ is a Deligne–Mumford stack, then we write $\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$ for the **small étale site** of $\mathfrak{X}$, i.e., the 2-category of Deligne–Mumford stacks with an étale map to $\mathfrak{X}$, equipped with the étale topology. We write $\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}}$ for the full subcategory on the affine schemes; this inherits the structure of a site for the étale topology.

7.1.1 Derived sheaves

Definition 7.4. Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{C}$ be an ∞-category with limits. An **étale sheaf** $\mathcal{F}$ on $\mathfrak{X}$ with values in $\mathcal{C}$ is a functor

$$\mathcal{F}\colon (\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}})^{\mathrm{op}} \longrightarrow \mathcal{C}$$

that satisfies the following conditions.

(1) The functor $\mathcal{F}$ preserves finite products.

(2) Let

$$\begin{tikzcd} \mathfrak{Z} \arrow[rr] \arrow[dr] & & \mathfrak{Y} \arrow[dl] \\ & \mathfrak{X} & \end{tikzcd}$$

be a diagram where all maps are étale and where the top map is an étale cover, defining a cover $\mathfrak{Z} \to \mathfrak{Y}$ in $\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$. Let $\mathfrak{Z}^{[\bullet]}$ denote the Čech nerve of this map, regarded as an augmented simplicial object of $\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$. Then the induced map

$$\mathcal{F}(\mathfrak{Y}) \longrightarrow \operatorname{Tot} \mathcal{F}(\mathfrak{Z}^{[\bullet]})$$

is an isomorphism in $\mathcal{C}$.

We write $\operatorname{Shv}(\mathfrak{X};\mathcal{C})$ for the full subcategory of $\operatorname{Fun}((\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}})^{\mathrm{op}}, \mathcal{C})$ on the étale sheaves.

If $\mathcal{C}$ is a 1-category with limits, then this recovers the usual notion. In general, by [SAG, Proposition A.3.3.1], this agrees with the usual definition of a $\mathcal{C}$-valued sheaf on the site $\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$.

Remark 7.5. If $\mathcal{C}$ is a compactly generated presentably symmetric monoidal ∞-category, then $\operatorname{Shv}(\mathfrak{X};\mathcal{C})$ naturally acquires a presentably symmetric monoidal structure; see [SAG, Section 1.3.4]. Roughly speaking, this tensor product is given by the sheafification of the levelwise tensor product.

Although we are interested in the values of a sheaf on an arbitrary stack over $\mathfrak{X}$, the sheaf is uniquely determined by its values on affine schemes over $\mathfrak{X}$.

Proposition 7.6. *Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{C}$ be a presentable ∞-category. Let $i\colon \mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}} \to \mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$ denote the inclusion. Then restriction i^* and right Kan extension i_* along i induce an adjoint equivalence*

$$\operatorname{Shv}(\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}};\, \mathcal{C}) \underset{i_*}{\overset{i^*}{\rightleftarrows}} \operatorname{Shv}(\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}};\, \mathcal{C}).$$

If $\mathcal{C}$ is a compactly generated presentably symmetric monoidal ∞-category, then this is naturally a symmetric monoidal equivalence.

In particular, if $\mathcal{F}$ is a sheaf on $\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}}$, then the right Kan extension $i_*\mathcal{F}$ agrees with $\mathcal{F}$ on affine schemes over $\mathfrak{X}$.

Proof. First, we need to check that the right Kan extension functor i_* takes values in sheaves. If $\mathcal{C} = \mathscr{S}$, then this follows from [BGH20, Lemma 3.12.7]. The case of a general presentable $\mathcal{C}$ follows by taking the Lurie tensor product with $\mathcal{C}$. In fact, loc. cit. also shows that i_* is fully faithful between these sheaf-categories, so that the counit $i^*i_* \to \mathrm{id}$ is an isomorphism.

To prove that this adjunction is an adjoint equivalence, it remains to be checked that the unit is also an isomorphism. Let $\mathcal{F}$ be a sheaf on $\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$, and let $\mathfrak{Y} \to \mathfrak{X}$ be an étale map. To show that the unit map

$$\mathcal{F}(\mathfrak{Y}) \longrightarrow (i_*i^*\mathcal{F})(\mathfrak{Y})$$

is an isomorphism, we choose an étale cover $\operatorname{Spec} A \to \mathfrak{Y}$, and let $\operatorname{Spec} A^{[\bullet]}$ denote the associated Čech nerve.[1] By naturality, the unit map sits in a commutative diagram

$$\begin{tikzcd} \mathcal{F}(\mathfrak{Y}) \arrow[r] \arrow[d] & (i_* i^* \mathcal{F})(\mathfrak{Y}) \arrow[d] \\ \operatorname{Tot} \mathcal{F}(\operatorname{Spec} A^{[\bullet]}) \arrow[r] & \operatorname{Tot}((i_* i^* \mathcal{F})(\operatorname{Spec} A^{[\bullet]})). \end{tikzcd}$$

The downward maps are isomorphisms because $\mathcal{F}$ and $i_* i^* \mathcal{F}$ are sheaves. It therefore suffices to show that the map of cosimplicial objects $\mathcal{F}(\operatorname{Spec} A^{[\bullet]}) \to (i_* i^* \mathcal{F})(\operatorname{Spec} A^{[\bullet]})$ is levelwise an isomorphism. As the cover $\operatorname{Spec} A^{[\bullet]}$ consists entirely of affine schemes, this follows from the fact that the counit is an isomorphism.

Finally, the functor i^* is strong symmetric monoidal, as follows from the definition of the symmetric monoidal structure on sheaves; see Remark 7.5. This proves the claim about symmetric monoidality. ■

Notation 7.7. Let $\mathfrak{Y} \to \mathfrak{X}$ be an étale map. We write

$$\Gamma(\mathfrak{Y}, -) \colon \operatorname{Shv}(\mathfrak{X}; \mathcal{C}) \xrightarrow{\mathcal{F} \mapsto \mathcal{F}(\mathfrak{Y})} \mathcal{C}.$$

In the special case where $\mathfrak{Y} = \mathfrak{X}$ and the map is the identity map, we call this the **global sections functor**. If the stack $\mathfrak{X}$ is understood, then we may also denote the global sections functor by $\Gamma(-)$.

Concretely, one may compute sections of a sheaf $\mathcal{F}$ on $\mathfrak{Y}$ by choosing an étale cover $\operatorname{Spec} A \to \mathfrak{Y}$, and then using that

$$\Gamma(\mathfrak{Y}, \mathcal{F}) \cong \operatorname{Tot} \mathcal{F}(\operatorname{Spec} A^{[\bullet]}).$$

Definition 7.8. Let $\mathfrak{X}$ be a Deligne–Mumford stack, let $\mathcal{F}$ be an étale sheaf of spectra on $\mathfrak{X}$, and let n be an integer. The ***n*-th homotopy sheaf** of $\mathcal{F}$ is the étale sheaf $\pi_n \mathcal{F}$ of abelian groups on $\mathfrak{X}$ given by the étale sheafification of the presheaf $\pi_n \circ \mathcal{F}$.

Unless otherwise indicated, the notation $\pi_n \mathcal{F}$ will always mean the homotopy sheaf, and we reserve $\pi_n \circ \mathcal{F}$ for the functor composition. We warn the reader that it is possible that $\pi_n \mathcal{F}$ is identically zero, even though the values of $\mathcal{F}$ do not all have vanishing π_n. (This issue disappears for hypersheaves.)

[1] Here it is crucial that we assume $\mathfrak{X}$ to be separated, so that the Čech nerve consists of affine schemes; see Remark 7.2. As explained in [BGH20, Section 3.12], for sheaves of spaces on a site with a basis, right Kan extension from the basis to the entire site need not be essentially surjective on sheaves. In general, this is only true on hypersheaves.

7.1.2 Derived stacks

So far, we have only considered derived sheaves on an ordinary stack, without turning the stack itself into an object of derived algebraic geometry. For our purposes, we can turn a stack into a derived object by lifting its structure sheaf to $\mathbf{E}_\infty$-rings. We copy the following definition from Meier [Mei21, Appendix B].

Definition 7.9 ([Mei21], Definition B.1). A **derived stack** is a pair $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ where $\mathfrak{X}$ is a Deligne–Mumford stack and $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ is an étale sheaf of $\mathbf{E}_\infty$-rings on $\mathfrak{X}$, that satisfies the following conditions.

(a) There is an isomorphism $\pi_0\, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}} \cong \mathcal{O}_{\mathfrak{X}}$, where $\pi_0\, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ denotes the 0-th homotopy sheaf of $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$.

(b) For every n, the functor

$$(\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}})^{\mathrm{op}} \longrightarrow \mathrm{Ab}, \quad \operatorname{Spec} A \longmapsto \pi_n(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}(\operatorname{Spec} A))$$

is a quasi-coherent sheaf on $\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}}$.

We will refer to $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ as the **structure sheaf** of the derived stack $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$.

Meier compares this notion to Lurie's *nonconnective spectral Deligne–Mumford stacks* of [SAG, Definition 1.4.4.2], showing that it is a special type of the latter.

Proposition 7.10 ([Mei21], Lemma B.2). *Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ be an étale sheaf of $\mathbf{E}_\infty$-rings on $\mathfrak{X}$. Then $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ is a derived stack if and only if $(\mathrm{Shv}(\mathfrak{X}; \mathscr{S}), \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ is a nonconnective spectral Deligne–Mumford stack in the sense of Lurie. Moreover, every nonconnective spectral Deligne–Mumford stack whose underlying ∞-topos is 1-localic arises in this way.*

This characterisation leads to the existence of pullbacks of derived stacks under certain conditions.

Definition 7.11. Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{O}$ be an étale sheaf of $\mathbf{E}_\infty$-rings on $\mathfrak{X}$. We say that $\mathcal{O}$ is **even-periodic** if

(1) the homotopy sheaf $\pi_2\, \mathcal{O}$ is an invertible $\pi_0\, \mathcal{O}$-module;

(2) the homotopy sheaf $\pi_1\, \mathcal{O}$ vanishes;

(3) for every n, the natural map

$$\pi_2\, \mathcal{O} \otimes_{\pi_0 \mathcal{O}} \pi_n\, \mathcal{O} \longrightarrow \pi_{n+2}\, \mathcal{O}$$

is an isomorphism.

Proposition 7.12 ([Mei21], Lemma B.3). *Let*

$$(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}) \longrightarrow (\mathfrak{Z}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{Z}}) \longleftarrow (\mathfrak{Y}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{Y}})$$

be a span of derived stacks all of whose structure sheaves are even-periodic. Then the pullback of this span in nonconnective spectral Deligne–Mumford stacks is again a derived stack $(\mathfrak{P}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{P}})$*, with* $\mathcal{O}^{\mathrm{top}}_{\mathfrak{P}}$ *an even-periodic sheaf. Moreover, if one of the underlying maps of Deligne–Mumford stacks is flat, then the underlying Deligne–Mumford stack* $\mathfrak{P}$ *is equivalent to the pullback of the underlying Deligne–Mumford stacks*

$$\mathfrak{X} \longrightarrow \mathfrak{Z} \longleftarrow \mathfrak{Y}.$$

Next, we recall the notion of a quasi-coherent sheaf over a derived stack. In a few places later in this chapter, we need this definition in slightly bigger generality, where the structure sheaf is not lifted to spectra, but to another monoidal ∞-category. Accordingly, we phrase this definition in this setting. Note that by [SAG, Proposition 2.2.4.3], this matches with the definition of quasi-coherent sheaves over (nonconnective) spectral Deligne–Mumford stacks.

Definition 7.13. Let $\mathcal{C}$ be a compactly generated presentably symmetric monoidal ∞-category. Let $\mathfrak{X}$ be a classical Deligne–Mumford stack, and let $\mathcal{O}$ be an étale sheaf on $\mathfrak{X}$ with values in $\mathrm{CAlg}(\mathcal{C})$. We say that an $\mathcal{O}$-module M in $\mathrm{Shv}(\mathfrak{X};\mathcal{C})$ is **quasi-coherent** if for every map of affine schemes $U \to V$ in the small étale site of $\mathfrak{X}$, the natural map of $\mathcal{O}(U)$-modules

$$M(V) \otimes_{\mathcal{O}(V)} \mathcal{O}(U) \longrightarrow M(U)$$

is an isomorphism. We write $\mathrm{QCoh}(\mathfrak{X}, \mathcal{O}) \subseteq \mathrm{Mod}_{\mathcal{O}}(\mathrm{Shv}(\mathfrak{X};\mathcal{C}))$ for the full subcategory on the quasi-coherent $\mathcal{O}$-modules.

7.1.3 Even-periodic refinements

We will focus on a particular kind of derived stack, namely ones that are even-periodic in a precise sense.

First, recall that associated to a formal group law F over a ring R, the *associated formal group* $\widehat{\mathbf{G}}_F$ represents the functor sending an R-algebra to its set of nilpotent elements. This set becomes a group under the multiplication given by $x +_F y = F(x, y)$. Using this notation, one can write down the modern interpretation of the Landweber exact functor theorem, as can be found in [EC2, Theorem 0.0.1], for example.

Notation 7.14. Write $\mathfrak{M}_{\mathrm{fg}}$ for the moduli stack of formal groups. Write ω for the line bundle defined on $\mathfrak{M}_{\mathrm{fg}}$ that sends a formal group $p\colon G \to X$ over a scheme X to the line bundle $p_*\Omega^1_{G/X}$. Alternatively, ω sends a formal group $G \to X$ to the dual of $\ker(G(R[t]/t^2) \to G(R))$, i.e., to the dual of the Lie algebra of G.

Warning 7.15. The stack $\mathfrak{M}_{\mathrm{fg}}$ cannot be upgraded to a derived stack: it is not a Deligne–Mumford stack, or even an Artin stack. An fpqc sheaf of $\mathbf{E}_\infty$-rings lifting the structure sheaf in a similar way as in Definition 7.9 is unlikely to exist; see [Goe09]. There is a different way in which $\mathfrak{M}_{\mathrm{fg}}$ *can* be upgraded to a spectral algebro-geometric stack; see [Gre21].

Theorem 7.16 (Landweber exact functor theorem). *Let R be a classical commutative ring with a formal group law F; this determines a map of stacks* $\operatorname{Spec} R \to \mathfrak{M}_{\mathrm{fg}}$. *If this map is flat, then there exists an even-periodic homotopy-commutative ring spectrum A equipped with an isomorphism of rings $\pi_0 A \cong R$ and an isomorphism of formal groups over R*

$$\operatorname{Spf} A^0(\mathbf{CP}^\infty) \cong \widehat{\mathbf{G}}_F.$$

Moreover, $\pi_{2k}A$ can be naturally identified with the line bundle $\omega_R^{\otimes k}$ associated to $\widehat{\mathbf{G}}_F$ over $\operatorname{Spec} R$.

The derived stacks we study in this chapter locally look like they came from this theorem.

Definition 7.17. Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be a derived stack, and let $f\colon \mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ be a flat map of stacks. We say that $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ is an **even-periodic refinement** (of $\mathfrak{X}$ and f) if $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ is *locally Landweber exact*, meaning that for each étale map of stacks $\operatorname{Spec} A \to \mathfrak{X}$, the $\mathbf{E}_\infty$-ring $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}(\operatorname{Spec} A)$ represents the Landweber-exact even-periodic cohomology theory associated to the composite

$$\operatorname{Spec} A \longrightarrow \mathfrak{X} \xrightarrow{f} \mathfrak{M}_{\mathrm{fg}}.$$

Often we leave the map f implicit, and simply refer to an even-periodic refinement $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$.

Lemma 7.18. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be an even-periodic refinement. Then $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ is even-periodic in the sense of Definition* 7.11.

Proof. All conditions may be checked locally on affine schemes, where they follow immediately from Theorem 7.16. ■

There are many examples of even-periodic refinements. We will focus on two main applications.

Example 7.19. Let E be an even-periodic $\mathbf{E}_\infty$-ring with an action of a finite group G. Then the quotient $\operatorname{Spec}(E)/G$ is an even-periodic refinement of $\operatorname{Spec}(\pi_0 E)/G$. Common examples include the Morava E-theories E_n and the action by various finite subgroups of the Morava stabiliser group. An integral example of this form is $\operatorname{Spec}(\mathrm{KU})/C_2$, where KU carries the complex-conjugation action. ▲

Example 7.20. Let $\overline{\mathfrak{M}}_{\mathrm{ell}}$ be the compactification of the moduli stack of elliptic curves. Then the sheaf $\mathcal{O}^{\mathrm{top}}$ is an even-periodic refinement of $\overline{\mathfrak{M}}_{\mathrm{ell}}$ by design; see [Goe10]. More general examples of this flavour include anything coming out of *Lurie's theorem*, see [Dav24b, Section 5], such as the spectra of topological automorphic forms of Behrens–Lawson [BL10]. ▲

7.2 The descent spectral sequence

If $\mathfrak{X}$ is a stack equipped with a sheaf of spectra $\mathcal{F}$, then the *descent spectral sequence* (DSS) for $\mathcal{F}$ attempts to compute the homotopy groups of the global sections $\Gamma(X, \mathcal{F})$ by means of the homotopy sheaves $\pi_* \mathcal{F}$. In this section we give a definition of the descent spectral sequence as a filtered spectrum. Note the similarities with the construction of the Adams spectral sequence from Section 2.5.

Although we can make the following definition for a general sheaf of spectra on a stack, we will only apply it in the case where the sheaf is quasi-coherent. Recall the cosimplicial décalage functor $\mathrm{Déc}^{\Delta}$ from Section 2.5.1.

Construction 7.21. Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{F}$ be an étale sheaf of spectra on $\mathfrak{X}$. Choose an étale cover $\operatorname{Spec} A \to \mathfrak{X}$ of $\mathfrak{X}$, and let $\operatorname{Spec} A^{[\bullet]} \to \mathfrak{X}$ be the Čech nerve associated to the cover $\operatorname{Spec} A \to \mathfrak{X}$;[2] note that this is an augmented simplicial object over $\mathfrak{X}$. We define the **cosimplicial descent filtration** for the pair $(\mathfrak{X}, \mathcal{F})$ as the filtered spectrum

$$\mathrm{DSS}(\mathfrak{X}, \mathcal{F}) = \mathrm{Déc}^{\Delta}(\mathcal{F}(\operatorname{Spec} A^{[\bullet]})).$$

Its underlying spectral sequence is called the **descent spectral sequence**. We index this to start on the second page. Since $\mathrm{Déc}^{\Delta}$ is lax symmetric monoidal by Remark 2.74, it follows that if $\mathcal{F}$ carries an $\mathcal{O}$-algebra structure over an ∞-operad $\mathcal{O}$, then $\mathrm{DSS}(\mathfrak{X}, \mathcal{F})$ is naturally a filtered $\mathcal{O}$-algebra.

Remark 7.22. We apply a décalage by definition in order to give the resulting object better monoidality properties; see Remark 2.79. Unfortunately, this does not match with the notational conventions we used for the Adams spectral sequence in Definition 2.81, but it does match with the synthetic conventions (see Theorem 4.71 (2)).

A priori, the filtered spectrum $\mathrm{DSS}(\mathfrak{X}, \mathcal{F})$ depends on a choice of cover of $\mathfrak{X}$. We now show that this is not the case, and also show that it has the properties we expect of the descent spectral sequence. In the next section we will define a version that does not depend on a choice of cover, so we content ourselves with this for the moment.

Lemma 7.23. *Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{F}$ be an étale sheaf spectra on $\mathfrak{X}$ such that for every n, the homotopy sheaf $\pi_n\, \mathcal{F}$ is a quasi-coherent $\mathcal{O}_{\mathfrak{X}}$-module.*

[2]Note that such an affine étale cover and an affine étale Čech nerve exist as we assume our Deligne–Mumford stacks are qc and separated, respectively.

(1) *The filtered spectrum* $\mathrm{DSS}(\mathfrak{X},\mathcal{F})$ *is a complete filtration of* $\Gamma(\mathfrak{X},\mathcal{F})$*, i.e., the limit of* $\mathrm{DSS}(\mathfrak{X},\mathcal{F})$ *is zero and its colimit is isomorphic to* $\Gamma(X,\mathcal{F})$*. Moreover, if* $\mathcal{F}$ *carries an* $\mathcal{P}$*-algebra structure over an* ∞*-operad* $\mathcal{P}$*, then the isomorphism is one of* $\mathcal{P}$*-algebras.*

(2) *The second page of the spectral sequence underlying* $\mathrm{DSS}(\mathfrak{X},\mathcal{F})$ *takes the form*

$$\mathrm{E}_2^{n,s} \cong \mathrm{H}^s(\mathfrak{X},\, \pi_{n+s}\,\mathcal{F}).$$

Proof. Pick an étale cover $\operatorname{Spec} A \to \mathfrak{X}$. As $\mathcal{F}$ is an étale sheaf, we have an isomorphism

$$\Gamma(\mathfrak{X},\mathcal{F}) \xrightarrow{\cong} \operatorname{Tot}\mathcal{F}(\operatorname{Spec} A^{[\bullet]}).$$

The first claim therefore follows from Proposition 2.76.

A complex computing the sheaf cohomology of $\pi_{n+s}\,\mathcal{F}$ can be constructed by taking an affine étale cover of $\mathfrak{X}$ and then evaluating $\mathcal{F}$ on the Čech nerve of this cover. By definition, this is precisely the second page of the spectral sequence associated to $\mathrm{DSS}(\mathfrak{X},\mathcal{F})$, and the second result follows. ∎

Lemma 7.24. *Let* $(\mathfrak{X},\mathcal{F})$ *be as in Construction* 7.21*. Choose two affine étale covers* $\operatorname{Spec} A_1 \to \mathfrak{X}$ *and* $\operatorname{Spec} A_2 \to \mathfrak{X}$ *of* $\mathfrak{X}$*, and write* $\mathrm{DSS}_1(\mathfrak{X},\mathcal{F})$ *and* $\mathrm{DSS}_2(\mathfrak{X},\mathcal{F})$ *for the associated filtered spectra à la Construction* 7.21*. Then there is a preferred isomorphism of filtered spectra*

$$\mathrm{DSS}_1(\mathfrak{X},\mathcal{F}) \cong \mathrm{DSS}_2(\mathfrak{X},\mathcal{F}).$$

Moreover, if $\mathcal{F}$ *carries an* $\mathcal{P}$*-algebra structure over an* ∞*-operad* $\mathcal{P}$*, then this isomorphism is naturally one of filtered* $\mathcal{P}$*-algebras.*

Proof. Define a third affine étale cover $\operatorname{Spec} B$ as the common pullback of stacks:

$$\begin{tikzcd} \operatorname{Spec} B \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner"] & \operatorname{Spec} A_1 \arrow[d] \\ \operatorname{Spec} A_2 \arrow[r] & \mathfrak{X}. \end{tikzcd}$$

We then have an identification of Čech nerves

$$\operatorname{Spec} B^{[\bullet]} \cong \operatorname{Spec} A_1^{[\bullet]} \times_{\mathfrak{X}} \operatorname{Spec} A_2^{[\bullet]}.$$

Note that the pullbacks above are indeed affine, because $\mathfrak{X}$ is assumed to be separated. We obtain morphisms of cosimplicial spectra

$$\mathcal{F}(\operatorname{Spec} A_1^{[\bullet]}) \longrightarrow \mathcal{F}(\operatorname{Spec} B^{[\bullet]}) \longleftarrow \mathcal{F}(\operatorname{Spec} A_2^{[\bullet]})$$

respecting any $\mathcal{P}$-structure on $\mathcal{F}$. The desired comparison morphisms of filtered spectra come from applying the décalage functor to these cosimplicial objects:

$$\mathrm{DSS}_1(\mathfrak{X},\mathcal{F}) \longrightarrow \mathrm{Déc}^{\Delta}(\mathcal{F}(\operatorname{Spec} B^{[\bullet]})) \longleftarrow \mathrm{DSS}_2(\mathfrak{X},\mathcal{F}).$$

We show that the left morphism is an equivalence; the same argument applies to the right morphism, which would therefore conclude the proof. Because the filtered spectra involved are complete by Proposition 2.76, it suffices to prove that it is an isomorphism on associated graded; see Proposition 3.30. By Lemma 7.23, the homotopy groups of the associated graded of both the source and target are naturally isomorphic to the sheaf cohomology of $\mathcal{F}$. ■

7.3 The synthetic lift of an even-periodic refinement

With a definition of the descent spectral sequence in hand, we will now lift these filtered spectra to MU-synthetic spectra. The reason we work with MU-synthetic spectra is because, locally on affines, an even-periodic refinement is consists of complex-orientable ring spectra, so that on affines, ν is an easily-understood operation (through the results of Section 4.4).

If $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}})$ is an even-periodic refinement, then we would like to define $\mathcal{O}^{\mathrm{syn}}$ as the sheafification of the composite $\nu \circ \mathcal{O}^{\mathrm{top}}$. This is a nontrivial operation as ν does not preserve all limits. To compute this sheafification effectively, we rely on the following result.

Proposition 7.25. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be an even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module. Then the composite functor*

$$(\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}})^{\mathrm{op}} \xrightarrow{\ \mathcal{F}\ } \mathrm{Sp} \xrightarrow{\ \nu\ } \mathrm{Syn}_{\mathrm{MU}}$$

is a sheaf.

Proof. We check the two conditions of Definition 7.4. The functor $\nu \circ \mathcal{F}$ preserves finite products, as both ν and $\mathcal{F}$ do. For the second condition, we have to show that if $\operatorname{Spec} B^{[\bullet]} \to \operatorname{Spec} A$ is the Čech nerve of an étale cover $\operatorname{Spec} B \to \operatorname{Spec} A$, then the map

$$\nu\, \mathcal{F}(\operatorname{Spec} A) \longrightarrow \operatorname{Tot}\Big(\nu\, \mathcal{F}(\operatorname{Spec} B^{[\bullet]})\Big) \tag{7.26}$$

is an isomorphism. We will do this by calculating the bigraded homotopy groups of both sides.

Let us omit Spec from the notation, writing $\mathcal{F}(A)$ instead of $\mathcal{F}(\operatorname{Spec} A)$, and likewise for B. The homotopy of the left-hand side is computed by Proposition 4.65 to be $\pi_*\mathcal{F}(A)[\tau]$. Likewise, we have that $\pi_{*,*}(\nu\mathcal{F}(B^{[\bullet]})) \cong \pi_*\mathcal{F}(B^{[\bullet]})[\tau]$. To show that the map (7.26) is an isomorphism, we use the Tot spectral sequence internal to synthetic spectra. By our previous computations, we see that normalised cochain complex computing its first page (see Proposition 2.71) is obtained by tensoring the normalised cochain complex of $\pi_*\mathcal{F}(B^{[\bullet]})$ with $\mathbf{Z}[\tau]$. As $\mathbf{Z}[\tau]$ is flat over $\mathbf{Z}$, we find that the first page of this synthetic Tot spectral sequence is $\mathbf{Z}[\tau]$ tensored with the first page of the Tot spectral sequence for $\mathcal{F}(B^{[\bullet]})$. Hence the second page of the synthetic Tot spectral sequence is given by the cohomology of the quasi-coherent sheaf $\pi_*\mathcal{F}[\tau]$ (this is a quasi-coherent sheaf by [SAG, Proposition 2.2.6.1]). However, because $\operatorname{Spec} A$ is affine, the higher cohomology groups of the quasi-coherent sheaf $\pi_*\mathcal{F}$ vanish, implying that the edge map

$$\pi_*\mathcal{F}(A)[\tau] \xrightarrow{\cong} \mathrm{H}^*(\operatorname{Spec} A,\, \pi_*\mathcal{F}[\tau])$$

induced by (7.26) is an isomorphism. ■

Via Proposition 7.6, we can extend this sheaf on affines to a sheaf on all étale covers of $\mathfrak{X}$.

Definition 7.27. Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be an even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module.

(1) We write $\mathcal{F}^{\mathrm{syn}}$ for the right Kan extension of the sheaf $\nu\mathcal{F}$ on the site $\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}}$ of Proposition 7.25 to the small étale site $\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}}$ of $\mathfrak{X}$. Since ν is lax symmetric monoidal, it follows from Proposition 7.6 that the functor

$$\mathrm{QCoh}(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}) \longrightarrow \mathrm{Shv}(\mathfrak{X}; \mathrm{Syn}_{\mathrm{MU}}), \quad \mathcal{F} \longmapsto \mathcal{F}^{\mathrm{syn}}$$

is naturally lax symmetric monoidal.

(2) In the special case $\mathcal{F} = \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$, we denote $\mathcal{F}^{\mathrm{syn}}$ by $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$. As $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ is a sheaf of $\mathbf{E}_\infty$-rings, the sheaf $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$ is naturally a sheaf of synthetic $\mathbf{E}_\infty$-rings.

As a result, the functor $\mathcal{F} \mapsto \mathcal{F}^{\mathrm{syn}}$ naturally lifts to one landing in modules over $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$. We will see momentarily that this in fact lands in quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$-modules.

Remark 7.28. It follows from Propositions 7.6 and 7.25 that for any étale map $\operatorname{Spec} A \to \mathfrak{X}$, we have

$$\mathcal{F}^{\mathrm{syn}}(\operatorname{Spec} A) \cong \nu(\mathcal{F}(\operatorname{Spec} A)).$$

Combining this with the sheaf property, it follows that if $\mathfrak{Y} \to \mathfrak{X}$ is étale, then we may compute $\mathcal{F}^{\mathrm{syn}}(\mathfrak{Y})$ by choosing an étale cover $\operatorname{Spec} A \to \mathfrak{Y}$, letting $\operatorname{Spec} A^{[\bullet]}$ denote the resulting Čech nerve, and using

$$\mathcal{F}^{\mathrm{syn}}(\mathfrak{Y}) \cong \mathrm{Tot}(\nu\mathcal{F}(\operatorname{Spec} A^{[\bullet]})).$$

Beware that in general this differs from, and is preferable to, $\nu(\mathcal{F}(\mathfrak{Y}))$.

Our main object of study is the global sections of $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$, which we will denote by either $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(\mathfrak{X})$ or $\Gamma(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}})$.

7.3.1 Basic properties

First, we collect some basic properties of this sheaf of synthetic spectra. Let us fix an even-periodic refinement $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ for now.

Proposition 7.29. *Let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module. Then the functor $\mathcal{F}^{\mathrm{syn}}$ takes values in νMU-local and τ-complete synthetic spectra.*

Proof. The classes of νMU-local and τ-complete synthetic spectra are both closed under limits, so it suffices to check the statement on affine schemes. For every étale map $\operatorname{Spec} A \to \mathfrak{X}$, the spectrum $\mathcal{F}(\operatorname{Spec} A)$ is a module over $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}(\operatorname{Spec} A)$. The latter is a complex oriented ring spectrum, so it follows that $\mathcal{F}(\operatorname{Spec} A)$ admits a homotopy MU-module structure. This means its synthetic analogue is both νMU-local and τ-complete, where the latter property follows from Theorem 4.71. ∎

Next, we discuss to quasi-coherence. The following subtlety will make another appearance in Section 7.4.1.

Remark 7.30. As $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$ takes values in νMU-local synthetic spectra, one could define the ∞-category

$$\mathrm{QCoh}(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}})$$

in the sense of Definition 7.13 by taking $\mathcal{C}$ to be either $\mathrm{Syn}_{\mathrm{MU}}$ or $\widehat{\mathrm{Syn}}_{\mathrm{MU}}$ (the νMU-local subcategory; see Notation 4.54). The two resulting symmetric monoidal ∞-categories coincide. Indeed, on affines, the values of $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$ admit homotopy νMU-algebra structures, so any module over $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$ is locally νMU-local. The property of being νMU-local is closed under limits, so the sheaf condition implies that every module over $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$ with values in $\mathcal{C} = \mathrm{Syn}_{\mathrm{MU}}$ automatically lands in the νMU-local subcategory. In particular, the global sections functor

$$\Gamma\colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}) \longrightarrow \mathrm{Syn}_{\mathrm{MU}}$$

lands in the νMU-local subcategory in $\widehat{\mathrm{Syn}}_{\mathrm{MU}}$.

Let $\mathcal{F}$ be a module over $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$. It follows that $\mathcal{F}^{\mathrm{syn}}$ is naturally a module over $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$.

Proposition 7.31. *Let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module. Then $\mathcal{F}^{\mathrm{syn}}$ is a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}$-module.*

Proof. We ought to show that for all maps $V \to U$ of affine schemes over $\mathfrak{X}$, the map

$$\mathcal{F}^{\mathrm{syn}}(U) \otimes_{\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(U)} \mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(V) \longrightarrow \mathcal{F}^{\mathrm{syn}}(V)$$

is an isomorphism. By Remark 7.28, this is the natural map

$$\nu(\mathcal{F}(U)) \otimes_{\nu(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}(U))} \nu(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}(V)) \longrightarrow \nu(\mathcal{F}^{\mathrm{top}}(V)).$$

Using a Tor spectral sequence internal to $\mathrm{Syn}_{\mathrm{MU}}$, this immediately follows from the fact that

$$\mathcal{F}(U) \otimes_{\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(U)} \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(V) \xrightarrow{\cong} \mathcal{F}^{\mathrm{top}}(V)$$

is an isomorphism, combined with Proposition 4.65. ■

As a result, the construction of Definition 7.27 naturally lifts to a functor

$$\mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}) \longrightarrow \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}), \quad \mathcal{F} \longmapsto \mathcal{F}^{\mathrm{syn}}.$$

7.3.2 The signature

Computing the signature of $\mathcal{F}^{\mathrm{syn}}$ proceeds in much the same way as the computation of the signature of ν from Section 4.4. In fact, the following result was used implicitly in Construction 4.70.

Proposition 7.32. *Let $X^\bullet$ be a cosimplicial spectrum. If each X^i admits the structure of a homotopy E-module, then there is an isomorphism of filtered spectra*

$$\sigma(\mathrm{Tot}(\nu X^\bullet)) \cong \mathrm{Déc}^\Delta X^\bullet.$$

In particular, if each X^i admits a homotopy E-module structure, then the signature spectral sequence of $\mathrm{Tot}(\nu X^\bullet)$ is isomorphic to the Tot spectral sequence for $X^\bullet$. Moreover, this isomorphism is naturally a lax symmetric monoidal transformation in $X^\bullet$.

Proof. Using that σ preserves limits and using Corollary 4.67, we obtain natural monoidal isomorphisms of filtered spectra

$$\sigma(\mathrm{Tot}(\nu X^\bullet)) \cong \mathrm{Tot}(\sigma(\nu X^\bullet)) \cong \mathrm{Tot}(\mathrm{Wh}\, X^\bullet) = \mathrm{Déc}^\Delta X^\bullet.$$ ■

Corollary 7.33. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be an even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}$-module. Let $f\colon \mathfrak{Y} \to \mathfrak{X}$ be an étale map of Deligne–Mumford stacks. Then there is an isomorphism of filtered spectra, as lax symmetric monoidal natural transformation in $\mathcal{F}$,*

$$\sigma\, \mathcal{F}^{\mathrm{syn}}(\mathfrak{Y}) \cong \mathrm{DSS}(\mathfrak{X},\, f_*\mathcal{F}).$$

In particular, the signature spectral sequence of $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ is isomorphic to the DSS for $(\mathfrak{X},\, \mathcal{F})$, and we have an isomorphism

$$\pi_{n,s}(\mathcal{F}^{\mathrm{syn}}(\mathfrak{Y})/\tau) \cong \mathrm{H}^s(\mathfrak{X},\, f_*\, \pi_{n+s}\, \mathcal{F}).$$

Bear in mind that this is the version of the descent spectral sequence where we have applied décalage by definition.

Proof. This is a special case of Proposition 7.32, using Remark 7.28. The final isomorphism follows from combining this with Lemma 7.23 (2). ■

Recall the notion of evenness from Section 5.3.

Corollary 7.34. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be an even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}$-module.*

(1) *Suppose that the homotopy sheaves $\pi_k \mathcal{F}$ vanish when k is odd. Then the synthetic spectrum $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ is even.*

(2) *Suppose that for every étale map* $\operatorname{Spec} A \to \mathfrak{X}$, *the homology* $\mathrm{MU}_*(\mathcal{F}(\operatorname{Spec} A))$ *is concentrated in even degrees. Then the sheaf $\mathcal{F}^{\mathrm{syn}}$ takes values in even synthetic spectra.*

In particular, the sheaf $\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$ takes values in even synthetic spectra.

Proof. For the first claim, we use the isomorphism

$$\pi_{n,s}(\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})/\tau) \cong \mathrm{H}^s(\mathfrak{X}, \pi_{n+s}\mathcal{F}),$$

from which it follows that $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ satisfies condition (b) of Proposition 5.30.

For the second claim, recall from Corollary 5.35 that the class of even MU-synthetic spectra is closed under limits (bear in mind that MU-synthetic spectra are cellular; see Section 5.1). By Corollary 5.33, the assumption translates to $\nu \circ \mathcal{F}$ being even on affine schemes over $\mathfrak{X}$, which by Remark 7.28 is the value of $\mathcal{F}^{\mathrm{syn}}$ on affines. ∎

Remark 7.35. The notion of evenness is a first indication that the descent spectral sequence is preferable to the ANSS: in general, $\nu(\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(\mathfrak{X}))$ need not be even, despite $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}$ being an even-periodic sheaf. Geometrically, evenness implies that, after p-completion, $\mathcal{F}^{\mathrm{syn}}$ is a sheaf of **C**-motivic spectra (see Theorem 5.40). In terms of calculations, it means that the underlying spectral sequence only has odd-length differentials. We will see in Chapter 9 that this is preferable over the ANSS for $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}$.

Finally, we discuss the difference between $\mathcal{F}^{\mathrm{syn}}$ and the composite functor $\nu \circ \mathcal{F}$ (i.e., without any sheafification applied). Proposition 4.59 tells us that $\mathcal{F}^{\mathrm{syn}}$ is levelwise a synthetic lift of $\mathcal{F}$. More precisely, we obtain the following result; note that we use $\circ$ to denote functor composition, without any sheafification being applied.

Proposition 7.36. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be an even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}$-module.*

(1) *The natural map $\nu \circ \mathcal{F} \to \mathcal{F}^{\mathrm{syn}}$ factors through an isomorphism $\nu \circ \mathcal{F} \cong \tau_{\geqslant 0} \circ \mathcal{F}^{\mathrm{syn}}$.*

(2) *The natural map $\nu \circ \mathcal{F} \to \mathcal{F}^{\mathrm{syn}}$ becomes an isomorphism after applying the functor* $\tau^{-1} \colon \mathrm{Syn}_{\mathrm{MU}} \to \mathrm{Sp}$.

Proof. This is a special case of Proposition 4.59, using Remark 7.28. ∎

Note that this in particular shows that $\nu(\mathcal{F}(\mathfrak{X}))$ is the best approximation of $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ by a synthetic analogue.

We further learn that the functor ν preserves the limit defining $\mathcal{F}(\mathfrak{X})$ if and only if $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ is connective. While this property of being connective may first appear to be a purely abstract condition on a synthetic spectrum, the preceding discussion shows that it has concrete, non-synthetic implications. Indeed, combining Theorem 4.71 and Proposition 7.36 with Corollary 7.33, we learn that if $\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}(\mathfrak{X})$ is connective, then there is a natural monoidal isomorphism of filtered spectra

$$\text{Déc}\,\mathrm{ASS}_{\mathrm{MU}}(\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(\mathfrak{X})) \cong \mathrm{DSS}(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}).$$

In fact, the connectivity condition itself turns out to have a non-synthetic description, and is a purely algebro-geometric condition on $\mathfrak{X}$: this is the goal of the next section; see in particular Theorem 7.52 and Corollary 7.55.

7.4 The homology of synthetic global sections

Recall from Proposition 7.36 that the natural map $\nu(\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(\mathfrak{X})) \to \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}(\mathfrak{X})$ is a connective cover in $\mathrm{Syn}_{\mathrm{MU}}$. To compute the connectivity of the target, and thus determine whether the map is an isomorphism, we need to study the $\nu\mathrm{MU}$-homology of $\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}(\mathfrak{X})$. This is the goal of this section. Our approach is to show that the signature of the synthetic spectrum $\nu\mathrm{MU} \otimes \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}(\mathfrak{X})$ is the descent spectral sequence for the sheaf $\mathrm{MU} \otimes \mathcal{O}^{\mathrm{top}}$. This requires proving that we may commute tensoring with $\nu\mathrm{MU}$ past the synthetic global sections; this is the difficult part, and is the subject of the more technical Section 7.4.1. In the end, we can answer the connectivity question in purely algebro-geometric terms; see Theorem 7.52.

7.4.1 Global sections and colimits

Let $\nu\mathrm{MU} \otimes \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$ denote the sheafification of levelwise tensoring $\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$ with $\nu\mathrm{MU}$. The goal of this section is to identify the global sections of this sheaf with $\nu\mathrm{MU} \otimes \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}(\mathfrak{X})$; in other words, to show that tensoring with $\nu\mathrm{MU}$ preserves the limit defining these global sections. To state the main result, we need some terminology.

Definition 7.37. An even-periodic refinement $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}})$ is **tame** if the underlying map $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ is tame; see [MM15, Definition 2.28] for this relative definition and [AOV08] for further discussion.

Proposition 7.38. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be a tame even-periodic refinement. Let $\mathcal{F}$ be a synthetic quasi-coherent $\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$-module. Let M be a synthetic spectrum that admits a homotopy $\nu\mathrm{MU}$-module structure. Then the natural limit-comparison map*

$$M \otimes \Gamma(\mathfrak{X},\, \mathcal{F}) \longrightarrow \Gamma(\mathfrak{X},\, M \otimes \mathcal{F}) \tag{7.39}$$

is an isomorphism.

We will deduce this as a consequence of a more general result about the interplay of synthetic global sections and colimits. To state this result, we introduce some notation. Recall from Remark 7.30 that every synthetic quasi-coherent sheaf over $\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$ takes values in νMU-local synthetic spectra. However, it will matter for us whether we regard the global sections functor Γ as landing in (see Notation 4.54)

$$\mathrm{Syn}_{\mathrm{MU}} \qquad \text{or} \qquad \widehat{\mathrm{Syn}}_{\mathrm{MU}}.$$

This distinction matters because the inclusion $\widehat{\mathrm{Syn}}_{\mathrm{MU}} \subseteq \mathrm{Syn}_{\mathrm{MU}}$ does not preserve all colimits. To allow for a distinction between these, we will write

$$\hat{\Gamma} \colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{syn}}) \longrightarrow \widehat{\mathrm{Syn}}_{\mathrm{MU}}$$

for the global sections functor considered as landing in νMU-local synthetic spectra. We stress that the values of $\hat{\Gamma}$ and Γ are the same; the difference is only in what we consider the target category to be. We will also write $- \mathbin{\hat{\otimes}} -$ for the νMU-localisation of the tensor product of synthetic spectra.

Theorem 7.40. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be a tame even-periodic refinement. Then the νMU-local synthetic global sections functor*

$$\hat{\Gamma} \colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{syn}}) \longrightarrow \widehat{\mathrm{Syn}}_{\mathrm{MU}}$$

preserves small colimits.

This is an analogue of results of Mathew–Meier [MM15] in the synthetic setting, although our proof will heavily rely on their results and their methods. Before we prove this result, let us show how Proposition 7.38 follows from this.

Proof of Proposition 7.38. Note that both sides of (7.39) are νMU-local: the source is a νMU-module, and the target is local by Remark 7.30. It therefore suffices to show that the map

$$S \mathbin{\hat{\otimes}} \Gamma(\mathfrak{X}, \mathcal{F}) \longrightarrow \Gamma(\mathfrak{X}, S \otimes \mathcal{F}) \tag{7.41}$$

is an isomorphism for all $S \in \widehat{\mathrm{Syn}}_{\mathrm{MU}}$, because the case $S = M$ retrieves the map (7.39). (Note that the sheaf $S \otimes \mathcal{F}$ coincides with $S \mathbin{\hat{\otimes}} \mathcal{F}$, by Remark 7.30.) For the rest of this proof, we therefore work in the νMU-local synthetic category, so that all colimits are implicitly νMU-localised.

Let $\mathcal{C}$ denote the full subcategory of those $S \in \widehat{\mathrm{Syn}}_{\mathrm{MU}}$ for which (7.41) is an isomorphism. This contains the bigraded spheres $\mathbf{S}^{n,s}$ for all n and s. We next claim that $\mathcal{C}$ is closed under colimits. Suppose we have a diagram S_α in $\mathcal{C}$. Then the map (7.41) in the case $S = \operatorname{colim} S_\alpha$ can be factored as

$$\operatorname*{colim}_{\alpha} S_\alpha \mathbin{\hat{\otimes}} \hat{\Gamma}(\mathfrak{X}, \mathcal{F}) \longrightarrow \operatorname*{colim}_{\alpha} \hat{\Gamma}(\mathfrak{X}, S_\alpha \mathbin{\hat{\otimes}} \mathcal{F}) \longrightarrow \hat{\Gamma}(\mathfrak{X}, \operatorname*{colim}_{\alpha} S_\alpha \mathbin{\hat{\otimes}} \mathcal{F}).$$

The first of these maps is an isomorphism, being a colimit of isomorphisms, and the second is an isomorphism as a consequence of Theorem 7.40.

The ∞-category $\widehat{\mathrm{Syn}}_{\mathrm{MU}}$ is cellular in the sense that it is generated under colimits by the bigraded spheres. Hence $\mathcal{C}$ is equal to all of $\widehat{\mathrm{Syn}}_{\mathrm{MU}}$, finishing the argument. ■

Remark 7.42. Looking at the proof, we see that we proved a stronger result: we showed that for an arbitrary $S \in \mathrm{Syn}_{\mathrm{MU}}$, the natural map

$$S \otimes \Gamma(\mathfrak{X}, \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}) \longrightarrow \Gamma(\mathfrak{X}, S \otimes \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}})$$

is νMU-localisation.

The rest of this subsection is devoted to the proof of Theorem 7.40. We begin by pointing out why we go to the trouble of distinguishing between the νMU-local and the ordinary synthetic global sections: considered as landing in the entire category $\mathrm{Syn}_{\mathrm{MU}}$, the global sections do *not* preserve colimits.

Warning 7.43. Consider the even-periodic refinement $\mathrm{Spec}(\mathrm{KU})/C_2$ of the affine map $\mathrm{B}C_2 \to \mathfrak{M}_{\mathrm{fg}}$. The corresponding ∞-category of quasi-coherent synthetic sheaves may be identified with $\mathrm{Mod}_{\nu\mathrm{KU}}(\mathrm{Syn}_{\mathrm{MU}})^{hC_2}$, as in [MM15, Proposition 2.16]. We claim that the (non-localised) global sections functor

$$\Gamma\colon \mathrm{Mod}_{\nu\mathrm{KU}}(\mathrm{Syn}_{\mathrm{MU}})^{hC_2} \longrightarrow \mathrm{Mod}_{\nu\mathrm{KO}}(\mathrm{Syn}_{\mathrm{MU}})$$

does not preserve all colimits, as can be seen by the non-nilpotence of $\eta \in \pi_{1,1}\, \nu\mathrm{KO}$. Indeed, the unit $\mathcal{O}^{\mathrm{syn}} \in \mathrm{Mod}_{\nu\mathrm{KU}}(\mathrm{Syn}_{\mathrm{MU}})^{hC_2}$ has a class $\eta \in \pi_{1,1}\, \mathcal{O}^{\mathrm{syn}}$ lifting $\eta \in \pi_{1,1}\, \nu\mathrm{KO}$ along the global sections functor, by linearity over the synthetic sphere. The forgetful functor

$$\mathrm{Mod}_{\nu\mathrm{KU}}(\mathrm{Syn}_{\mathrm{MU}})^{hC_2} \longrightarrow \mathrm{Mod}_{\nu\mathrm{KU}}(\mathrm{Syn}_{\mathrm{MU}})$$

is conservative, preserves colimits, and sends η to 0, hence

$$\mathcal{O}^{\mathrm{syn}}[\eta^{-1}] = 0 \quad \text{in } \mathrm{Mod}_{\nu\mathrm{KU}}(\mathrm{Syn}_{\mathrm{MU}})^{hC_2}.$$

The colimit-comparison map

$$\Gamma(\mathcal{O}^{\mathrm{syn}})[\eta^{-1}] = (\nu\mathrm{KO})[\eta^{-1}] \longrightarrow \Gamma(\mathcal{O}^{\mathrm{syn}}[\eta^{-1}])$$

therefore cannot be an isomorphism, since $\eta \in \pi_{1,1}\, \nu\mathrm{KO}$ is not nilpotent. Note however that the νMU-localisation of $(\nu\mathrm{KO})[\eta^{-1}]$ *does* vanish, since η induces the zero map on νMU-homology.

Warning 7.44. A similar analysis shows that the synthetic global sections functor

$$\Gamma\colon \mathrm{QCoh}(\mathfrak{M}_{\mathrm{ell}}, \mathcal{O}^{\mathrm{syn}}) \longrightarrow \mathrm{Mod}_{\mathcal{O}^{\mathrm{syn}}(\mathfrak{M}_{\mathrm{ell}})}(\mathrm{Syn}_{\mathrm{MU}})$$

is *not* an equivalence. Again, one reason is that $\pi_{1,1}\, \mathcal{O}^{\mathrm{syn}}(\mathfrak{M}_{\mathrm{ell}})$ contains the nonnilpotent element η, a fact which will follow from our later computations in Section 8.5.

Proving Theorem 7.40 can be done one diagram at a time, where it boils down to checking that the colimit-comparison map is an isomorphism. This we can do after inverting and after killing τ, which we do in Lemma 7.45 and Lemma 7.50 respectively. The τ-invertible case reduces to the spectral case considered by Mathew–Meier. We write Γ^{top} for the global sections functor

$$\Gamma^{\mathrm{top}} \colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}) \longrightarrow \mathrm{Sp}.$$

Lemma 7.45. *Let $\mathcal{F} \colon I \to \mathrm{QCoh}(\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}})$ be a diagram. Then the colimit-comparison map*

$$\operatorname{colim} \hat{\Gamma}(\mathcal{F}) \longrightarrow \hat{\Gamma}(\operatorname{colim} \mathcal{F})$$

becomes, after τ-inversion, the map

$$\operatorname{colim} \Gamma^{\mathrm{top}}(\mathcal{F}^{\tau=1}) \longrightarrow \Gamma^{\mathrm{top}}(\operatorname{colim} \mathcal{F}^{\tau=1}).$$

Moreover, if $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ is tame, then this last map is an isomorphism.

Proof. The first claim follows from Proposition 4.59 and by using that the τ-inversion is symmetric monoidal. (Note that symmetric monoidality also implies that it preserves quasi-coherence.) The second claim follows from [MM15, Theorem 4.14]. ∎

The case of killing τ is more delicate, because the synthetic global sections do not preserve all colimits by Warnings 7.43 and 7.44. This is where it becomes important that we work $\nu\mathrm{MU}$-locally. Our proof for the mod τ version is inspired by the proofs of [MM15, Propositions 3.9 and 4.11].

We begin by giving a more concrete description of $C\tau \otimes \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$.

Notation 7.46. Write $\mathcal{G}$ for the functor

$$\mathcal{G} \colon \mathrm{Sp} \longrightarrow \mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}, \quad X \longmapsto \mathrm{MU}_*(X).$$

We can consider $\mathcal{G}(X)$ as an object of $\mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}})$, and we will do so without altering the notation. If $\mathcal{O}$ is a sheaf of spectra on a stack $\mathfrak{X}$, then we denote by $\mathcal{G}(\mathcal{O})$ the sheafification in $\mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}})$ of the composite

$$(\mathrm{DM}^{\text{ét}}_{/\mathfrak{X}})^{\mathrm{op}} \xrightarrow{\ \mathcal{O}\ } \mathrm{Sp} \xrightarrow{\ \mathcal{G}\ } \mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}).$$

Lemma 7.47. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be an even-periodic refinement. Then the functors $\mathcal{G}(\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ and $C\tau \otimes \mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}$ are naturally isomorphic. Moreover, via this isomorphism, the functor $C\tau \otimes \hat{\Gamma}$ can be identified with the global sections functor*

$$\Gamma \colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{G}(\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})) \longrightarrow \mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}})$$

in the sense of Definition 7.13 for $\mathcal{C} = \mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}})$.

Proof. Both follow from the fact that the functor $C\tau \otimes -$ is symmetric monoidal and preserves limits (as $C\tau$ is dualisable), and that by Proposition 4.47 we have a natural isomorphism $C\tau \otimes \nu(-) \cong \mathcal{G}$. ■

Let $f\colon \mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ be the map underlying our even-periodic refinement, let L denote the Lazard ring, and let $q\colon \operatorname{Spec} L \to \mathfrak{M}_{\mathrm{fg}}$ denote the cover. We form the pullback

$$\begin{tikzcd} \mathfrak{Y} \arrow[r, "\bar{f}"] \arrow[d, "\bar{q}"'] \arrow[dr, phantom, "\lrcorner", very near start] & \operatorname{Spec} L \arrow[d, "q"] \\ \mathfrak{X} \arrow[r, "f"] & \mathfrak{M}_{\mathrm{fg}}. \end{tikzcd} \tag{7.48}$$

We remind the reader that q is a faithfully flat and affine map of stacks. The map $\bar{q}$ is therefore faithfully flat and affine as well. As the map $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ is tame by assumption, it follows that $\mathfrak{Y}$ is of finite cohomological dimension; see [MM15, Proposition 2.29].

The key step now is to show that $\mathcal{G}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ is pushed forward (in the derived sense) from $\mathfrak{Y}$. In the following statement, the pullback and pushforward functors denote their derived variants, landing in the derived ∞-category of graded comodules. Let us also write ω for the usual line bundle on $\mathfrak{M}_{\mathrm{fg}}$ from Notation 7.14.

Lemma 7.49. *With $\mathfrak{X}$ as above, there are isomorphisms of quasi-coherent $\mathcal{O}_{\mathfrak{X}}$-modules*

$$\mathcal{G}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}) \cong \bar{q}_*\bar{q}^*\,\omega_{\mathfrak{X}}^{\otimes */2}.$$

Proof. We will prove a chain of isomorphisms

$$\mathcal{G}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}) \cong f^*q_*\,\omega_{\operatorname{Spec} L}^{\otimes */2} \cong \bar{q}_*\bar{f}^*\,\omega_{\operatorname{Spec} L}^{\otimes */2} \cong \bar{q}_*\bar{q}^*\,\omega_{\mathfrak{X}}^{\otimes */2}.$$

For the first isomorphism, both sides are sheaves, so it suffices to show they agree on affines. As f is flat and q is flat and affine, f^* and q_* are exact functors, so the derived sheaf $f^*q_*\,\omega_{\operatorname{Spec} L}^{\otimes */2}$ sends an affine $\operatorname{Spec} R \to \mathfrak{X}$ to the (2-periodic) $\mathrm{MU}_*\mathrm{MU}$-comodule algebra appearing as the pullback

$$\operatorname{Spec} R \times_{\mathfrak{M}_{\mathrm{fg}}} \operatorname{Spec} L$$

with its comodule structure from the pullback descent datum for the cover q, regarded as a cochain complex concentrated in cohomological degree zero. This pullback is computed as MU_*A by [MM15, Proposition 2.4] where A is the Landweber theory associated to the flat map $\operatorname{Spec} R \to \mathfrak{M}_{\mathrm{fg}}$.

For the second isomorphism, we refer to [SAG, Corollary 2.5.4.6], which says that the canonical Beck–Chevalley transformation $f^*q_* \to \bar{q}_*\bar{f}^*$ is an isomorphism. The third isomorphism is clear from the isomorphism $\bar{f}^*q^* \cong \bar{q}^*f^*$ following from the commutative diagram (7.48). ■

Lemma 7.50. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be as above. If $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ is tame, then the functor*

$$\Gamma\colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{G}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})) \longrightarrow \mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}).$$

preserves small colimits.

The proof uses manipulations of sheaf cohomology. This is also the reason why we have to restrict to the νMU-local subcategory: in the derived ∞-category of comodules, homotopy groups compute sheaf cohomology, which is not true in general for the stable comodule ∞-category.

Proof. The functor is an exact functor between stable ∞-categories, so it preserves finite colimits. It suffices therefore to show that it preserves small coproducts. We follow closely the argument of [MM15, Proposition 3.9], and we make use of the descent spectral sequence internal to synthetic spectra. If $\mathcal{F}$ is a quasi-coherent sheaf over $\mathcal{G}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$, then this is of the form

$$\mathrm{E}_2^{n,w,s} = \mathrm{H}^s(\mathfrak{X},\, \pi_{n+s,\,w-s}(\mathcal{F})) \implies \pi_{n,w}\,\Gamma(\mathfrak{X},\, \mathcal{F}). \tag{7.51}$$

The E_2-page of this SS commutes with coproducts in $\mathcal{F}$ since $\mathfrak{X}$ is quasi-compact and separated (see [MM15, Lemma 3.10]), hence so does the E_∞-page. If the E_∞-page is concentrated in finitely many rows, then the filtrations are finite, and it follows that the natural map

$$\bigoplus_i \Gamma(\mathfrak{X},\, \mathcal{F}_i) \longrightarrow \Gamma(\mathfrak{X},\, \bigoplus_i \mathcal{F}_i)$$

induces an isomorphism on bigraded homotopy groups and is therefore an isomorphism.

It suffices to show that the E_2-page of (7.51) is concentrated in finitely many rows $s < N$ for fixed N independent of $\mathcal{F}$. From Lemma 7.49, we learn that the homotopy sheaves of $\mathcal{F}$ are pushed forward from $\mathfrak{Y}$:

$$\begin{aligned} \pi_{*,*}(\mathcal{F}) &\cong \pi_{*,*}(\mathcal{F} \mathbin{\hat{\otimes}} \mathcal{G}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})) \\ &\cong \pi_{*,*}(\mathcal{F} \mathbin{\hat{\otimes}} \bar{q}_* \bar{f}^*\, \omega_{\operatorname{Spec} L}^{\otimes */2}) \\ &\cong \pi_{*,*}(\bar{q}_*(\bar{q}^*\mathcal{F} \mathbin{\hat{\otimes}} \bar{f}^*\, \omega_{\operatorname{Spec} L}^{\otimes */2})) \\ &\cong \bar{q}_*\, \pi_{*,*}(\bar{q}^*\mathcal{F} \mathbin{\hat{\otimes}} \bar{f}^*\, \omega_{\operatorname{Spec} L}^{\otimes */2})). \end{aligned}$$

Here the first isomorphism is Lemma 7.49, the second is the projection formula, and the third follows since $\bar{q}$ is affine and flat. Finally, since $\bar{q}$ is affine and flat, it follows (see [Car22, Lemma 6.2.10] for example) that

$$\mathrm{H}^s(\mathfrak{X},\, \bar{q}_*\, \pi_{*,*}(\bar{q}^*\mathcal{F} \mathbin{\hat{\otimes}} \bar{f}^*\, \omega_{\operatorname{Spec} L}^{\otimes */2})) \cong \mathrm{H}^s(\mathfrak{Y},\, \pi_{*,*}(\bar{q}^*\mathcal{F} \mathbin{\hat{\otimes}} \bar{f}^*\, \omega_{\operatorname{Spec} L}^{\otimes */2})).$$

This proves the claim, because $\mathfrak{Y}$ has finite cohomological dimension. ∎

Proof of Theorem 7.40. Let $\mathcal{F}\colon I \to \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}})$ be a diagram. The natural colimit-comparison map

$$\operatorname{colim} \hat{\Gamma}(\mathcal{F}) \longrightarrow \hat{\Gamma}(\operatorname{colim} \mathcal{F})$$

becomes an isomorphism after inverting τ and after tensoring with $C\tau$ by Lemmas 7.45 and 7.50, respectively. By Proposition 3.79, this implies that it is an isomorphism. ■

7.4.2 Comparing the DSS and the ANSS

The connection to the algebraic geometry of $\mathfrak{X}$ allows us to answer the question of the connectivity of $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(\mathfrak{X})$ in purely algebro-geometric terms. More specifically, the fact that $\mathfrak{Y}$ has finite cohomological dimension implies that $\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(\mathfrak{X})$ is bounded below, and we can in fact determine the precise bound, as follows.

Theorem 7.52. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be a tame even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module. Let $\mathfrak{Y}$ and $\bar{q}\colon \mathfrak{Y} \to \mathfrak{X}$ denote the pullback as in* (7.48). *Then for every $n \geqslant 0$, the following are equivalent:*

(a) *the synthetic spectrum $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ is $(-n)$-connective;*

(b) *the sheaf cohomology groups $\mathrm{H}^s(\mathfrak{Y}, \bar{q}^*\, \pi_t\mathcal{F})$ vanish whenever $s > n$.*

In particular, $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ is $(-N)$-connective, where N is the cohomological dimension of $\mathfrak{Y}$.

The proof requires some preparation.

Lemma 7.53. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ and $\mathcal{F}$ be as above. Then there is an isomorphism of quasi-coherent $\mathcal{O}_{\mathfrak{X}}$-modules*

$$\pi_t(\mathrm{MU} \otimes \mathcal{F}) \cong \bar{q}_*\bar{q}^*(\pi_t\, \mathcal{F}).$$

Proof. If A is a Landweber exact ring spectrum and M is a homotopy A-module, then Quillen's theorem implies that there is an isomorphism of graded $\mathrm{MU}_*\mathrm{MU}$-comodules

$$\mathrm{MU}_*(M) \cong \mathrm{MU}_*\mathrm{MU} \otimes_{\mathrm{MU}_*} M_*.$$

Applying this on affines to $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ and $\mathcal{F}$ yields the result. ■

The results from the previous section imply the following.

Proposition 7.54. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be a tame even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module. Then there is a natural isomorphism of filtered spectra* (*as lax symmetric monoidal functors in $\mathcal{F}$*)

$$\sigma(\nu\mathrm{MU} \otimes \mathcal{F}^{\mathrm{syn}}(\mathfrak{X})) \cong \mathrm{DSS}(\mathfrak{X}, \mathrm{MU} \otimes \mathcal{F}).$$

Moreover, the underlying spectral sequence converges strongly, and is of the form

$$\mathrm{E}_2^{n,s} \cong \mathrm{H}^s(\mathfrak{Y}, \bar{q}^*\, \pi_{n+s}\, \mathcal{F}) \implies \mathrm{MU}_n(\mathcal{F}(\mathfrak{X})).$$

Proof. We first claim that the sheaves of synthetic spectra

$$\nu\mathrm{MU} \otimes \mathcal{F}^{\mathrm{syn}} \qquad \text{and} \qquad (\mathrm{MU} \otimes \mathcal{F})^{\mathrm{syn}}$$

are naturally isomorphic. For this, it is enough to construct an isomorphism between their restrictions to $\mathrm{Aff}^{\text{ét}}_{/\mathfrak{X}}$; see Proposition 7.6. On affine schemes, the sheaf $\mathcal{F}^{\mathrm{syn}}$ is given by $\nu \circ \mathcal{F}$ and likewise for $(\mathrm{MU} \otimes \mathcal{F})^{\mathrm{syn}}$ (see Remark 7.28), so that the natural isomorphism of Example 4.10 provides the isomorphism we need. Combining this with Proposition 7.38, we obtain a chain of isomorphisms

$$\nu\mathrm{MU} \otimes \mathcal{F}^{\mathrm{syn}}(\mathfrak{X}) \cong \Gamma(\mathfrak{X},\, \nu\mathrm{MU} \otimes \mathcal{F}^{\mathrm{syn}}) \cong \Gamma(\mathfrak{X},\, (\mathrm{MU} \otimes \mathcal{F})^{\mathrm{syn}}).$$

The first result then follows from Corollary 7.33.

Note that the spectral sequence converges conditionally by Proposition 7.29. To prove strong convergence, it suffices to show that it has a horizontal vanishing line, as this implies collapse at a finite page. Using the isomorphism from Lemma 7.53, and [Car22, Lemma 6.2.10] to see that the resulting sheaf cohomology can be computed over $\mathfrak{Y}$, we see that the second page of the spectral sequence is given by

$$\mathrm{E}_2^{n,s} \cong \mathrm{H}^s(\mathfrak{X},\, \bar{q}_*\bar{q}^*(\pi_{n+s}\,\mathcal{F})) \cong \mathrm{H}^s(\mathfrak{Y},\, \bar{q}^*\, \pi_{n+s}\,\mathcal{F}).$$

Since $\mathfrak{X}$ is tame, the stack $\mathfrak{Y}$ has finite cohomological dimension by [MM15, Proposition 2.29], so the DSS for $\mathrm{MU} \otimes \mathcal{F}$ has a horizontal vanishing line. Finally, the colimit of this filtration is $\mathrm{MU} \otimes \mathcal{F}(\mathfrak{X})$ by Proposition 7.38. ■

Proof of Theorem 7.52. Recall that an MU-synthetic spectrum S is $(-n)$-connective if and only if $\nu\mathrm{MU}_{*,s}(S)$ vanishes for all $s > n$. Proposition 7.54 gives us the signature of $\nu\mathrm{MU} \otimes \mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$, telling us that the $\mathbf{Z}[\tau]$-module $\nu\mathrm{MU}_{*,*}(\mathcal{F}^{\mathrm{syn}}(\mathfrak{X}))$ captures the DSS for $\mathrm{MU} \otimes \mathcal{F}$. More precisely, by the strong convergence of this spectral sequence, the Omnibus Theorem applies. By Theorem 4.77 (4), it follows that (a) holds if and only if the DSS for $\mathrm{MU} \otimes \mathcal{F}$ has no permanent cycles in filtrations $s > n$.

Because the second page of this DSS is given by

$$\mathrm{E}_2^{n,s} \cong \mathrm{H}^s(\mathfrak{Y},\, \bar{q}^*\, \pi_{n+s}\,\mathcal{F}),$$

we see that (b) is equivalent to the vanishing of the second page of the DSS in filtrations $s > n$. Therefore (b) clearly implies (a). To prove the converse, suppose that the group $\mathrm{H}^s(\mathfrak{Y},\, \bar{q}^*\, \pi_t\,\mathcal{F})$ does not vanish for certain s and t, where $s > n$. This group appears in the DSS in bidegree $(t-s,\, s)$. If this group contains a permanent cycle, then by the above, (a) does not hold. If, on the other hand, all elements in this degree support differentials, then this in particular means there is an $r \geqslant 2$ such that the group $\mathrm{H}^{s+r}(\mathfrak{Y},\, \bar{q}^*\, \pi_{t+r-1}\,\mathcal{F})$ does not vanish. Since $s + r > n$, we again find that (a) does not hold. ■

It follows that, since $\mathrm{Syn}_{\mathrm{MU}}$ is cellular and σ is therefore conservative, the map[3] obtained by applying σ

$$\mathrm{D\acute{e}c}^{\Delta}\, \mathrm{ASS}^{\Delta}_{\mathrm{MU}}(\mathcal{F}(\mathfrak{X})) \longrightarrow \mathrm{DSS}(\mathfrak{X}, \mathcal{F})$$

is an isomorphism if *and only if* the sheaf cohomology groups $\mathrm{H}^s(\mathfrak{Y}, \bar{q}^* \pi_t \mathcal{F})$ vanish whenever $s > 0$.

There are some practical situations where the conditions of Theorem 7.52 are satisfied.

Corollary 7.55. *Let $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ be an even-periodic refinement such that the underlying map $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ is affine. Let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}$-module. Then the natural map of synthetic spectra*

$$\nu(\mathcal{F}(\mathfrak{X})) \longrightarrow \mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$$

is an isomorphism. In particular, this induces an isomorphism of filtered spectra

$$\mathrm{D\acute{e}c}^{\Delta}\, \mathrm{ASS}^{\Delta}_{\mathrm{MU}}(\mathcal{F}(\mathfrak{X})) \cong \mathrm{DSS}(\mathfrak{X}, \mathcal{F}).$$

Proof. As the map $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ is affine, the pullback $\mathfrak{Y} = \mathfrak{X} \times_{\mathfrak{M}_{\mathrm{fg}}} \operatorname{Spec} L$ is an affine scheme. In particular, $\mathfrak{Y}$ has cohomological dimension zero. Affine maps are tame, so Theorem 7.52 applies, showing that the synthetic spectrum $\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$ is connective. By Proposition 7.36, we see that the natural map of synthetic spectra

$$\nu(\mathcal{F}(\mathfrak{X})) \longrightarrow \mathcal{F}^{\mathrm{syn}}(\mathfrak{X})$$

is an isomorphism. Applying $\sigma\colon \mathrm{CAlg}(\mathrm{Syn}_{\mathrm{MU}}) \to \mathrm{CAlg}(\mathrm{FilSp})$ to this map results in the desired isomorphism of filtered spectra; indeed, the signatures of these synthetic spectra are computed by Theorem 4.71 and Corollary 7.33, respectively. ∎

This identification of spectral sequences, and other similar statements, have appeared in the literature before:

- In [Mat16, Corollary 5.3], Mathew shows that the ANSS for tmf takes the form of a descent spectral sequence (or rather, its E_2-page is given by the cohomology of a particular Hopf algebroid). For a further discussion, see Chapter 11.
- In unpublished works, Meier [Mei17, Theorem 4.6] and Devalapurkar [Dev18, Corollary 5] prove the above statement as an isomorphism of spectral sequences, but without addressing the multiplicative structures.

[3] Although the notation appears to suggest otherwise, remember that by definition we include an application of décalage in $\mathrm{DSS}(\mathfrak{X}, \mathcal{F})$; see Remark 7.22.

7.5 Inverting elements

In this short section, we discuss the problem of commuting synthetic global sections with the inversion of an element. This is not possible in general, as evidenced by Warnings 7.43 and 7.44. It is possible if the element is detected in filtration 0, because in this case the map is detected in νMU-homology. Again, we rely heavily on the results of Mathew–Meier [MM15].

Proposition 7.56. *Let $(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}})$ be an even-periodic refinement, and let $\mathcal{F}$ be a quasi-coherent $\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$-module.*

(1) *Let $f \in \mathrm{H}^0(\mathfrak{X}, \omega_{\mathfrak{X}}^{\otimes d})$ be an element, and write $D(f)$ for its nonvanishing locus on $\mathfrak{X}$. Regarding f as an element of $\pi_{2d,0}(\mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(\mathfrak{X})/\tau)$, the natural map*

$$(\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})/\tau)[f^{-1}] \xrightarrow{\cong} \mathcal{F}^{\mathrm{syn}}(D(f))/\tau$$

is an isomorphism.

(2) *Let $g \in \pi_{2d,0}\, \mathcal{O}_{\mathfrak{X}}^{\mathrm{syn}}(\mathfrak{X})$ be an element that is detected in filtration 0, and write $\overline{g} \in \mathrm{H}^0(\mathfrak{X}, \omega_{\mathfrak{X}}^{\otimes d})$ for its image under the edge homomorphism. Suppose that the underlying stack $\mathfrak{X}$ is noetherian, and that $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ is tame. Then the natural map*

$$\mathcal{F}^{\mathrm{syn}}(\mathfrak{X})[g^{-1}] \xrightarrow{\cong} \mathcal{F}^{\mathrm{syn}}(D(\overline{g}))$$

is an isomorphism.

Proof. For item (1), we check that the map induces an isomorphism on bigraded homotopy groups. As filtered colimits commute with bigraded homotopy groups, we calculate (using Corollary 7.33) the left-hand side to be the E_2-page of the DSS for $\mathcal{F}(\mathfrak{X})$ with f inverted. This agrees with the E_2-page of the DSS for $\mathcal{F}$ at $D(f)$. Indeed, sheaf cohomology commutes with filtered colimits in the abelian 1-category of quasi-coherent sheaves; see [Stacks, Tag 0GQV]. As f is of filtration zero, it comes from a map in the abelian 1-category of quasi-coherent sheaves. This E_2-page in turn agrees with the homotopy groups of $C\tau \otimes \mathcal{F}^{\mathrm{syn}}(D(f))$, and we see that the map induces the natural isomorphism on bigraded homotopy groups.

For item (2), it suffices to check that the map is an isomorphism after tensoring with $C\tau$ and after τ-inversion; see Proposition 3.79. The former of the two follows from item (1) for $f = \overline{g}$, because tensoring with $C\tau$ preserves colimits. The latter of the two follows from from [MM15, Theorem 4.14]; note that this result requires $\mathfrak{X}$ to be noetherian and the map $\mathfrak{X} \to \mathfrak{M}_{\mathrm{fg}}$ to be tame. Indeed, this result says that the global sections functor

$$\Gamma(\mathfrak{X}, -)\colon \mathrm{QCoh}(\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}) \longrightarrow \mathrm{Mod}_{\mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}(\mathfrak{X})}(\mathrm{Sp})$$

preserves colimits; in particular, we obtain the chain of natural isomorphisms

$$\Gamma(\mathfrak{X}, \mathcal{F})[f^{-1}] \cong \Gamma(\mathfrak{X}, \mathcal{F}[\overline{f}^{-1}]) \cong \Gamma(D(\overline{f}), \mathcal{F}). \qquad \blacksquare$$

7.6 Homotopy fixed-point spectral sequences

Let E be an even-periodic Landweber-exact $\mathbf{E}_\infty$-ring with an action by a finite group G; that is, E is a functor $BG \to \mathrm{CAlg}(\mathrm{Sp})$. Then there is a canonical morphism of stacks

$$\mathrm{Spec}(\pi_0 E)/G \longrightarrow \mathfrak{M}_{\mathrm{fg}}$$

that admits a canonical even-periodic refinement given by $\mathrm{Spec}(E)/G$; see [MM15, Proposition 2.15]. It is well known that the homotopy fixed-point spectral sequence (HFPSS) for E can be naturally identified with the DSS for the global sections of this even-periodic refinement, which we write as $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$ for a moment. Indeed, the cobar resolution computing the homotopy fixed points E^{hG} takes the form of the cosimplicial object

$$E^\bullet = \left(E \rightrightarrows \prod_G E \;\substack{\longrightarrow\\[-2pt]\longrightarrow\\[-2pt]\longrightarrow}\; \prod_{G^2} E \cdots \right) \tag{7.57}$$

so E^{hG} is given by $\mathrm{Tot}\, E^\bullet$. Notice that we have $E = \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(\mathrm{Spec}\, \pi_0 E)$ by definition, and that there are identifications

$$\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}\left(\coprod_{G^n} \mathrm{Spec}\, \pi_0 E \right) \cong \prod_{G^n} E$$

$$\coprod_G \mathrm{Spec}\, \pi_0 E \cong \mathrm{Spec}\, \pi_0 E \times_{\mathrm{Spec}(\pi_0 E)/G} \mathrm{Spec}\, \pi_0 E,$$

the second coming from the fact that $\mathrm{Spec}\, \pi_0 E \to \mathrm{Spec}(\pi_0 E)/G$ is an étale G-torsor. Inductively, we now see that (7.57) is naturally equivalent to the cosimplicial $\mathbf{E}_\infty$-ring $\mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(\mathrm{Spec}(\pi_0 E)^\bullet)$ used to define the DSS for the global sections of $(\mathfrak{X}, \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}})$, where $(\pi_0 E)^\bullet$ comes from the Čech nerve of above G-torsor.

In particular, combining this discussion with Proposition 7.32 we obtain the following corollary.

Corollary 7.58. *Let G be a finite group and let $E \in \mathrm{Fun}(BG, \mathrm{CAlg}(\mathrm{Sp}))$ whose underlying ring spectrum is Landweber exact. Then, using the notation above, there is a natural isomorphism of cosimplicial $\mathbf{E}_\infty$-rings*

$$E^\bullet \cong \mathcal{O}^{\mathrm{top}}_{\mathfrak{X}}(\mathrm{Spec}(\pi_0 E)^\bullet).$$

In particular, the HFPSS for E^{hG} and the DSS for $\mathfrak{X}$ are isomorphic, and the signature of the synthetic $\mathbf{E}_\infty$-ring $(\nu E)^{hG}$ is the HFPSS for E.

Our framework gives a straightforward criterion that allows one to often identify this DSS (and hence also the HFPSS) with the ANSS of E^{hG}.

Theorem 7.59. *Let G be a finite group and let $E \in \mathrm{Fun}(\mathrm{B}G, \mathrm{CAlg}(\mathrm{Sp}))$ whose underlying ring spectrum is Landweber exact. If, for every field-valued point x of $\operatorname{Spec} \pi_0 E$, the subgroup of G stabilising x acts faithfully on the formal group associated to x, then the canonical map*

$$\nu(E^{hG}) \longrightarrow (\nu E)^{hG}$$

is an isomorphism of synthetic $\mathbf{E}_\infty$-rings. In particular, the signature of $(\nu E)^{hG}$ is the ANSS of E^{hG}.

Proof. As remarked in [MM15, Section 6.2], the morphism

$$\operatorname{Spec}(\pi_0 E)/G \longrightarrow \mathfrak{M}_{\mathrm{fg}}$$

is affine precisely when the conditions of the theorem hold. The claim now follows immediately from Theorem 7.52. ■

This theorem is a highly structured lift of the following useful corollary, which we now see is immediate from Theorem 7.52.

Corollary 7.60. *Under the conditions of Theorem 7.59, the ANSS of E^{hG} is isomorphic to the HFPSS of E from the second page on.*

Example 7.61. Fix a height n formal group Γ over a perfect field k of characteristic p, and fix a finite subgroup G of the corresponding Morava stabilizer group $\mathbf{G} := \mathrm{Gal}(k/\mathbf{F}_p) \rtimes \mathrm{Aut}(\Gamma)$. By the Goerss–Hopkins–Miller theorem, see [EC2, Section 5], the even-periodic Landweber-exact $\mathbf{E}_\infty$-ring E_n admits an action of G by $\mathbf{E}_\infty$-ring maps. The conditions of Theorem 7.59 apply, and we get an equivalence of synthetic $\mathbf{E}_\infty$-rings

$$\nu(\mathrm{EO}_n(G)) \cong \nu(\mathrm{E}_n)^{hG},$$

where $\mathrm{EO}_n(G) := \mathrm{E}_n^{hG}$ is a higher real K-theory. ▲

Remark 7.62. In general, regardless of whether or not $\operatorname{Spec}(\pi_0 E)/G \to \mathfrak{M}_{\mathrm{fg}}$ is affine, the signature of the synthetic spectrum $\nu(E)^{hG}$ is the HFPSS of E by Corollary 7.58. This will not generally coincide with the ANSS of E^{hG}, making $\nu(E)^{hG}$ often preferable over $\nu(E^{hG})$ as a synthetic lift of E^{hG}. The unit of the synthetic $\mathbf{E}_\infty$-ring $\nu(E)^{hG}$ still provides a map of multiplicative spectral sequences

$$\mathrm{ANSS}(\mathbf{S}) \longrightarrow \mathrm{HFPSS}(E)$$

suitable for detection arguments, for example. From this point of view, the ANSS of E^{hG} has no advantage over the HFPSS.

Example 7.63. Consider KU with a trivial C_2-action. The spectral Deligne–Mumford stack $\mathrm{Spec}(\mathrm{KU})/C_2$ is an even-periodic refinement of the morphism

$$f\colon \mathrm{B}C_2 \longrightarrow \mathfrak{M}_{\mathrm{fg}}$$

that picks out $\widehat{\mathbf{G}}_m$ with a trivial C_2-action. The morphism f is not faithful, and hence does not satisfy the conditions of Theorem 7.59.

In fact, the corresponding homotopy fixed-point spectral sequence does not coincide with the Adams–Novikov spectral sequence of $\mathrm{KU}^{hC_2} \cong \mathrm{KU}^{\mathrm{B}C_{2+}}$. Indeed, the latter is complex-oriented as a KU-module, and hence its ANSS is concentrated in filtration zero. On the other hand, the cohomology groups $\mathrm{H}^s(C_2, \pi_*\mathrm{KU})$ are nontrivial for positive values of s, hence the homotopy fixed point spectral sequence is nontrivial in positive filtrations.

Since the ANSS of $\mathrm{KU}^{\mathrm{B}C_{2+}}$ collapses on the zero-line, the ANSS provides no computational leverage: computing the E_2-page already requires knowing $\pi_*(\mathrm{KU} \otimes \mathrm{B}C_{2+})$. The HFPSS, on the other hand, has an E_2-page that is algebraically computable using only $\pi_*\mathrm{KU}$ with its C_2-action as input. ▲

Chapter 8

Synthetic modular forms

In this chapter, we define and begin our study of *synthetic modular forms*. The results of the previous chapter allow us to compare these to topological modular forms, showing that SMF is isomorphic to $\nu\,\mathrm{TMF}$, while Smf is different from $\nu\,\mathrm{Tmf}$.

The version Smf is the more fundamental one to study. After summarising the basic properties of SMF and Smf in Sections 8.1 and 8.2, we then focus on Smf alone, and begin our computation of its homotopy groups; in other words, the computation of the descent spectral sequence for Tmf. While the full computation is deferred to the next chapter, we here provide the basics for that computation:

- in Section 8.3, we compute the second page of the DSS for Tmf, in other words, the sheaf cohomology $\mathrm{H}^*(\overline{\mathfrak{M}}_{\mathrm{ell}}, \omega^{\otimes *})$;
- in Section 8.4, we compute a *transfer map* in order to conclude that the modular form $2c_6$ is a permanent cycle;
- in Section 8.5, we show that the unit map $\mathbf{S} \to \mathrm{Smf}$ detects a number of important elements in the sphere in low dimensions.

The reader solely interested in the computation of Smf may, after reading Section 8.1, skip forward to Chapter 9, and refer back to the above results as needed.

8.1 Definition and basic properties

Recall from [DFHH, Section 12] that there exists an even-periodic refinement $\mathcal{O}^{\mathrm{top}}$ on the *compactification of the moduli stack of elliptic curves* $\overline{\mathfrak{M}}_{\mathrm{ell}}$. Using this, we define the $\mathbf{E}_\infty$-rings

$$\mathrm{Tmf} = \mathcal{O}^{\mathrm{top}}(\overline{\mathfrak{M}}_{\mathrm{ell}}) \qquad \text{and} \qquad \mathrm{TMF} = \mathcal{O}^{\mathrm{top}}(\mathfrak{M}_{\mathrm{ell}}),$$

where $\mathfrak{M}_{\mathrm{ell}} \subseteq \overline{\mathfrak{M}}_{\mathrm{ell}}$ is the moduli stack of elliptic curves. See [EC2, Section 7] for an alternative construction of $\mathcal{O}^{\mathrm{top}}$ on $\mathfrak{M}_{\mathrm{ell}}$.

Definition 8.1. We write $\mathcal{O}^{\mathrm{syn}}$ for the étale sheaf of synthetic $\mathbf{E}_\infty$-rings on $\overline{\mathfrak{M}}_{\mathrm{ell}}$ lifting $\mathcal{O}^{\mathrm{top}}$ via Definition 7.27. We define

$$\mathrm{Smf} = \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}}) \qquad \text{and} \qquad \mathrm{SMF} = \mathcal{O}^{\mathrm{syn}}(\mathfrak{M}_{\mathrm{ell}}).$$

Our general results from Chapter 7 specialise to the following.

Proposition 8.2.

(1) *The synthetic* $\mathbf{E}_\infty$*-rings* Smf *and* SMF *are both* νMU*-local,* τ*-complete, and even.*

(2) *The synthetic* $\mathbf{E}_\infty$*-rings* Smf *and* SMF *are synthetic lifts of* Tmf *and* TMF*, respectively.*

(3) *We have isomorphisms of filtered* $\mathbf{E}_\infty$*-rings*

$$\sigma\,\mathrm{Smf} \cong \mathrm{DSS}(\overline{\mathfrak{M}}_{\mathrm{ell}}, \mathcal{O}^{\mathrm{top}}) \qquad \textit{and} \qquad \sigma\,\mathrm{SMF} \cong \mathrm{DSS}(\mathfrak{M}_{\mathrm{ell}}, \mathcal{O}^{\mathrm{top}}).$$

Proof. This follows from Proposition 7.29, Corollary 7.34, Proposition 7.36, and Corollary 7.33, respectively. ■

Later, in Proposition 10.7, we will give an alternative description of Smf in terms of SMF and ν tmf. We postpone this until later because, as explained in Remark 10.8, this requires computational input.

Remark 8.3. We will ignore the case of connective tmf, as there is no clear way to define this object from the spectral algebraic geometry used here. The first main problem is that the stack $\mathfrak{M}_{\mathrm{cub}}$ associated to tmf is not Deligne–Mumford, as the automorphism group of the cuspidal cubic given by $y^2 = x^3$ is not an étale group scheme. The second problem is that the map $\mathfrak{M}_{\mathrm{cub}} \to \mathfrak{M}_{\mathrm{fg}}$ is not flat. To see this, we notice that if it were flat, then $\mathrm{MU}_*\mathrm{tmf}$ would be flat over $\pi_*\mathrm{MU}$, which we can see fails once we know that

$$\mathrm{MU}_*\mathrm{tmf} \cong \mathbf{Z}[a_1, a_2, a_3, a_4, a_6, e_n \mid n \geqslant 4];$$

see [Rez07, Proposition 20.4]. One can of course simply *define* smf to be ν tmf. This definition is partially justified by the observation that the map $\mathfrak{M}_{\mathrm{cub}} \to \mathfrak{M}_{\mathrm{fg}}$ is affine, so that Corollary 7.55 suggests that ν tmf should be the preferred synthetic lift of tmf. (As explained in Example 5.43, the MU-synthetic spectrum ν tmf is the one studied by Gheorge–Isaksen–Krause–Ricka [GIKR21] under the name *motivic modular forms*.) We think it is an interesting problem to reconstruct ν tmf purely from the synthetic $\mathbf{E}_\infty$-ring Smf.

Remark 8.4. In a similar vein, ν ko is a preferred MU-synthetic model for ko as the map $\mathfrak{M}_{\mathrm{quad}} \to \mathfrak{M}_{\mathrm{fg}}$ (where $\mathfrak{M}_{\mathrm{quad}}$ is the stack associated to ko; see [DFHH, Section 9]) is also affine. In this case, Carrick and Davies show [CD24, Proposition 3.14] by direct computation that there are isomorphisms

$$\nu\,\mathrm{KO} \cong (\nu\,\mathrm{KU})^{hC_2} \qquad \text{and} \qquad \nu\,\mathrm{ko} \cong \tau_{\geqslant 0}^{\uparrow}\,\nu\,\mathrm{KO},$$

which further confirm this suspicion that $\nu\,\mathrm{ko}$ is the preferred synthetic lift of ko. We repeat the warning from [CD24, Warning 3.16] that these isomorphisms do *not* hold in $\mathrm{Syn}_{\mathbf{F}_2}$.

Our formalism also extends to synthetic modular forms with level structures.

Variant 8.5. It is also possible to alter Definition 8.1 to accommodate moduli stacks of elliptic curves with *level structures*. Indeed, for a congruence subgroup $\Gamma \subseteq \mathrm{SL}_2(\mathbf{Z})$, write $\mathfrak{M}_\Gamma$ for the moduli stack of elliptic curves with Γ-level structure. The map $\mathfrak{M}_\Gamma \to \mathfrak{M}_{\mathrm{ell}}$ is étale, and we define

$$\mathrm{SMF}_\Gamma := \mathcal{O}^{\mathrm{syn}}(\mathfrak{M}_\Gamma).$$

This is a synthetic $\mathbf{E}_\infty$-ring that lifts TMF_Γ and whose signature is the DSS for TMF_Γ. For compactifications of $\mathfrak{M}_\Gamma$, we need to go to the *log étale* site of $\overline{\mathfrak{M}}_{\mathrm{ell}}$ as in [HL16]. This is due to the structure map $\overline{\mathfrak{M}}_\Gamma \to \overline{\mathfrak{M}}_{\mathrm{ell}}$ generally not being étale, but only log étale. In particular, replacing the small étale site of a Deligne–Mumford stack $\mathfrak{X}$ with its *log étale site*, for a chosen log structure on $\mathfrak{X}$, all of the results of this article follow *mutatis mutandis*. This leads to a synthetic $\mathbf{E}_\infty$-ring

$$\mathrm{Smf}_\Gamma := \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_\Gamma)$$

that, likewise, lifts Tmf_Γ and whose signature is the DSS for Tmf_Γ.

8.2 Comparison to topological modular forms

Our earlier results let us understand SMF in more familiar terms.

Proposition 8.6. *The canonical limit-comparison map of synthetic $\mathbf{E}_\infty$-rings*

$$\nu\,\mathrm{TMF} = \nu(\mathcal{O}^{\mathrm{top}}(\mathfrak{M}_{\mathrm{ell}})) \xrightarrow{\cong} \mathcal{O}^{\mathrm{syn}}(\mathfrak{M}_{\mathrm{ell}}) = \mathrm{SMF}$$

is an isomorphism.

Proof. The map $\mathfrak{M}_{\mathrm{ell}} \to \mathfrak{M}_{\mathrm{fg}}$ defined by the formal group associated to the universal elliptic curve over $\mathfrak{M}_{\mathrm{ell}}$ is affine: see the first paragraph in the proof of [MM15, Theorem 7.2]. The above isomorphism then follows from Corollary 7.55. ∎

In other words, ν preserves the limit $\mathcal{O}^{\mathrm{top}}(\mathfrak{M}_{\mathrm{ell}})$ defining TMF. While the synthetic spectrum SMF is therefore not new, we will continue to denote it by SMF rather than $\nu\,\mathrm{TMF}$, in order to emphasise its construction and the fact that its signature naturally gives the DSS for TMF.

Remark 8.7. For any congruence subgroup Γ, the map of synthetic $\mathbf{E}_\infty$-rings $\nu\,\mathrm{TMF}_\Gamma \to \mathrm{SMF}_\Gamma$ is also an isomorphism. Indeed, the map of stacks $\mathfrak{M}_\Gamma \to \mathfrak{M}_{\mathrm{ell},\mathbf{Z}[\frac{1}{N}]}$ is finite étale, and hence affine, so Corollary 7.55 applies again.

Although the affine map $\mathfrak{M}_{\mathrm{ell}} \to \mathfrak{M}_{\mathrm{fg}}$ extends over the compactification $\overline{\mathfrak{M}}_{\mathrm{ell}}$, the resulting map is *not* affine, but merely quasi-affine. We can then use Theorem 7.52 to determine how far away the map $\nu\,\mathrm{Tmf} \to \mathrm{Smf}$ is from being an isomorphism, or in other words, how far away the ANSS for Tmf is from its DSS.

Proposition 8.8. *The synthetic $\mathbf{E}_\infty$-ring* Smf *is (-1)-connective but not connective. In particular, there is a cofibre sequence of synthetic spectra*

$$\nu\,\mathrm{Tmf} \longrightarrow \mathrm{Smf} \longrightarrow \Sigma^{-1}(\pi_{-1}^{\heartsuit}\mathrm{Smf}).$$

Moreover, viewing $\pi_{-1}^{\heartsuit}(\mathrm{Smf})$ as an element of $\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}$*, we have an isomorphism of graded comodules*

$$\pi_{-1}^{\heartsuit}(\mathrm{Smf}) \cong \mathrm{H}^1(\overline{\mathfrak{M}}_{\mathrm{ell}} \times_{\mathfrak{M}_{\mathrm{fg}}} \operatorname{Spec} L,\ \omega^{\otimes */2}).$$

In particular, Smf is a synthetic lift of Tmf different from $\nu\,\mathrm{Tmf}$, and $\nu\,\mathrm{Tmf}$ is the best approximation to Smf by a synthetic analogue (simply because all synthetic analogues are connective). As explained in Section 6.3 and to be explored in detail in Chapter 9, for computational reasons, Smf is preferable to $\nu\,\mathrm{Tmf}$.

Proof. In [Mat16, Proposition 5.1], Mathew computes the DSS for $\mathrm{MU} \otimes \mathcal{O}_{\mathfrak{X}}^{\mathrm{top}}$ for $\mathfrak{X} = \overline{\mathfrak{M}}_{\mathrm{ell}}$, finding that it is concentrated in filtrations 0 and 1. Therefore Proposition 7.54 shows that $\nu\mathrm{MU}_{*,1}(\mathrm{Smf})$ is nonzero while $\nu\mathrm{MU}_{*,s}(\mathrm{Smf})$ vanishes for all $s \geqslant 2$. More specifically, it shows that we have an isomorphism

$$\nu\mathrm{MU}_{*,1}(\mathrm{Smf}) \cong \mathrm{H}^1(\overline{\mathfrak{M}}_{\mathrm{ell}} \times_{\mathfrak{M}_{\mathrm{fg}}} \operatorname{Spec} L,\ \omega^{\otimes(*+1)/2}).$$

In particular, Smf is (-1)-connective but not connective.

By Proposition 7.36, the cofibre of the map $\nu\,\mathrm{Tmf} \to \mathrm{Smf}$ is the (-1)-truncation $\tau_{\leqslant -1}\,\mathrm{Smf}$. Since Smf is (-1)-connective, this cofibre is therefore the desuspension of the discrete object $\pi_{-1}^{\heartsuit}(\mathrm{Smf})$, considered as an object of $\mathrm{Syn}_{\mathrm{MU}}^{\heartsuit} \simeq \mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}$. The last claim follows from the isomorphism mentioned in Warning 4.27, which says that if S is a synthetic spectrum, then we have an isomorphism of graded $\mathrm{MU}_*\mathrm{MU}$-comodules

$$\nu\mathrm{MU}_{*,s}(S) \cong \pi_{-s}^{\heartsuit}(S)[s],$$

where $[s]$ denotes the s-fold grading shift. ■

Recall that $\mathfrak{M}_{\mathrm{ell}}$ is the nonvanishing locus of the function Δ on $\overline{\mathfrak{M}}_{\mathrm{ell}}$. As a result, using Section 7.5, one might expect that TMF can be obtained from Tmf by inverting

Δ, and likewise for SMF and Smf. While this is true, proving this requires an additional computational input, namely that a suitable power of Δ lifts to an element of $\pi_{*,*}$ Smf (and consequently to π_* Tmf). Because of this, we will prove this result later in Proposition 10.5.

8.3 The second page

We now turn to the computation of $\pi_{*,*}(\mathrm{Smf}/\tau)$, or in other words, the cohomology of the stack $\overline{\mathfrak{M}}_{\mathrm{ell}}$. These computations were done by Bauer [Bau08] and Konter [Kon12]. Our goal in this section then is two-fold: first, to justify the use of these results without running into any circularity issues, and second, to review these results and phrase them in our current synthetic language.

In [Bau08], Bauer computes the cohomology of the moduli stack of *cubic curves*, and uses this as the second page of the ANSS for tmf. Although the latter use of this computation needs a justification for $\mathrm{MU}_*\mathrm{tmf}$, the computation of the cohomology of the moduli of cubic curves is an algebro-geometric computation that is independent of topological considerations. Let us briefly summarise Bauer's setup.

Definition 8.9 ([Bau08], Section 3). The **cubic Hopf algebroid** is the graded Hopf algebroid (A, Γ) given by the data

$$A = \mathbf{Z}[a_1, a_2, a_3, a_4, a_6], \qquad |a_i| = 2i,$$

$$\Gamma = A[r, s, t], \qquad |r| = 4, \quad |s| = 2, \quad |t| = 6,$$

equipped with the left unit $\eta_L \colon A \to \Gamma$ given by the standard inclusion, the right unit $\eta_R \colon A \to \Gamma$ and comultiplication $\psi \colon \Gamma \to \Gamma \otimes_A \Gamma$ defined as in [Bau08, Section 3].

Remark 8.10. The cubic Hopf algebroid is referred to as the *elliptic curve Hopf algebroid* by Bauer in [Bau08].

Bauer gives simplifications of this Hopf algebroid when one inverts 2 or 3; see [Bau08, Section 4]. Using this, in Sections 4, 5 and 7 of op. cit., Bauer computes the cohomology of this graded Hopf algebroid locally at every prime.

Definition 8.11. Recall the equivalence between Hopf algebroids and algebraic stacks, as given in [Nau07, Theorem 8] for example. Write $\mathfrak{M}(A, \Gamma)$ for the stack corresponding to the cubic Hopf algebroid (A, Γ) under this equivalence. The stack $\mathfrak{M}(A, \Gamma)$ carries a $\mathbf{G}_m$-action arising from the grading of (A, Γ). The stack resulting as the quotient of this $\mathbf{G}_m$-action is called the **moduli stack of cubic curves**, which we denote by $\mathfrak{M}_{\mathrm{cub}}$.

Remark 8.12. In [Kon12], Konter writes $\overline{\mathfrak{M}}^+_{\mathrm{ell}}$ for what we call $\mathfrak{M}_{\mathrm{cub}}$.

There is a map $\mathfrak{M}_{\mathrm{cub}} \to \mathfrak{M}_{\mathrm{fg}}$ arising from a map of the corresponding Hopf algebroids; see [Bau08, Section 3]. Taking the derived pushforward of the structure sheaf along this map allows us to view (A, Γ) as defining an $\mathbf{E}_\infty$-algebra in $\mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}})$. Through the equivalence from Theorem 4.53

$$L_{\nu\mathrm{MU}}\mathrm{Mod}_{C\tau}(\mathrm{Syn}_{\mathrm{MU}}) \simeq \mathcal{D}(\mathrm{grComod}_{\mathrm{MU}_*\mathrm{MU}}), \tag{8.13}$$

we can thus regard (A, Γ) as an $\mathbf{E}_\infty$-algebra in $C\tau$-modules in $\mathrm{Syn}_{\mathrm{MU}}$. We will denote this by (A, Γ) again in this section.

Our goal now is to import Bauer's computations about (A, Γ) to Smf/τ. This passage from $\mathfrak{M}_{\mathrm{cub}}$ to $\overline{\mathfrak{M}}_{\mathrm{ell}}$ was done by Konter in [Kon12], which we now briefly recall. As we would also like to import Bauer's Massey products, we are interested in these objects not just as $C\tau$-modules, but rather as $\mathbf{E}_\infty$-algebras over $C\tau$.

There are homogeneous quantities c_4 and Δ in A, see [Sil, page 42], and inverting these elements yields Hopf algebroids $(A[c_4^{-1}], \Gamma[c_4^{-1}])$ and $(A[\Delta^{-1}], \Gamma[\Delta^{-1}])$ and their corresponding stacks $\mathfrak{M}_{\mathrm{cub}}[c_4^{-1}]$ and $\mathfrak{M}_{\mathrm{cub}}[\Delta^{-1}]$.

Proposition 8.14. *There is a pushout diagram of stacks*

$$\overline{\mathfrak{M}}_{\mathrm{ell}} = \mathfrak{M}_{\mathrm{cub}}[c_4^{-1}] \cup_{\mathfrak{M}_{\mathrm{cub}}[c_4^{-1}, \Delta^{-1}]} \mathfrak{M}_{\mathrm{cub}}[\Delta^{-1}],$$

and we have an equivalence of stacks

$$\mathfrak{M}_{\mathrm{cub}}[\Delta^{-1}] \simeq \mathfrak{M}_{\mathrm{ell}}.$$

Proof. See [Kon12, Section 1.4]. ■

Using Bauer's aforementioned computations and the above pushout, Konter [Kon12, Sections 3, 4.1, 4.2, 5.1, and 5.2] computes the cohomology of $\overline{\mathfrak{M}}_{\mathrm{ell}}$ locally at every prime.

Through the equivalence of (8.13), Corollary 7.33 gives an identification of the $\mathbf{E}_\infty$-$C\tau$-algebra Smf/τ with the derived pushforward of the structure sheaf along $\overline{\mathfrak{M}}_{\mathrm{ell}} \to \mathfrak{M}_{\mathrm{fg}}$. As a result, the map $\overline{\mathfrak{M}}_{\mathrm{ell}} \to \mathfrak{M}_{\mathrm{cub}}$ gives rise to a map of $\mathbf{E}_\infty$-$C\tau$-algebras

$$(A, \Gamma) \longrightarrow \mathrm{Smf}/\tau.$$

The aforementioned computations of Bauer and Konter in particular imply the following.

Corollary 8.15 ([Bau08; Kon12]). *The map of* $\mathbf{E}_\infty$-$C\tau$*-algebras*

$$\Phi\colon (A, \Gamma) \longrightarrow \mathrm{Smf}/\tau$$

induces a split injection on bigraded homotopy groups, and even an isomorphism on $\pi_{n,s}$ *when* $5s \leqslant n + 12$.

If one inverts 2, then the map Φ on homotopy groups is an isomorphism in all bidegrees where the cohomology of (A, Γ) does not vanish. This is not the case at the prime 2, due to the presence of h_1-towers; the bulk of these h_1 towers appear outside the range $5s \leqslant n + 12$, explaining the above formulation. We refer to either [Kon12, Figures 10 and 25] or Figures A.1 and A.2 below for a graphic demonstration of this at the primes 3 and 2, respectively.

8.4 Synthetic transfer maps

In this section, we describe how to obtain *transfer maps* from synthetic modular forms with level structure. We then use this to prove that a few elements, most notably $2c_6$, is a permanent cycle in the DSS for Tmf; see Corollary 8.23. We closely follow [Dav24a, Section 1.3] for the setup of these transfers; for other references on algebraic transfers in this context, see [Mei23, Section 5.6] for a discussion in the case of periodic TMF, and see [Stacks, Tag 03SH] for the algebraic theory.

In what follows, we use $\mathcal{H}\mathrm{om}(-,-)$ to denote the internal hom of quasi-coherent sheaves.

Definition 8.16. Let $\mathfrak{X}$ be a Deligne–Mumford stack, and let $\mathcal{A}$ be a quasi-coherent sheaf of algebras on $\mathfrak{X}$ that is locally free of finite rank. Note that the natural map $\mathcal{H}\mathrm{om}_{\mathcal{O}_{\mathfrak{X}}}(\mathcal{A}, \mathcal{O}_{\mathfrak{X}}) \otimes \mathcal{A} \to \mathcal{H}\mathrm{om}_{\mathcal{O}_{\mathfrak{X}}}(\mathcal{A}, \mathcal{A})$ is an isomorphism, because it is locally so. The composite

$$\mathcal{A} \longrightarrow \mathcal{H}\mathrm{om}_{\mathcal{O}_{\mathfrak{X}}}(\mathcal{A}, \mathcal{A}) \cong \mathcal{H}\mathrm{om}_{\mathcal{O}_{\mathfrak{X}}}(\mathcal{A}, \mathcal{O}_{\mathfrak{X}}) \otimes \mathcal{A} \longrightarrow \mathcal{O}_{\mathfrak{X}}$$

is called the **algebraic transfer map**.

Notation 8.17. If $f\colon \mathfrak{Y} \to \mathfrak{X}$ is a finite map of Deligne–Mumford stacks, then the map $\mathcal{O}_{\mathfrak{X}} \to f_*\mathcal{O}_{\mathfrak{Y}}$ is a map of $\mathcal{O}_{\mathfrak{X}}$-algebras with target locally free of finite rank. We write

$$f_!\colon f_*\mathcal{O}_{\mathfrak{Y}} \longrightarrow \mathcal{O}_{\mathfrak{X}}$$

for the resulting algebraic transfer map. More generally, if $\mathcal{F}$ is a quasi-coherent $\mathcal{O}_{\mathfrak{X}}$-module, then we also write $f_!$ for the resulting maps

$$f_*\mathcal{O}_{\mathfrak{Y}} \otimes_{\mathcal{O}_{\mathfrak{X}}} \mathcal{F} \longrightarrow \mathcal{F} \qquad \text{and} \qquad \mathrm{H}^s(\mathfrak{X}, f_*\mathcal{O}_{\mathfrak{Y}} \otimes_{\mathcal{O}_{\mathfrak{X}}} \mathcal{F}) \longrightarrow \mathrm{H}^s(\mathfrak{X}, \mathcal{F}).$$

The map $f_!$ is characterised by étale locally on $\mathfrak{X}$ being given by the d-fold sum map, where d is the rank of f; see [Stacks, Tag 03SH].

Our next goal is to lift these transfer maps to the level of $\mathcal{O}^{\mathrm{syn}}$, and to show that this map induces the algebraic transfer map on mod τ homotopy groups. The lift to $\mathcal{O}^{\mathrm{top}}$ is described by Davies in [Dav24a, Proposition 1.18], whose proof we closely imitate.

We will construct transfer maps for every finite log étale map of stacks. We work with log étale maps because we would like to accomodate level structures on Smf, but maps such as $\overline{\mathfrak{M}}_1(3) \to \overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}[\frac{1}{3}]$ are not finite étale, but only log étale. The main theorem of [HL16] states that there exists a log étale sheaf of $\mathbf{E}_\infty$-rings

$$\mathcal{O}^{\mathrm{top}} \colon (\mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}})^{\mathrm{op}} \longrightarrow \mathrm{CAlg}$$

whose global sections are Tmf. Imitating Definition 8.1, in what follows, we write $\mathcal{O}^{\mathrm{syn}}$ for the sheafification of the composite

$$(\mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}})^{\mathrm{op}} \xrightarrow{\mathcal{O}^{\mathrm{top}}} \mathrm{CAlg} \xrightarrow{\nu} \mathrm{CAlg}(\mathrm{Syn}_{\mathrm{MU}}).$$

Further, we write

$$\mathrm{Span}(\mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}}, \text{all}, \text{fin})$$

the span ∞-category where the backwards maps are all maps of stacks in the small log étale site for $\overline{\mathfrak{M}}_{\mathrm{ell}}$, and the forward maps are the finite log étale maps of stacks.

Theorem 8.18. *The log étale sheaf $\mathcal{O}^{\mathrm{syn}}$ post-composed with the forgetful functor to* $\mathrm{Syn}_{\mathrm{MU}}$ *admits a unique factorisation*

$$\mathbf{O}^{\mathrm{syn}} \colon \mathrm{Span}(\mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}}, \text{all}, \text{fin}) \longrightarrow \mathrm{Syn}_{\mathrm{MU}}.$$

Moreover, if $f \colon \mathfrak{Y} \to \mathfrak{X}$ is a finite log étale morphism in this site, then the map

$$f_! := \mathbf{O}^{\mathrm{syn}}(f) \colon \mathcal{O}^{\mathrm{syn}}(\mathfrak{Y}) \longrightarrow \mathcal{O}^{\mathrm{syn}}(\mathfrak{X})$$

induces the algebraic transfer map of f from Definition 8.16 *on $\pi_{n,s}(C\tau \otimes -)$:*

$$f_! \colon \mathrm{H}^s(\mathfrak{X}, f_*\omega_{\mathfrak{Y}}^{\otimes(n+s)/2}) \longrightarrow \mathrm{H}^s(\mathfrak{X}, \omega_{\mathfrak{X}}^{\otimes(n+s)/2}).$$

As with the proof of [Dav24a, Theorem 1.14], we obtain the above result by directly applying [BH21, Corollary C.13].

Proof. Consider the composite

$$(\mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}})^{\mathrm{op}} \xrightarrow{\mathcal{O}^{\mathrm{syn}}} \mathrm{CAlg}(\mathrm{Syn}_{\mathrm{MU}}) \longrightarrow \mathrm{Syn}_{\mathrm{MU}}$$

equipped with the wide subcategory of the domain on the finite log étale morphisms. In the notation of [BH21, Corollary C.13], we take $\mathcal{D}$ to be $\mathrm{Syn}_{\mathrm{MU}}$ with the cartesian monoidal structure. Since $\mathrm{Syn}_{\mathrm{MU}}$ is stable, the ∞-category $\mathrm{CAlg}(\mathcal{D})$ (a.k.a. $\mathbf{E}_\infty$-monoids in $\mathrm{Syn}_{\mathrm{MU}}$) is equivalent to $\mathrm{Syn}_{\mathrm{MU}}$. We claim this data satisfies the conditions of [BH21, Corollary C.13]. Indeed, the category $\mathcal{C} = \mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}}$ is extensive, meaning it admits finite coproducts, coproducts are disjoint, and finite coproducts decompositions are preserved under base-change, as this is true for any slice category of stacks over a base stack which is closed under coproducts and summands. The other hypotheses hold because finite log étale morphisms are log étale locally finite disjoint unions. By loc. cit., we obtain a unique lifting of the above composite to the desired functor

$$\mathbf{O}^{\mathrm{syn}} \colon \mathrm{Span}(\mathrm{DM}^{\text{log-ét}}_{/\overline{\mathfrak{M}}_{\mathrm{ell}}}, \mathrm{all}, \mathrm{fin}) \longrightarrow \mathrm{Syn}_{\mathrm{MU}}.$$

Next, we turn to the identification of the effect of $f_!$ on mod τ homotopy groups. By Corollary 7.33, we have an isomorphism

$$\pi_{n,s}(\mathcal{O}^{\mathrm{syn}}_{\mathfrak{X}}(\mathfrak{X})/\tau) \cong \mathrm{H}^s(\mathfrak{X}, \omega_{\mathfrak{X}}^{\otimes(n+s)/2}).$$

From the definition of the transfer from Notation 8.17, it is clear that it suffices to show the claim for $n+s=0$, i.e., to show that the effect on

$$\mathrm{H}^s(\mathfrak{X}, f_*\mathcal{O}_{\mathfrak{Y}}) \longrightarrow \mathrm{H}^s(\mathfrak{X})$$

is given by the algebraic transfer map.

First, we deal with the case where $\mathfrak{Y}$ and $\mathfrak{X}$ are affine, so that f is of the form $\operatorname{Spec} B \to \operatorname{Spec} A$. Choose an étale cover $q\colon \operatorname{Spec} C \to \operatorname{Spec} A$ such that the pullback along f is a disjoint union:

$$\begin{tikzcd}
\operatorname{Spec}(C)^{\sqcup d} \arrow[r] \arrow[d] & \operatorname{Spec} C \arrow[d, "q"] \\
\operatorname{Spec} B \arrow[r, "f"] & \operatorname{Spec} A.
\end{tikzcd}$$

On the map $\operatorname{Spec}(C)^{\sqcup d} \to \operatorname{Spec} C$, the algebraic transfer is given by multiplication by d, by the aforementioned [Stacks, Tag 03SH], and it is easy to see that the transfer from $\mathcal{O}^{\mathrm{syn}}$ is the same. The map $\operatorname{Spec} C \to \operatorname{Spec} A$ is faithfully flat, so that

$$\mathrm{H}^s(\operatorname{Spec} C, q^*\mathcal{O}_{\operatorname{Spec} A}) \cong \mathrm{H}^s(\operatorname{Spec} A),$$

and likewise for $\operatorname{Spec}(C)^{\sqcup d}$ covering $\operatorname{Spec} B$. Since the algebraic transfer and the transfer from $\mathcal{O}^{\mathrm{syn}}$ have the same effect on $\operatorname{Spec}(C)^{\sqcup d} \to \operatorname{Spec} C$, it follows that they agree on $\operatorname{Spec} B \to \operatorname{Spec} A$ as well.

In the general case, let $\mathfrak{Y} \to \mathfrak{X}$ be a finite log-étale map of Deligne–Mumford stacks over $\overline{\mathfrak{M}}_{\mathrm{ell}}$. Choose an étale cover $\operatorname{Spec} A \to \mathfrak{X}$; pulling this back along f, we obtain an étale cover $\operatorname{Spec} B \to \mathfrak{Y}$. Both the algebraic and the synthetic transfer maps are natural, so that they both assemble to a maps between the Čech nerve of this cover:

$$\begin{tikzcd} \operatorname{Spec} B^\bullet \arrow[r] \arrow[d] & \operatorname{Spec} A^\bullet \arrow[d] \\ \mathfrak{Y} \arrow[r, "f"] & \mathfrak{X}. \end{tikzcd}$$

Between these covers, the two transfer maps agree. It follows that they agree in the general case as well. ∎

8.4.1 A computation of a transfer

Using the finite log étale map $p\colon \overline{\mathfrak{M}}_1(3) \times \operatorname{Spec} \mathbf{Z}_{(2)} \to \overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}_{(2)}$ and using Section 7.5, we obtain a synthetic transfer map $p_!\colon \mathrm{Smf}_1(3)_{(2)} \to \mathrm{Smf}_{(2)}$, which in turn induces the algebraic transfer map

$$p_!\colon \mathrm{H}^s(\overline{\mathfrak{M}}_1(3) \times \operatorname{Spec} \mathbf{Z}_{(2)}, \omega^*) \longrightarrow \mathrm{H}^s(\overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}_{(2)}, \omega^*).$$

To compute this map in a specific range, we will compare this to a transfer map defined using Hopf algebroids. We use the simplification of the cubic Hopf algebroid at the prime 2 described by Bauer.

Definition 8.19 ([Bau08], Section 7). Let (A', Γ') be the **2-primary cubic Hopf algebroid**, defined by the data

$$A' = \mathbf{Z}_{(2)}[a_1, a_3], \qquad \Gamma' = A'[s,t]/\sim,$$

where the latter denotes the quotient by the relations

$$s^4 - 6st + a_1 s^3 - 3a_1 t - 3a_3 s = 0,$$

$$s^6 - 27t^2 + 3a_1 s^5 - 9a_1 s^2 t + 3a_1^2 s^4 - 9a_1^2 st + a_1^3 s^3 - 27a_3 t = 0,$$

and with left unit, right unit, and comultiplication defined such that the evident quotient

$$(A_{(2)}, \Gamma_{(2)}) \longrightarrow (A', \Gamma')$$

is a map of Hopf algebroids.

By [Bau08, Section 7], the above quotient map of Hopf algebroids induces an equivalence of stacks, and hence induces an isomorphism on cohomology groups. The left A'-module Γ' corresponds to the pushforward of the structure sheaf of $\overline{\mathfrak{M}}_1(3) \times \operatorname{Spec} \mathbf{Z}_{(2)}$ to $\overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}_{(2)}$ along the map p above; see [Mat16, Section 3]. We compute the algebraic transfer along f by computing the trace map $\operatorname{Tr}\colon \Gamma' \to A'$.

In more detail, first note that the left A'-module Γ' is free of rank 8 with basis

$$\{1, s, s^2, s^3, t, st, s^2t, s^3t\}.$$

For $x \in \Gamma'$, multiplication by x on Γ' is an A'-linear endomorphism. The associated trace map $\mathrm{Tr}\colon \Gamma' \to A'$ agrees with the algebraic transfer associated to $p_!$, as follows directly from the definition of the algebraic transfer from Definition 8.16. One can compute the traces of multiplication by s and t with respect to the above basis, resulting in the following.

Proposition 8.20. *The trace map* $\mathrm{Tr}\colon \Gamma' \to A'$ *is the map of left* A'*-modules determined by the following formulas:*

$$\begin{aligned} \mathrm{Tr}(1) &= 8 & \mathrm{Tr}(t) &= \tfrac{1}{3}a_1^3 - 4a_3 \\ \mathrm{Tr}(s) &= -4a_1 & \mathrm{Tr}(st) &= -2a_1a_3 \\ \mathrm{Tr}(s^2) &= 2a_1^2 & \mathrm{Tr}(s^2t) &= -\tfrac{1}{3}a_1^5 + 9a_1^2a_3 \\ \mathrm{Tr}(s^3) &= -a_1^3 & \mathrm{Tr}(s^3t) &= \tfrac{1}{3}a_1^6 - 7a_1^3a_3 - 27a_3^2. \end{aligned}$$

At the prime 2, the formulas from [Sil, Section III.1] for c_4, c_6 and Δ simplify to

$$c_4 = a_1^4 - 24a_1a_3, \qquad c_6 = -a_1^6 + 36a_1^3a_3 - 216a_3^2, \qquad \Delta = a_3^3(a_1^3 - 27a_3), \tag{8.21}$$

as do the formulas for η_R from [Bau08, Section 3]:

$$\eta_R(a_1) = a_1 + 2s \qquad \text{and} \qquad \eta_R(a_3) = a_3 + \tfrac{1}{3}a_1s^2 + \tfrac{1}{3}a_1^2s + 2t. \tag{8.22}$$

This immediately implies the following corollary.

Corollary 8.23. *The ideal*

$$(8, 2c_4, 2c_6) \subseteq \mathrm{H}^0(\overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}_{(2)}, \omega^{\otimes *})$$

consists of permanent cycles in the DSS for $\mathrm{Tmf}_{(2)}$.

Proof. The map of synthetic spectra $p_!\colon \mathrm{Smf}_1(3)_{(2)} \to \mathrm{Smf}_{(2)}$ induces a natural map of DSSs, which by Corollary 8.15 can be computed in the range $5s \leqslant n + 12$ using Proposition 8.20. As $\overline{\mathfrak{M}}_1(3)$ has cohomological dimension 1, the DSS for $\mathrm{Tmf}_1(3)_{(2)}$ collapses on the E_2-page, so all classes are permanent cycles. In particular, any class in the image of the algebraic transfer map is a permanent cycle. The formulas from Proposition 8.20 together with (8.21) and (8.22) then show that the elements 8, $2c_4$, and c_6 lie in the image of the transfer map:

$$\mathrm{Tr}(1) = 8, \qquad \mathrm{Tr}(3\eta_R(a_1a_3)) = 2c_4, \qquad \mathrm{Tr}(-\eta_R(a_3^2)) = 2c_6. \qquad \blacksquare$$

A similar analysis at the prime 3 for the DSS for TMF can be found in [Mei23, Section 5.6]. We only need the 2-local argument for our computations.

Remark 8.24. The precomposition of the transfer $\mathrm{Tr}\colon \Gamma' \to A'$ with the right unit $\eta_R\colon A' \to \Gamma'$ is a kind of $\Gamma_1(3)$-Hecke operator. In other words, the computations above are a computation of the effect of the spectral Hecke operator $\mathrm{T}_{\Gamma_1(3)}$ on $\mathrm{TMF}_{(2)}$ defined in [Dav24a, Definition 2.6] on the second page of its DSS on the submodule of holomorphic modular forms.

8.5 Classes from the sphere and detection

By construction, the synthetic $\mathbf{E}_\infty$-ring Smf comes with a map $\mathbf{S} \to \mathrm{Smf}$, which in particular gives us a map from the ANSS for $\mathbf{S}$ to the DSS for Tmf. In this section, we prove that various elements in the ANSS for $\mathbf{S}$ are detected in the DSS for Tmf via this map. We will freely use the computation of the homotopy groups of Smf/τ from the previous section.

We do this to define a number of elements in Smf and to show that these are nonzero. Specifically, we pin down certain elements in Smf by defining them as the image of pre-defined elements in the sphere. Particularly in the 2-local case, we require the sphere to define these elements because, in Smf, it is not a priori clear that they are permanent cycles, and because it is a priori not clear if these can be defined uniquely, due to the presence of elements in high filtration. Such problems do not arise in the sphere.

Definition 8.25. In 2-local Smf, we use the following notation.

- We write $\eta \in \pi_{1,1}\,\mathrm{Smf}_{(2)}$ for the image of $\eta \in \pi_{1,1}\,\mathbf{S}_{(2)}$.
- We write $\nu \in \pi_{3,1}\,\mathrm{Smf}_{(2)}$ for the image of $\nu \in \pi_{3,1}\,\mathbf{S}_{(2)}$.
- We write $\varepsilon \in \pi_{8,2}\,\mathrm{Smf}_{(2)}$ for the image of $\varepsilon \in \pi_{8,2}\,\mathbf{S}_{(2)}$.
- We write $\kappa \in \pi_{14,2}\,\mathrm{Smf}_{(2)}$ for the image of $\kappa \in \pi_{14,2}\,\mathbf{S}_{(2)}$.

In 3-local Smf, we use the following notation.

- We write $\alpha \in \pi_{3,1}\,\mathrm{Smf}_{(3)}$ for the image of $\alpha \in \pi_{3,1}\,\mathbf{S}_{(3)}$.
- We write $\beta \in \pi_{10,2}\,\mathrm{Smf}_{(3)}$ for the image of $\beta \in \pi_{10,2}\,\mathbf{S}_{(3)}$.

The case of $\bar{\kappa}$ is more complicated, both because $\bar{\kappa}$ is not uniquely defined in the sphere, and because there is a filtration jump when going from the sphere to Smf. For this reason, we postpone a definition of it to Section 8.5.2; see Notation 8.34 in particular.

In all of the above definitions, it is not a priori clear if the resulting classes in Smf are nonzero. Our first goal is to show that this is in fact the case, by showing that mod τ these elements are mapped to nonzero classes.

8.5.1 Height one detection results

To show that Smf detects v_1-periodic elements, we show that an appropriate version of Smf with level structure detects v_1.

Lemma 8.26.

(1) *The map* $\mathbf{S}/\tau \to \mathrm{Smf}_1(3)/\tau$ *detects the classes*

$$\widetilde{v}_1 \in \pi_{2,0}\,\mathbf{S}/(\tau,2) \cong \mathbf{F}_2 \quad \textit{and} \quad \widetilde{v}_1^2 \in \pi_{4,0}\,\mathbf{S}/(\tau,4) \cong \mathbf{F}_2,$$

where $\widetilde{v}_1$ *and* $\widetilde{v}_1^2$ *are generators.*

(2) *The natural map of synthetic* $\mathbf{E}_\infty$*-rings* $\mathbf{S}/\tau \to \mathrm{Smf}_1(2)/\tau$ *detects the nonzero class* $\widetilde{v}_1 \in \pi_{4,0}\,\mathbf{S}/(\tau,3)$.

Proof. The DSS for $\mathrm{Tmf}_1(3)$ collapses on the E_2-page, as the stack $\overline{\mathfrak{M}}_1(3)$ is a weighted projective stack (see [Mei22, Example 2.1]) so their DSSs are concentrated in filtrations 0 and 1. This yields the calculation $\pi_*\,\mathrm{Tmf}_1(3) \cong A \oplus \widehat{A}$, where

$$A = \mathbf{Z}[\tfrac{1}{2}][a_1,a_3], \qquad \text{where } |a_i| = 2i,$$

and where $\widehat{A} = A/(a_1^\infty, a_3^\infty)$ is the torsion A-module with $\mathbf{Z}[\frac{1}{2}]$-basis given by

$$\frac{1}{a_1^i a_3^j} \quad \text{with } i,j>0, \qquad \text{where } \left|\frac{1}{a_1^i a_3^j}\right| = -2i-6j-1;$$

this is essentially [MR09, Corollary 3.3]. In positive degrees, these elements admit a description in terms of modular forms of level $\Gamma_1(3)$. By [HL16, Theorem 6.2], evaluation at the cusp gives a map of $\mathbf{E}_\infty$-rings

$$\mathrm{Tmf}_1(3) \longrightarrow \mathrm{KU}[\tfrac{1}{3}]$$

that sends $a_1 \in \pi_2\,\mathrm{Tmf}_1(3)$ to u, where u denotes the Bott periodicity generator of $\pi_2\,\mathrm{KU}[\frac{1}{3}]$. Let a_1 denote the nonzero class in $\pi_{2,0}\,\mathrm{Smf}_1(3)/(\tau,2) \cong \mathbf{F}_2$. Because the mod 2 reduction of u detects $\widetilde{v}_1$ on Adams–Novikov E_2-pages, see [CD24, Theorem 5.4] for example, the reduction of the synthetic lift of a_1, must also detect $\widetilde{v}_1$. Similar arguments show that $\widetilde{v}_1^2$ is detected by $\mathrm{Smf}_1(2)/(\tau,4)$. The arguments at the prime 3 are the same. ■

Using this, we show that Smf detects the corresponding v_1-periodic elements in the divided α-family.

Proposition 8.27. *The unit map of synthetic* $\mathbf{E}_\infty$*-rings* $\mathbf{S}/\tau \to \mathrm{Smf}/\tau$ *detects the following classes:*

$$\alpha \in \pi_{3,1}\,\mathbf{S}_{(3)}/\tau, \qquad h_1 \in \pi_{1,1}\,\mathbf{S}_{(2)}/\tau, \qquad h_2 \in \pi_{3,1}\,\mathbf{S}_{(2)}/\tau, \qquad c \in \pi_{8,2}\,\mathbf{S}_{(2)}/\tau.$$

We assume the reader is familiar with this Adams–Novikov notation, as well as the ANSS for $\mathbf{S}$ in low degrees; see, e.g., [Rav78, Table 2].

Proof. The class η can be defined as $\partial(\widetilde{v}_1)$, where $\partial\colon \mathbf{S}/2 \to \Sigma\mathbf{S}_{(2)}$ is the boundary map associated to the mod 2 Moore spectrum, and $\widetilde{v}_1 \in \pi_2\,\mathbf{S}/2$ is the class inducing multiplication by v_1 on K(1)-homology. The class $\widetilde{v}_1$ is detected on the ANSS for $\mathbf{S}$ in filtration 0, i.e., in the bigraded homotopy group $\pi_{2,0}\,\mathbf{S}/(\tau,2)$.

We now consider the following commutative diagram of abelian groups

$$\begin{tikzcd}
 & \pi_{2,0}\,\mathbf{S}/(\tau,2) \arrow[r, "\partial"] \arrow[d, "h"] & \pi_{1,1}\,\mathbf{S}_{(2)}/\tau \arrow[d] & \\
\pi_{2,0}\,\mathrm{Smf}_{(2)}/\tau \arrow[r, "q"] & \pi_{2,0}\,\mathrm{Smf}/(\tau,2) \arrow[r, "\partial"] \arrow[d] & \pi_{1,1}\,\mathrm{Smf}_{(2)}/\tau \arrow[r, "2"] & \pi_{1,1}\,\mathrm{Smf}_{(2)}/\tau \\
 & \pi_{2,0}\,\mathrm{Smf}_1(3)/(\tau,2) & &
\end{tikzcd}$$

where the vertical maps are induced by the maps of synthetic $\mathbf{E}_\infty$-rings $\mathbf{S} \to \mathrm{Smf} \to \mathrm{Smf}_1(3)$. First, notice that $h(v_1) \neq 0$ as the image of v_1 is nonzero in $\pi_{2,0}\,\mathrm{Smf}_1(2)/(\tau,3)$ by Lemma 8.26.

Next, by Figure A.2 and the exactness of the rows above, we see that $\pi_{2,0}\,\mathrm{Smf}_{(2)}/\tau$ is zero and the groups $\pi_{*,s}\,\mathrm{Smf}_{(2)}/\tau$ are 2-torsion for $s > 0$. In particular, the multiplication by 2 map above is zero, and we see that ∂ in the middle row is an isomorphism. From the commutativity of the above square, this implies that the nonzero class in $\pi_{1,1}\,\mathrm{Smf}_{(2)}/\tau$ detects $\eta = \partial(\widetilde{v}_1)$.

For ν, we use that $\nu = \partial(\widetilde{v}_1^2)$ coming from $\mathbf{S}/4$, and in the η-case above. The α-case follows similarly, *mutatis mutandis*.

For the claim about c, note that in $\mathbf{S}_{(2)}/\tau$, we have the relation $h_2^3 = h_1 \cdot c$. As both h_2^3 and h_1 are nonzero in $\mathrm{Smf}_{(2)}/\tau$, the conclusion follows. ∎

In particular, the elements α, η, ν, and ε in Smf are nonzero.

8.5.2 The classes κ and $\bar{\kappa}$

The height 2 classes κ and $\bar{\kappa}$ are more subtle than the height 1 elements encountered above. We begin with κ.

Proposition 8.28. *The natural map of synthetic $\mathbf{E}_\infty$-rings $\mathbf{S}_{(2)}/\tau \to \mathrm{Smf}_{(2)}/\tau$ detects the mod τ reduction of the class $\kappa \in \pi_{14,2}\,\mathbf{S}_{(2)}$. In particular, the image of the class $\kappa \in \pi_{14,2}\,\mathbf{S}_{(2)}$ in $\mathrm{Smf}_{(2)}$ is also nonzero.*

Proof. In [Isa19, Table 23], Isaksen shows that the nonzero class $\eta\kappa$ lies in $\langle 2\nu, \nu, \varepsilon\rangle$ in $\pi_{14}\,\mathbf{S}$. As $\pi_{14,2}\,\mathbf{S}$ is τ-torsion free, we can lift this containment to $\eta\kappa \in \langle 2\nu, \nu, \varepsilon\rangle$ in the synthetic sphere. Using Proposition B.14 and Corollary 8.15, we see that modulo τ we have a nonzero class $h_1 d \in \langle 2h_2, h_2, c\rangle$ in the cohomology of the 2-primary cubic Hopf algebroid; see [Bau08, Appendix A].

As $\pi_{12,2}\,\mathrm{Smf}_{(2)}/\tau$ and $\pi_{7,1}\,\mathrm{Smf}_{(2)}/\tau$ both vanish (see Figure A.2) the above Massey product has zero indeterminacy and hence cannot contain zero. In particular, we see that $\eta\kappa$ cannot vanish in $\pi_{15,3}\,\mathrm{Smf}_{(2)}$. As η is detected in $\pi_{1,1}\,\mathrm{Smf}_{(2)}$, this implies that $\kappa \in \pi_{14,2}\,\mathrm{Smf}_{(2)}$ is also nonzero and reduces to the generator d in $\pi_{14,2}\,\mathrm{Smf}_{(2)}/\tau \cong \mathbf{F}_2$. ∎

Next, we turn to $\bar\kappa$. We begin by pinning down what exactly we mean by $\bar\kappa$, because in the sphere spectrum, this is only well defined up to a factor of $\nu^2\kappa$. We then pin down what we mean by $\bar\kappa$ as an element in the synthetic sphere, after which we then specify how this gives rise to an element in Smf. As we need various facts about $\bar\kappa$ relating to Toda brackets, we focus on this in the sphere as well.

Definition 8.29. We write $\bar\kappa$ for any choice of element in the Toda bracket $\langle \kappa, 2, \eta, \nu\rangle$ in the (non-synthetic) $\pi_{20}\,\mathbf{S}_{(2)}$. Similarly, we write $\bar\kappa$ for any choice of element inside the synthetic Toda bracket $\langle \kappa, 2, \eta, \nu\rangle$ in $\pi_{20,2}\,\mathbf{S}_{(2)}$.

A priori, it is not clear that either of these brackets are nonempty or nonzero. This follows from some classical facts.

Lemma 8.30. *The indeterminacy of the Toda bracket $\langle \kappa, 2, \eta, \nu\rangle$ in the non-synthetic sphere spectrum $\pi_{20}\,\mathbf{S}_{(2)}$ is equal to the subgroup generated by $\nu^2\kappa$. Moreover, this bracket does not contain zero.*

Proof. Using the formula for the indeterminacy of a four-fold Toda bracket as presented by Kochman [Koc, Theorem 2.3.1 (b)], one can compute the desired indeterminacy. We verified this using the Massey product calculator of [`ext-rs`].

For the final statement, we use the computation of Kochman [Koc, Lemma 5.3.8 (e)], which states that $\langle \kappa, 2, \eta, \nu\rangle$ contains a generator of $\pi_{20}\,\mathbf{S}_{(2)} \cong \mathbf{Z}/8$. This, combined with the fact that $\nu^2\kappa$ is the nonzero 2-torsion element in this group [Koc, Theorem 5.3.1 (a)], shows that this bracket does not contain zero. ∎

The naturality of Toda brackets then implies that the synthetic Toda bracket $\langle \kappa, 2, \eta, \nu\rangle$ also cannot contain zero. Moreover, the fact that $\pi_{20,2}\,\mathbf{S}_{(2)}$ is τ-power torsion-free allows us to lift relations on any choice of $\bar\kappa$ as well.

Lemma 8.31. *For any choice of $\bar\kappa \in \langle \kappa, 2, \eta, \nu\rangle \subseteq \pi_{20,2}\,\mathbf{S}_{(2)}$, we have $4\bar\kappa = \tau^2\nu^2\kappa$.*

Proof. The proof of Lemma 8.30 shows that the equation $4\bar\kappa = \nu^2\kappa$ holds in $\pi_{20}\,\mathbf{S}_{(2)}$. As the τ-inversion map is injective in degree $(20, 2)$, we obtain the desired equation in $\pi_{20,2}\,\mathbf{S}_{(2)}$. ∎

Lemma 8.32. *For any choice of $\bar{\kappa} \in \langle \kappa, 2, \eta, \nu \rangle \subseteq \pi_{20,2}\,\mathbf{S}_{(2)}$, the class $\nu^3\bar{\kappa} \in \pi_{29,5}\,\mathbf{S}_{(2)}$ is τ-power torsion.*

Proof. This follows from the classical fact that $\pi_{29}\,\mathbf{S}_{(2)} = 0$; see [Rav04, Figure A3.2]. ■

The ill-definedness of $\bar{\kappa}$ persists in Smf, where there is yet another issue. In the sphere, the class $\bar{\kappa}$ has Adams–Novikov filtration 2, while in Smf, it turns out to have filtration 4. We now make precise what we mean by this, which will allow us to pin down what what we mean by $\bar{\kappa}$ as an element of $\pi_{*,*}\,\mathrm{Smf}$ going forward.

Proposition 8.33. *For every choice of $\bar{\kappa} \in \langle \kappa, 2, \eta, \nu \rangle \subseteq \pi_{20,2}\,\mathbf{S}_{(2)}$, there is a unique element y in $\pi_{20,4}\,\mathrm{Smf}_{(2)}$ such that $\tau^2 \cdot y \in \pi_{20,2}\,\mathrm{Smf}_{(2)}$ is the image of $\bar{\kappa} \in \pi_{20,2}\,\mathbf{S}_{(2)}$ under $\mathbf{S}_{(2)} \to \mathrm{Smf}_{(2)}$. Moreover, the reduction of y to $\mathrm{Smf}_{(2)}/\tau$ is a generator for the group $\pi_{20,4}\,\mathrm{Smf}_{(2)}/\tau \cong \mathbf{Z}/8$.*

Proof. We will implicitly localise at the prime 2 in this proof. Write $x \in \pi_{20,2}\,\mathrm{Smf}$ for the image of $\bar{\kappa} \in \pi_{20,2}\,\mathbf{S}$ under the unit map $\mathbf{S} \to \mathrm{Smf}$. We claim that x is uniquely τ^2-divisible. Indeed, we look at the long exact sequence induced by the cofibre sequence

$$\Sigma^{-1,1}\,\mathrm{Smf}/\tau^2 \longrightarrow \Sigma^{0,-2}\,\mathrm{Smf} \xrightarrow{\tau^2} \mathrm{Smf} \longrightarrow \mathrm{Smf}/\tau^2.$$

Part of this long exact sequence reads

$$\pi_{21,1}\,\mathrm{Smf}/\tau^2 \longrightarrow \pi_{20,4}\,\mathrm{Smf} \xrightarrow{\tau^2} \pi_{20,2}\,\mathrm{Smf} \longrightarrow \pi_{20,2}\,\mathrm{Smf}/\tau^2.$$

The homotopy of Smf/τ vanishes in bidegrees $(21,1)$, $(21,2)$, $(20,2)$, and $(20,3)$; see Figure A.2. Corollary 3.72 then implies that the homotopy of Smf/τ^2 vanishes in $(21,1)$ and $(20,2)$, so that multiplication by τ^2 induces an isomorphism

$$\tau^2\colon\ \pi_{20,4}\,\mathrm{Smf} \xrightarrow{\cong} \pi_{20,2}\,\mathrm{Smf}.$$

Write y for the τ^2-division of x. From Lemma 8.31, we learn that

$$4x = \tau^2 \cdot \kappa\nu^2.$$

Because multiplication by τ^2 in this bidegree in Smf is injective, we deduce that

$$4y = \kappa\nu^2.$$

From the multiplicative structure of Smf/τ, we see that $\kappa\nu^2$ reduces to 4 times a generator of $\pi_{20,4}\,\mathrm{Smf}/\tau \cong \mathbf{Z}/8$. Write g for a choice of a generator, and write $\bar{y}$ for the reduction mod τ of y. The relation above then tells us that $4g = 4\bar{y}$, i.e., $4(g - \bar{y}) = 0$. As such, $g - \bar{y}$ is an even multiple of g, i.e., $\bar{y}$ is an odd multiple of g, which says that it is a unit in $\mathbf{Z}/8$ away from g. This means that $\bar{y}$ is also a generator, as claimed. ■

This thesis is focussed on the case of Smf rather than the sphere, so from now on, we will use the notation $\bar{\kappa}$ to mean the element in filtration 4, as follows.

Notation 8.34.

- From now on, we fix a choice of preferred element in $\langle \kappa, 2, \eta, \nu \rangle \subseteq \pi_{20,2}\, \mathbf{S}_{(2)}$.
- We write $\bar{\kappa} \in \pi_{20,4}\, \mathrm{Smf}_{(2)}$ for the element uniquely determined by this choice using Proposition 8.33. As a result, the chosen element in $\pi_{20,2}\, \mathbf{S}_{(2)}$ maps to $\tau^2\bar{\kappa}$ in $\pi_{20,2}\, \mathrm{Smf}_{(2)}$.
- We write $g \in \pi_{20,4}\, \mathrm{Smf}_{(2)}/\tau$ for the reduction of $\bar{\kappa}$ to Smf/τ. By Proposition 8.33, this is a generator of $\pi_{20,4}\, \mathrm{Smf}_{(2)}/\tau \cong \mathbf{Z}/8$.

From this definition of $\bar{\kappa}$ in $\mathrm{Smf}_{(2)}$, we immediately obtain the following facts.

Corollary 8.35.

- (20,2) $\tau^2\bar{\kappa} \in \langle \kappa, 2, \eta, \nu \rangle$ *in* $\mathrm{Smf}_{(2)}$.
- (20,4) $4\bar{\kappa} = \nu^2\kappa$ *in* $\mathrm{Smf}_{(2)}$.
- (29,7) $\nu^3\bar{\kappa}$ *is* τ*-power torsion in* $\mathrm{Smf}_{(2)}$.

All of these facts will be crucial in the computations of the next chapter.

8.5.3 Toda bracket relations from the sphere

There are a handful of standard Toda bracket computations from the non-synthetic sphere which we would like to lift to Smf.

Proposition 8.36.

- (8,2) $\varepsilon \in \langle \nu, 2\nu, \eta \rangle$ *in* $\mathrm{Smf}_{(2)}$.
- (15,3) $\eta\kappa \in \langle \nu, 2\nu, \varepsilon \rangle$ *and* $\eta\kappa \in \langle 2\nu, \nu, \varepsilon \rangle$ *in* $\mathrm{Smf}_{(2)}$.
- (21,3) $\tau^2\eta\bar{\kappa} \in \langle \nu, 2\nu, \kappa \rangle$ *in* $\mathrm{Smf}_{(2)}$.

All of these relations will follow from the analogous relations in the sphere. We only state them in Smf to be consistent with our use of $\bar{\kappa}$ from Notation 8.34.

Proof. It suffices to prove these statements in the synthetic sphere. By [IWX23, Table 10], all of these statements hold in the non-synthetic sphere, except for those for $\eta\kappa$, which follow from [Isa19, Table 23].[1] ■

[1]Although [Isa19; IWX20] as a whole rely on the homotopy groups of tmf, this is only for computations in stems higher than those considered here. One can independently check these computations using the Massey product calculator of [`ext-rs`], for example.

Chapter 9

The homotopy groups of Smf

In this chapter, we compute the descent spectral sequence for Tmf locally at the primes 2 and 3, as well as after inverting 6. By far the most difficult part is the 2-local computation, which takes up almost the entirety of this chapter; the remainder is treated in Section 9.10.

Since our overall goal is the Gap Theorem, we focus on stems above -21. The exact same methods also yield the homotopy groups of Tmf in stems in and below dimension -21, but as they have no bearing on the Gap Theorem, we do not include these arguments here. We refer to [CDvN24, Section 6.9] for the computation of these remaining stems, which justifies certain arguments made by Konter [Kon12] on the matter.

In Appendix A, we provide spectral sequence charts that aid in reading the computation in this section. We also include several tables with information from this section, including relevant information from the sphere, important lifts of E_2-elements to Smf/τ^k for various k, relations and extensions in $\pi_{*,*}\mathrm{Smf}/\tau^k$ for various k, important values of the (truncated) total differential, and key Toda brackets. Another diagram of the same spectral sequence is given in [Kon12, Figures 26–27] with only minor typographical changes and omissions. We recommend that the reader keeps these charts and tables nearby throughout the arguments in this chapter.

Let us begin by outlining our approach. The E_2-page was computed in Section 8.3; see Figure A.2 or [Kon12, Figure 25] for a visual representation. As explained there, part of this diagram (the *connective region* of Definition 9.3 below) is computed in [Bau08, Section 7], and we can import multiplicative results from loc. cit. through Proposition B.14 and Corollary 8.15.

We will only explicitly state those differentials that propagate to all others by the Leibniz rule; let us introduce some terminology for this.

Definition 9.1. In a multiplicative spectral sequence, an **atomic differential** is a differential of the form $d_r(x) = y$, where x is an indecomposable element of the dga E_r and where $y \neq 0$.

It follows from the Leibniz rule that all of our atomic differentials yield all of the differentials in the *connective region* (see Definition 9.3 below); this is the region of Figures A.3 to A.6 below the blue line. The *meta-arguments* of Section 9.1 ensure both that all of the differentials above the blue line and to the right of the orange line follow from our atomic differentials, and that there are no differentials that cross this blue line.

We will work page by page through the DSS. This amounts to ticking off the following checklist at every page.

- Compute the atomic differentials.
- Check the conditions of the meta-arguments to deduce all other differentials whose sources lie in stems $n \geqslant -20$.
- Calculate as many lifts of elements, total differentials, extensions, and synthetic Toda brackets as will be necessary for future pages.

Finally, in order to start the computation, we need a few low-dimensional classes from the sphere. We import the following elements from the sphere, in the way explained in Section 8.5. For the convenience of the reader, let us recall these definitions.

Notation 9.2.

- We write $\eta \in \pi_{1,1}\,\mathrm{Smf}$ for the image of $\eta \in \pi_{1,1}\,\mathbf{S}$ under $\mathbf{S} \to \mathrm{Smf}$.
- We write $\nu \in \pi_{3,1}\,\mathrm{Smf}$ for the image of $\nu \in \pi_{3,1}\,\mathbf{S}$ under $\mathbf{S} \to \mathrm{Smf}$.
- We write $\varepsilon \in \pi_{8,2}\,\mathrm{Smf}$ for the image of $\varepsilon \in \pi_{8,2}\,\mathbf{S}$ under $\mathbf{S} \to \mathrm{Smf}$.
- We write $\kappa \in \pi_{14,2}\,\mathrm{Smf}$ for the image of $\kappa \in \pi_{14,2}\,\mathbf{S}$ under $\mathbf{S} \to \mathrm{Smf}$.
- We write $\bar{\kappa} \in \pi_{20,4}\,\mathrm{Smf}$ for the element defined in Notation 8.34.

Modulo τ, these elements reduce to h_1, h_2, c, d and g, respectively; see Proposition 8.27 and Section 8.5.2. In particular, these classes in the DSS for Tmf are permanent cycles.

9.1 Meta-arguments

There are two steps in our computation that amount to a repeated check on each page, so we condense these checks into a pair of meta-arguments here. For clarity of exposition, we separate these from the main page-by-page argument. However, we will need to make some forward references to the E_4-page. We will therefore be explicit in computing the E_4-page in our main argument without using the meta-arguments here in a circular way. This should not cause much confusion as the E_3 page is quite simple, with only one atomic differential.

In this section, we will frequently use terminology based on the way the relevant information appears in the spectral sequence chart, so we define these terms carefully first.

Definition 9.3. The **connective region** on the E_r-page of the descent spectral sequence for Tmf is the region in the plane consisting of those values of (n, s) on or below the line

$$s = \tfrac{1}{5}n + \tfrac{12}{5}.$$

The **nonconnective region** refers to the region above this line. A **line-crossing differential** is a differential $d_r(x) = y$ with the property that x is in the connective region and y is in the nonconnective region.

In the charts of Section A.2, we have drawn the line $s = \frac{1}{5}n + \frac{12}{5}$ in blue.

9.1.1 Line-crossing differentials

Our first meta-argument allows us to efficiently rule out the possibility of line-crossing differentials. There is a simple condition on atomic differentials we must check on each page — this guarantees that we only have to rule out line-crossing differentials through a finite range on a given page, which we check on each page by hand.

Proposition 9.4. *Let $r > 4$. Suppose we know that for all $r' < r$ and all atomic differentials $d_{r'}(x) = y$ in the connective region that the implication*

$$gy = 0 \text{ in } E_{r'} \implies gx = 0 \text{ in } E_{r'}$$

holds. Then every element in the connective region on E_r in filtration $\geqslant 4$ is divisible by g.

Proof. Every element in the connective region on E_4 in filtration $\geqslant 4$ is divisible by g, as follows from Proposition 9.12. Inductively, an element $z \in E_r$ can be written as $z = gx$ on $E_{r'}$. If the element z fails to be divisible by g on E_r, there must be a differential $d_{r'}(x) = y$ with $gy = 0$, but the assumptions preclude this possibility. ∎

Proposition 9.5. *Let $r > 4$. Suppose that any element in the connective region on E_r in filtration $\geqslant 4$ is divisible by g. Let n be the largest stem on E_r such that the nonconnective region has a nonzero class in bidegrees (n, s) for $r \leqslant s \leqslant r + 3$. If there are no line-crossing d_r-differentials through the $(n + 1)$-stem, then there are no line-crossing d_r-differentials.*

Proof. Any possible line-crossing differential whose source x is in filtration $\geqslant 4$ has $x = gx'$, hence if $d_r(x') = 0$, then $d_r(x) = 0$. We may therefore assume without loss of generality that x is in filtration $0 \leqslant s \leqslant 3$, and the result follows. ∎

9.1.2 Differentials in the S-region

Our second meta-argument allows us to conclude that all differentials in the S-region of Definition 9.7 are accounted for by application of the Leibniz rule. For this we will use the following.

Proposition 9.6. *If the* E_r*-page is* Δ^8*-torsion free in stems* $\geqslant n-1$*, and* Δ^8 *is a* d_r*-cycle, then all* d_r*-differentials in the nonconnective region with source in the n-stem are uniquely determined by the* d_r*-differentials in the connective region via the Leibniz rule.*

Proof. Let $x \in \mathrm{E}_r$ be an element in the n-stem in the nonconnective region. For a suitable power k, the product $\Delta^{8k}x$ lies in the connective region, and one may divide a nonzero differential $d_r(\Delta^{8k}x)$ by Δ^{8k} to obtain a nonzero differential in the nonconnective region. Conversely, also by Δ^8-torsion freeness, any nonzero differential with source in the n-stem in the nonconnective region determines a nonzero differential in the connective region. ■

The torsion-free condition of the previous proposition also follows from the checks in Propositions 9.4 and 9.5.

Definition 9.7. Let us denote by S the region consisting of stems $n > -21$, together with the (-21)-stem in filtration $s > 1$. In our spectral sequence charts of Section A.2, we indicate this region as the one to the right of the orange line.

Proposition 9.8. *Let* $r > 4$*. Suppose the conditions of Propositions* 9.4 *and* 9.5 *have been verified for* $d_{\leqslant r}$ *and suppose that* Δ^8 *is a* d_{r+1}*-cycle. Then* E_{r+1} *is* Δ^8*-torsion free in the region* S*.*

Remark 9.9. The region S has been chosen precisely to yield Theorem A. Much of the nonconnective region is actually Δ^8-torsion free, but this does not always hold. For example, 8 times the generator of $\pi_{-21,1}\,\mathrm{Smf}/\tau$ is Δ-torsion; see Figure A.2. We deal with these more subtle cases in [CDvN24, Section 6.9].

Proof of Proposition 9.8. We proceed by induction, noting that E_4 is Δ^8-torsion free in S, that the conditions of Propositions 9.4 and 9.5 hold for $d_{\leqslant 4}$, and finally, that by Proposition 9.14 the following pair of conditions holds for $r \leqslant 4$.

- A class $a \in \mathrm{E}_r^{n,s}$ satisfies
$$s < \tfrac{1}{5}(n-192) + \tfrac{12}{5}$$
if and only if $a = \Delta^8 b$, where b is in the connective region.
- A class $a \in \mathrm{E}_r^{n,s}$ in S satisfies $a = \Delta^8 b$ for b in the S-region if and only if $n \geqslant 171$.

We therefore assume by induction that E_r is Δ^8-torsion free, and that the two conditions above hold for all $r' \leqslant r$. Since E_r is Δ^8-torsion free, the E_{r+1}-page is Δ^8-torsion free in S unless there is a d_r-cycle x in S such that $d_r(y) = \Delta^8 x$ for some y, and y is not divisible by Δ^8. Given that the conditions of Propositions 9.4 and 9.5 hold for

$d_{\leqslant r}$, such a differential cannot be line-crossing, so it has both source and target in either the connective or nonconnective region. In the former case, it follows that $\Delta^8 y$ lives in a bidegree (n, s) satisfying

$$s < \tfrac{1}{5}(n - 192) + \tfrac{12}{5} \tag{9.10}$$

and hence so must x. This leads to a contradiction as then x is also divisible by Δ^8. In the latter case, it follows that $\Delta^8 y$ lives in a bidegree (n, s) satisfying $n \geqslant 192$, hence the same is true for x, showing again that x is Δ^8-divisible.

Therefore E_{r+1} is Δ^8-torsion free in S. To complete the induction, we need to establish the above two conditions for E_{r+1}. If b is in the connective region and $a = \Delta^8 b$, then it is clear from E_2 that (9.10) holds. Conversely, if $a \in E_{r+1}^{n,s}$ satisfies (9.10), then $a = \Delta^8 b$ on E_r by induction, and b is in the connective region for degree reasons. Since E_r is Δ^8-torsion free in S, and a is a d_r-cycle, so is b, so the division $a = \Delta^8 b$ carries to the E_{r+1}-page unless b is hit by a differential. But since Δ^8 is a d_r-cycle, it would also follow that a is hit, a contradiction. The condition in the nonconnective region is established by an analogous argument. ■

9.2 Page 3

9.2.1 Atomic differentials

Proposition 9.11 (3,1). $d_3(b) = h_1^4$.

Proof. By Proposition 8.27, the map

$$\mathbf{S} \longrightarrow \mathrm{Smf} \longrightarrow \mathrm{Smf}/\tau$$

sends η to the unique nonzero class $h_1 \in \pi_{1,1}\,\mathrm{Smf}/\tau$. There is a differential in the ANSS for $\mathbf{S}$ which hits η^4, see [Rav78, Table 2] for example, so there also has to be a differential in the signature spectral sequence of Smf hitting η^4. For degree reasons, the only things that can do this is a d_3 killing h_1^4, and $b \in \pi_{5,1}\,\mathrm{Smf}/\tau$ is the only potential source. ■

All other differentials follow from the Leibniz rule and the fact that η^4 is τ^2-torsion. For example, we have $d_3(c_6) = \eta^3 c_4$.

9.2.2 Meta-arguments

Computing homology with respect to the d_3-differentials, we have the following facts about the E_4-page. These serve as base cases for the induction arguments used in the meta-arguments of Section 9.1.

Proposition 9.12. *Every element $x \in E_4$ in the connective region of filtration $\geqslant 4$ is divisible by g.*

Proof. The computation of the connective region of the E_2-page is done by Bauer in [Bau08, Section 7]. A straightforward consequence is that every d_3-cycle x in the connective region of E_2 of filtration $\geqslant 4$ is divisible by either g or h_1^4. The claim then follows from the differential of Proposition 9.11. ■

Proposition 9.13. *There are no line-crossing d_3 differentials.*

Proof. For degree reasons, the only possible d_3's crossing the line have source of the form $g^k h_1^3$, which is a permanent cycle. ■

Proposition 9.14. *The class Δ^8 is a d_3-cycle, and E_4 is Δ^8-torsion free in the region S of Proposition 9.8. Moreover, the following properties hold.*

- *A class $a \in E_4^{n,s}$ satisfies*
$$s < \tfrac{1}{5}(n-192) + \tfrac{12}{5}$$
if and only if $a = \Delta^8 b$, where b is in the connective region.
- *A class $a \in E_4^{n,s}$ in the S-region satisfies $a = \Delta^8 b$ for b in the S-region if and only if $n \geqslant 171$.*

Proof. The class Δ^8 is a d_3-cycle for degree reasons. It follows from [Kon12, Section 5.1] that the region S is Δ^8-torsion free. The claim for E_4 then follows from the fact there is a single atomic d_3 that does not introduce Δ^8-torsion.

For the latter claims, it helps to consult a chart Figure A.2; a larger version of this chart appears in [Kon12, Figure 25]. The E_4-page is divided into g-periodic strips of width 24 and slope $1/5$. Multiplication by Δ maps one strip isomorphically to the next within the region S, which implies the above claims. ■

9.2.3 Hidden extensions

A crucial 2-extension is also generated on the E_3-page.

Lemma 9.15 (3,1). *We have an isomorphism*

$$\pi_{3,1}\,\mathrm{Smf}/\tau^{14} \cong \mathbf{Z}/8\langle\nu\rangle \qquad \textit{where} \quad \tau^2\eta^3 = 4\nu.$$

Proof. This follows from the sphere: we claim that

$$\pi_{3,1}\,\mathbf{S} \cong \mathbf{Z}/8\langle\nu\rangle.$$

Indeed, using the $\mathbf{F}_2$-Adams spectral sequence for the sphere, we learn that $4\nu = \eta^3$ in the non-synthetic homotopy group $\pi_3\,\mathbf{S}$. Because the Adams–Novikov spectral sequence for the sphere has no differentials hitting the 3-stem, we learn that $\pi_{3,*}\,\mathbf{S}$ (referring to the MU-synthetic sphere) is τ-torsion free. For degree reasons, one therefore has the relation $4\nu = \tau^2\eta^3$.

We defined the elements η and ν in Smf to be the images of these respective elements in $\nu\mathbf{S} \to \mathrm{Smf}$, so we learn that $4\nu = \tau^2\eta^3$ holds in Smf. It remains to show that ν generates $\pi_{3,1}\,\mathrm{Smf}/\tau^{14}$. This follows from the fact that $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau$ induces an isomorphism on $\pi_{3,1}$ using the d_3's already computed and Theorem 3.70. Since we know that ν reduces to h_2 by Proposition 8.27, this finishes the argument. ■

9.2.4 Lifts

On later pages, we will need to work with precisely defined lifts of elements from Smf/τ to higher Smf/τ^k. These will serve as the way to express hidden extensions and total differentials.

We begin by lifting Δ from Smf/τ to Smf/τ^4. It does not lift to Smf/τ^5, as Δ turns out to support a d_5. As such, the existence of this lift to Smf/τ^4 is a purely synthetic phenomenon. This lift will be absolutely crucial to all of our computations going forward; see Proposition 9.29, for example.

Lemma 9.16 (24,0). *The reduction map* $\mathrm{Smf}/\tau^4 \to \mathrm{Smf}/\tau$ *is an isomorphism on homotopy groups in degree* $(24, 0)$.

Proof. The element Δ generates $\pi_{24,0}\,\mathrm{Smf}/\tau$. We showed above that Δ is a d_3-cycle, so by evenness, it is also a d_4-cycle. This means that it lifts to $\pi_{24,4}\,\mathrm{Smf}/\tau^4$, so that the reduction map is surjective. By Corollary 3.72, it is also injective: the homotopy groups of Smf/τ vanish in bidegrees $(24, 1)$, $(24, 2)$ and $(24, 3)$. ■

Notation 9.17. We write $\Delta \in \pi_{24,0}\,\mathrm{Smf}/\tau^4$ for the unique lift of $\Delta \in \pi_{24,0}\,\mathrm{Smf}/\tau$ guaranteed by Lemma 9.16.

Our abuse of notation is mild, due to the uniqueness of the lift. We will freely consider Δ as an element of $\pi_{24,0}\,\mathrm{Smf}/\tau^4$ going forward.

9.3 Page 5

9.3.1 Atomic differentials

Proposition 9.18 (24,0). $d_5(\Delta) = \pm h_2 g$.

Proof. We know that $\nu^3\bar{\kappa}$ is τ-power torsion in Smf by Corollary 8.35. Since ν and $\bar{\kappa}$ project to h_2 and g mod τ respectively, and $h_2^3 g \neq 0$, it follows that $h_2^3 g$ must be the target of a differential. The only possibility is $d_5(\Delta h_2^2) = g h_2^3$. In particular, by the Leibniz rule, this gives us $d_5(\Delta) = \pm g h_2$. ■

9.3.2 Meta-arguments

Proposition 9.19. *The condition of Proposition* 9.4 *holds for* d_5. *Moreover,* Δ^8 *is a* d_5*-cycle.*

Proof. The condition of Proposition 9.4 may be checked directly for the atomic d_5's. Since $8d_5(\Delta) = 0$, the Leibniz rule implies that Δ^8 is a d_5-cycle. ■

Proposition 9.20. *There are no line-crossing d_5-differentials.*

Proof. By Proposition 9.14, we may invoke the meta-argument of Proposition 9.5, which implies that we only need to check this through the 17-stem. The only possible atomic d_5's crossing the line in this range have source h_2 or d, which are permanent cycles. ■

9.3.3 Lifts

We now pick up our task of lifting elements with more vigour. Unlike the class Δ in Smf/τ^4 from Notation 9.17, most of the classes here turn out to lift all the way to Smf, but we will not need this.

Lemma 9.21 (25,1), (97,1), (121,1). *The reduction maps* $\mathrm{Smf}/\tau^8 \to \mathrm{Smf}/\tau$ *induces an isomorphism on bigraded homotopy groups in degrees* $(25,1)$, $(97,1)$, *and* $(121,1)$.

The argument is a synthetic version of the statement that $h_1\Delta$ is a $d_{\leqslant 8}$-cycle for degree reasons: all potential targets either support or are hit by a shorter differential.

Proof. We start with the first map. The target $\pi_{25,1}\,\mathrm{Smf}/\tau$ is generated by $h_1\Delta$; we will first show that this generator lifts. This is equivalent to $\partial_1^8(h_1\Delta) = 0$, so it suffices to show that $\pi_{24,3}\,\mathrm{Smf}/\tau^7 = 0$. This follows from Theorem 3.70, as the class in $(24,4)$ supports a d_3, while the classes in $(24,8)$ are hit by a d_3.

It remains therefore only to show that the reduction map is injective. This too follows from an application of Theorem 3.70: the nonzero elements in $\pi_{25,5}\,\mathrm{Smf}/\tau$ support a d_3.

The other two cases follow in the exact same way. ■

Notation 9.22.

- (25,1) We write $\eta_1 \in \pi_{25,1}\,\mathrm{Smf}/\tau^8$ for the unique lift of $h_1\Delta \in \pi_{25,1}\,\mathrm{Smf}/\tau$.
- (97,1) We write $\eta_4 \in \pi_{97,1}\,\mathrm{Smf}/\tau^8$ for the unique lift of $h_1\Delta^4 \in \pi_{97,1}\,\mathrm{Smf}/\tau$.
- (121,1) We write $\eta_5 \in \pi_{121,1}\,\mathrm{Smf}/\tau^8$ for the unique lift of $h_1\Delta^5 \in \pi_{121,1}\,\mathrm{Smf}/\tau$.

Warning 9.23. The element η_5 does not lift beyond Smf/τ^{22}: as we will see later in Proposition 9.57, the element $h_1\Delta^5$ supports a d_{23}.

Next, we turn to Δ-multiples of h_2. Here we run into two problems. First, the lift from Smf/τ to Smf/τ^{14} is not uniquely defined. Second, in order to describe the group structure, we need to take the relation $4\nu = \tau^2\eta^3$ from Lemma 9.15 into account, but this is invisible to Smf/τ.

Both problems are solved by working with Smf/τ^4 instead of Smf/τ. For the first problem, we instead lift Δ-multiples of ν in Smf/τ^4. As the class Δ in Smf/τ^4 is uniquely determined, this procedure also uniquely specifies these lifts. For the second problem, we note that the relation $4\nu = \tau^2\eta^3$ is visible in Smf/τ^4. We can transport this by multiplying by powers of Δ.

Lemma 9.24. *We have isomorphisms of abelian groups*

$$\begin{aligned} \pi_{27,1}\,\mathrm{Smf}/\tau^4 &\cong \mathbf{Z}/8\langle \nu\Delta\rangle, \\ \pi_{51,1}\,\mathrm{Smf}/\tau^4 &\cong \mathbf{Z}/8\langle \nu\Delta^2\rangle, \\ \pi_{123,1}\,\mathrm{Smf}/\tau^4 &\cong \mathbf{Z}/8\langle \nu\Delta^5\rangle, \\ \pi_{147,1}\,\mathrm{Smf}/\tau^4 &\cong \mathbf{Z}/8\langle \nu\Delta^6\rangle. \end{aligned}$$

Proof. Lemma 9.15 implies that $\pi_{3,1}\,\mathrm{Smf}/\tau^4 \cong \mathbf{Z}/8\langle\nu\rangle$. We claim that multiplication by $\Delta \in \pi_{24,0}\,\mathrm{Smf}/\tau^4$ induces an isomorphism

$$\Delta\colon\ \pi_{3,1}\,\mathrm{Smf}/\tau^4 \xrightarrow{\cong} \pi_{27,1}\,\mathrm{Smf}/\tau^4.$$

To see this, note that multiplication by Δ induces an injection on Smf/τ in degrees $(3, 1+m)$ for $0 \leqslant m \leqslant 3$. Moreover, there are no $d_{\leqslant 4}$-differentials entering $\pi_{3,1+m}\,\mathrm{Smf}/\tau$, and the Δ-multiples of these elements in $\pi_{27,1+m}\,\mathrm{Smf}/\tau$ are also not hit by $d_{\leqslant 4}$-differentials. It follows that multiplication by Δ induces an injection on the relevant groups appearing in (3.71) in Theorem 3.70. As a result, Theorem 3.70 implies that multiplication by Δ on Smf/τ^4 is an injection on degree $(3,1)$. Using the same Theorem 3.70, we see that the target group is of the same size as the source, showing the map is indeed an isomorphism. The other cases follow similarly. ∎

Lemma 9.25. *The reduction map* $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau$ *is injective on homotopy groups in bidegrees* $(27,1)$, $(51,1)$, $(91,1)$, $(123,1)$ *and* $(147,1)$. *Moreover, these maps can be identified, respectively, with*

$$\begin{aligned} \mathbf{Z}/4\langle x\rangle &\longrightarrow \mathbf{Z}/8\langle \nu\Delta\rangle, & x &\longmapsto 2\nu\Delta, \\ \mathbf{Z}/8\langle y\rangle &\longrightarrow \mathbf{Z}/8\langle \nu\Delta^2\rangle, & y &\longmapsto \nu\Delta^2, \\ \mathbf{Z}/8\langle z\rangle &\longrightarrow \mathbf{Z}/8\langle \nu\Delta^2\rangle, & z &\longmapsto \nu\Delta^4, \\ \mathbf{Z}/4\langle w\rangle &\longrightarrow \mathbf{Z}/8\langle \nu\Delta^5\rangle, & w &\longmapsto 2\nu\Delta^5, \\ \mathbf{Z}/8\langle t\rangle &\longrightarrow \mathbf{Z}/8\langle \nu\Delta^6\rangle, & t &\longmapsto \nu\Delta^6. \end{aligned}$$

Finally, the reduction map $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau^{10}$ *is an isomorphism in bidegrees* $(51,1)$ *and* $(147,1)$.

Proof. We study the case of the first of the five maps; the other ones are similar. We have an exact sequence

$$\pi_{27,5}\,\mathrm{Smf}/\tau \xrightarrow{\tau^4} \pi_{27,1}\,\mathrm{Smf}/\tau^5 \longrightarrow \pi_{27,1}\,\mathrm{Smf}/\tau^4 \xrightarrow{\partial_4^5} \pi_{26,6}\,\mathrm{Smf}/\tau.$$

The group on the left vanishes. The last map is equal to the d_5-differential by Corollary 3.69: indeed, there are no shorter differentials entering or leaving bidegree $(26,6)$, so $E_5 = E_2$ in this bidegree. Since $d_5(2h_2\Delta) = 0$, this means that there is a unique lift of $2\nu\Delta$ from Smf/τ^4 to Smf/τ^5. Since $d_5(h_2\Delta) \neq 0$, the element $\nu\Delta$ does not lift. This means that $\pi_{27,1}\,\mathrm{Smf}/\tau^5$ is isomorphic to $\mathbf{Z}/4$ and is generated by this unique lift of $2\nu\Delta$.

Next, we have an exact sequence

$$\pi_{27,6}\,\mathrm{Smf}/\tau^9 \xrightarrow{\tau^5} \pi_{27,1}\,\mathrm{Smf}/\tau^{14} \longrightarrow \pi_{27,1}\,\mathrm{Smf}/\tau^5 \xrightarrow{\partial_5^{14}} \pi_{26,7}\,\mathrm{Smf}/\tau^9.$$

The outer two groups vanish by Theorem 3.70, so the middle map is an isomorphism.

Finally, the last claim is checked using Theorem 3.70. ■

As before, we use a subscript to denote the power of Δ present in its mod τ reduction. In some cases, the element cannot truly reduce to $h_2\Delta^i$, but has to differ by an element of small degree (in this case, 2). Note that our notation does not indicate this.

Notation 9.26.

- (27,1) We write $\nu_1 \in \pi_{27,1}\,\mathrm{Smf}/\tau^{14}$ for the unique lift of $2\nu\Delta \in \pi_{27,1}\,\mathrm{Smf}/\tau^4$.
- (51,1) We write $\nu_2 \in \pi_{51,1}\,\mathrm{Smf}/\tau^{14}$ for the unique lift of $\nu\Delta^2 \in \pi_{51,1}\,\mathrm{Smf}/\tau^4$.
- (99,1) We write $\nu_4 \in \pi_{99,1}\,\mathrm{Smf}/\tau^{14}$ for the unique lift of $\nu\Delta^4 \in \pi_{99,1}\,\mathrm{Smf}/\tau^4$.
- (123,1) We write $\nu_5 \in \pi_{123,1}\,\mathrm{Smf}/\tau^{14}$ for the unique lift of $2\nu\Delta^5 \in \pi_{123,1}\,\mathrm{Smf}/\tau^4$.
- (147,1) We write $\nu_6 \in \pi_{147,1}\,\mathrm{Smf}/\tau^{14}$ for the unique lift of $\nu\Delta^6 \in \pi_{147,1}\,\mathrm{Smf}/\tau^4$.

9.3.4 Relations

Lemma 9.27 (23,5). *Multiplication by $\bar\kappa$ induces an isomorphisms*

$$\bar\kappa\colon\ \pi_{3,1}\,\mathrm{Smf}/\tau^{14} \xrightarrow{\cong} \pi_{23,5}\,\mathrm{Smf}/\tau^{14}$$
$$\bar\kappa\colon\ \pi_{99,1}\,\mathrm{Smf}/\tau^{14} \xrightarrow{\cong} \pi_{119,5}\,\mathrm{Smf}/\tau^{14}$$

In particular, the reduction maps $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau$ in degrees $(23,5)$ and $(119,5)$ can be identified with

$$\mathbf{Z}/8\langle\nu\bar\kappa\rangle \twoheadrightarrow \mathbf{Z}/4\langle h_2 g\rangle, \qquad \nu\bar\kappa \longmapsto h_2 g,$$
$$\mathbf{Z}/8\langle\nu_4\bar\kappa\rangle \twoheadrightarrow \mathbf{Z}/4\langle h_2 g\Delta^4\rangle, \qquad \nu_4\bar\kappa \longmapsto h_2 g\Delta^4.$$

Proof. In the same way as in Lemma 9.24, one can show that $\bar\kappa$ induces an isomorphism on Smf/τ^4 in degrees $(3,1)$ and $(99,1)$. Using Theorem 3.70, we see that the reduction map $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau^4$ is injective in the relevant degrees, proving the claim. ■

9.3.5 Total differentials

For many of the later pages, it will be crucial to compute a truncated total differential on Δ. Given the d_5 of Proposition 9.18, a natural guess for $\partial_1^\infty(\Delta)$ would be $\nu\bar{\kappa}$. Because we at this point do not know the fate of the elements in very high filtration, we cannot compute the entire differential, but only a truncated version. While computing the 8-truncated version would be enough to deduce $d_7(\Delta^4)$, it is hardly any more work to at this point record the 14-truncated version, and this will be needed later on. As Δ lifts uniquely to Smf/τ^4, we can work with ∂_4^{14} instead. This is both easier to compute, taking values in Smf/τ^{10} rather than Smf/τ^{13}, and also gives more information, as we will see in our later computations.

In what follows, we think of $\nu\bar{\kappa}$ as defining an element in Smf/τ^{10} via the reduction map.

Proposition 9.28 (24,0). *We have $\partial_4^{14}(\Delta) = u \cdot \nu\bar{\kappa}$ in $\pi_{23,5}\,\mathrm{Smf}/\tau^{10}$, where $u \in (\mathbf{Z}/8)^\times$.*

Proof. First we show that

$$\partial_4^8(\Delta) = u \cdot \nu\bar{\kappa}$$

where $u \in (\mathbf{Z}/8)^\times$ is a unit. To show this, we first claim that we have a commutative diagram

$$\begin{tikzcd}
\pi_{24,0}\,\mathrm{Smf}/\tau^4 \arrow[r, "\partial_4^8"] \arrow[d] & \pi_{23,5}\,\mathrm{Smf}/\tau^4 \arrow[d] \\
\pi_{24,0}\,\mathrm{Smf}/\tau \arrow[r, "d_5"'] & \pi_{23,5}\,\mathrm{Smf}/\tau,
\end{tikzcd}$$

where the vertical maps are the reductions mod τ. Indeed, we use Corollary 3.69 together with the fact that the differentials d_2, d_3, d_4 vanish on $(24,0)$ and hit no elements in $(23,5)$, so that E_5 is the same as E_2 in these bidegrees. The group $\pi_{24,0}\,\mathrm{Smf}/\tau$ is the free (2-local) abelian group generated by Δ and c_4^3. The element c_4 is a $d_{\leqslant 10}$-cycle for degree reasons, so we can ignore this summand in further analysis.

The left vertical map is an isomorphism by Lemma 9.16, and the right vertical map is surjective by Lemma 9.27. In other words, restricting to the Δ-summand, the commutative diagram is of the form

$$\begin{tikzcd}
\mathbf{Z}_{(2)} \arrow[r] \arrow[dr, two heads] & \mathbf{Z}/8 \arrow[d, two heads] \\
& \mathbf{Z}/4.
\end{tikzcd}$$

This means the top horizontal map must send 1 to a unit in $\mathbf{Z}/8$, which is exactly the claim about $\partial_4^8(\Delta)$.

Using Theorem 3.70, we see that the reduction $\pi_{23,5}\,\mathrm{Smf}/\tau^{10} \to \pi_{23,5}\,\mathrm{Smf}/\tau^4$ is injective. It is also surjective, as $h_2 g$ is a d_5-boundary, so in particular a permanent cycle. The claim about $\partial_4^{14}(\Delta)$ therefore follows. ■

9.4 Page 7

9.4.1 Atomic differentials

Proposition 9.29.

- (24,0) $d_7(4\Delta) = h_1^3 g$.
- (48,0) $d_7(2\Delta^2) = h_1^3 g\Delta$.
- (72,0) $d_7(4\Delta^3) = h_1^3 g\Delta^2$.
- (96,0) $d_7(\Delta^4) = h_1^3 g\Delta^3$.
- (120,0) $d_7(4\Delta^5) = h_1^3 g\Delta^4$.
- (144,0) $d_7(2\Delta^6) = h_1^3 g\Delta^5$.
- (168,0) $d_7(4\Delta^7) = h_1^3 g\Delta^6$.

The atomic d_7-differentials on powers of Δ are difficult to deduce directly. Our approach essentially deduces these from the d_5 on Δ, through the use of the total differential on Δ from Proposition 9.28. In all of this, the lift of Δ to Smf/τ^4 (Notation 9.17) is the crucial input to make the following arguments work.

Proof. Let us start with $d_7(\Delta^4)$. First of all, from Proposition 9.28, we learn in particular that $\partial_4^8(\Delta) = u \cdot \nu\bar{\kappa}$. The synthetic Leibniz rule of Theorem 3.40 tells us that

$$\partial_4^8(\Delta^4) = 4\Delta^3 \cdot \partial_4^8(\Delta) = u \cdot 4\nu\bar{\kappa}\Delta^3.$$

Combining this with the relation $4\nu = \tau^2\eta^3$ from Lemma 9.15, we learn that

$$\partial_4^8(\Delta^4) = u \cdot \tau^2\eta^3\bar{\kappa}\Delta^3 = \tau^2\eta^3\bar{\kappa}\Delta^3,$$

where we use that η is 2-torsion to ignore the unit. From this, the differential $d_7(\Delta^4) = h_1^3 g\Delta^3$ follows by Proposition 3.37.

The remaining differentials follow similarly using the synthetic Leibniz rule for ∂_4^8 and the same relation $4\nu = \tau^2\eta^3$. ■

9.4.2 Meta-arguments

Proposition 9.30. *The condition of Proposition 9.4 holds for d_7. Moreover, Δ^8 is a d_7-cycle.*

Proof. The condition of Proposition 9.4 is checked directly as before. Since $2d_7(\Delta^4) = 0$, the Leibniz rule implies that Δ^8 is a d_7-cycle. ■

Proposition 9.31. *There are no line-crossing d_7-differentials.*

Proof. By Proposition 9.19, we may invoke the meta-argument of Proposition 9.5, which implies we only need to check for line-crossing differentials through the 25-stem. The only possible atomic d_7 crossing the line in this range has source $2c_6$, which is a permanent cycle because it is a transfer; see Corollary 8.23. ■

9.4.3 Total differentials

In the proof of Proposition 9.29, we computed the total differentials ∂_4^8 on powers of Δ using the synthetic Leibniz rule. The synthetic Leibniz rule does not apply to the higher truncations ∂_4^N for $N > 8$, so we have to manually compute these total differentials.

Proposition 9.32. *There is a unit $u \in (\mathbf{Z}/8)^\times$ such that*

- (48,0) $\partial_4^{14}(\Delta^2) = u \cdot \nu_1\bar{\kappa}$,
- (72,0) $\partial_4^{14}(\Delta^3) = u \cdot 3\nu_2\bar{\kappa}$,
- (144,0) $\partial_4^{14}(\Delta^6) = u \cdot 3\nu_5\bar{\kappa}$,
- (168,0) $\partial_4^{14}(\Delta^7) = u \cdot 7\nu_6\bar{\kappa}$.

Proof. Using Theorem 3.70, we see that $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau^4$ is injective in the relevant bidegrees. As ∂_4^{14} reduces to ∂_4^8 when taken mod τ^4, we are reduced to the computation of ∂_4^8 in these degrees. This now follows from Proposition 9.28 and the synthetic Leibniz rule Theorem 3.40. For example, we have

$$\partial_4^8(\Delta^2) = 2\Delta \cdot \partial_4^8(\Delta) = 2\Delta \cdot u \cdot \nu\bar{\kappa} = u \cdot 2\nu\Delta\bar{\kappa} \qquad \text{in } \pi_{47,5}\,\mathrm{Smf}/\tau^4,$$

and ν_1 is defined as a lift of $2\nu\Delta \in \pi_{27,1}\,\mathrm{Smf}/\tau^4$. The other cases are proved in the exact same way. ■

9.4.4 Relations

To later be able to define Toda brackets, we will need relations like $\tau^4\nu\bar{\kappa} = 0$. Such a relation is plausible because $h_2 g$ is the target of a d_5-differential. However, knowing this differential is not enough to deduce that $\tau^4\nu\bar{\kappa} = 0$ holds (see Warning 3.38), so we have to check this relation by hand.

Lemma 9.33.

- (23,1) $\tau^4\nu\bar{\kappa} = 0$ *in* Smf/τ^{12}.

- (119,1) $\tau^4\nu_4\bar{\kappa} = 0$ *in* Smf/τ^{12}.

Proof. Recall from Proposition 9.18 that $d_5(\Delta) = \pm h_2 g$. By Omnibus Theorem 3.62 (3), this means that there exists a τ^4-torsion lift of $h_2 g$ to $\pi_{23,5}\,\mathrm{Smf}/\tau^{12}$. We claim that this must be $\nu\bar{\kappa}$ itself. To see this, we first use Lemma 9.27 to conclude that $\nu\bar{\kappa}$ and $5\nu\bar{\kappa}$ are the only lifts of $h_2 g$. At least one of them is therefore τ^4-torsion; we claim that this implies that both are τ^4-torsion.

Since their difference is $4\nu\bar{\kappa} = \tau^2\eta^3\bar{\kappa}$, it suffices to establish that $\eta^3\bar{\kappa}$ in $\pi_{23,6}\,\mathrm{Smf}/\tau^{12}$ is τ^6-torsion. The element $h_1^3 g$ is hit by a d_7, so by Theorem 3.67, there exists a τ^6-torsion lift of it to Smf/τ^{12}. But $\eta^3\bar{\kappa}$ is the only element in $\pi_{23,7}\,\mathrm{Smf}/\tau^{12}$ that lifts $h_1^3 g$: the mod τ reduction map is injective in this bidegree by Theorem 3.70. We learn that $\eta^3\bar{\kappa}$ must therefore be this τ^6-torsion lift.

A similar argument applied to the differential $d_5(\Delta^5) = \pm h_2 g\Delta^4$ yields the relation $\tau^4\nu_4\bar{\kappa} = 0$. ∎

Lemma 9.34.

- (27,1) $2\nu_1 = \tau^2\eta^2\eta_1$ *in* Smf/τ^8.
- (123,1) $2\nu_5 = \tau^2\eta^2\eta_5$ *in* Smf/τ^8.

Proof. Using Lemma 9.15, we have the relation $4\nu\Delta = \tau^2\eta^3\Delta$ in Smf/τ^4. This means that $2\nu_1 \in \pi_{27,1}\,\mathrm{Smf}/\tau^{14}$ and $\tau^2\eta^2\eta_1 \in \pi_{27,1}\,\mathrm{Smf}/\tau^8$ reduce to the same element in Smf/τ^4. It is therefore enough to establish that $\pi_{27,1}\,\mathrm{Smf}/\tau^8 \to \pi_{27,1}\,\mathrm{Smf}/\tau^4$ is injective, which follows from Theorem 3.70. The exact same arguments apply to the second relation. ∎

Corollary 9.35.

- (47,1) $\tau^4\nu_1\bar{\kappa} = 0$ *in* Smf/τ^{12}.
- (143,1) $\tau^4\nu_5\bar{\kappa} = 0$ *in* Smf/τ^{12}.

Proof. It suffices to prove this statement in Smf/τ^6, as the reduction map is injective in this bidegree by Theorem 3.70. Similar to the argument of Lemma 9.24, we see that multiplication by $\bar{\kappa}$ induces an isomorphism from $\pi_{27,1}\,\mathrm{Smf}/\tau^4$ to $\pi_{47,5}\,\mathrm{Smf}/\tau^4$, and that $\pi_{47,5}\,\mathrm{Smf}/\tau^6 \to \pi_{47,5}\,\mathrm{Smf}/\tau^4$ is injective. Since $\pi_{47,5}\,\mathrm{Smf}/\tau^4$ is generated by $\bar{\kappa}\nu_1$, which lifts to Smf/τ^6, we therefore conclude that

$$\pi_{47,5}\,\mathrm{Smf}/\tau^6 \cong \mathbf{Z}/4\langle\bar{\kappa}\nu_1\rangle.$$

The only lifts of $2gh_2\Delta$ are therefore $\bar{\kappa}\nu_1$ and $3\bar{\kappa}\nu_1$. Because $2gh_2\Delta$ is the target of a d_5-differential, it must have a τ^4-torsion lift to Smf/τ^6. Using the relation from Lemma 9.34, we see that their difference is τ^2-divisible. In the same way as in Lemma 9.33, we can deduce from this that both lifts are τ^4-torsion, so in particular, $\tau^4\bar{\kappa}\nu_1 = 0$. The proof of the equality $\tau^4\nu_5\bar{\kappa} = 0$ follows from the same arguments. ∎

9.4.5 Lifts

We now turn to lifting powers of Δ times c. This is more difficult, as it involves ruling out a d_9 on them. We cannot use degree arguments: it turns out that the potential target for these d_9's each support a d_{11}. Instead we give a Toda bracket argument to show that these elements lift.

Lemma 9.36.

- (32,2) $\langle \nu_1, \nu, \eta \rangle$ *consists of one element in* Smf/τ^{14}, *which lifts* $c\Delta$ *in* Smf/τ.
- (104,2) $\langle \nu, 2\nu_4, \eta \rangle$ *consists of one element in* Smf/τ^{14}, *which lifts* $c\Delta^4$ Smf/τ.
- (128,2) $\langle \nu_1, \nu_4, \eta \rangle$ *consists of one element in* Smf/τ^{14}, *which lifts* $c\Delta^5$ *in* Smf/τ.

Proof. We use Proposition B.14 to verify these claims. One can easily check that the indeterminacy of the mod τ reductions of these brackets vanishes, as well as compute the values of the corresponding Massey products. An application of Theorem 3.70 tells us that the reduction map $\mathrm{Smf}/\tau^{14} \to \mathrm{Smf}/\tau$ is injective in these bidegrees. This means that the Toda brackets in Smf/τ^{14} therefore also consist of singletons, ending the argument. ■

Notation 9.37.

- (32,2) We write $\varepsilon_1 = \langle \nu_1, \nu, \eta \rangle$ in $\pi_{32,2}\,\mathrm{Smf}/\tau^{12}$.
- (104,2) We write $\varepsilon_4 = \langle \nu, 2\nu_4, \eta \rangle$ in $\pi_{104,2}\,\mathrm{Smf}/\tau^{12}$.
- (128,2) We write $\varepsilon_5 = \langle \nu_1, \nu_4, \eta \rangle$ in $\pi_{128,2}\,\mathrm{Smf}/\tau^{12}$.

9.4.6 Toda brackets

Lemma 9.38.

- (27,1) $\nu_1 = \langle \bar{\kappa}, \tau^4\nu, 2\nu \rangle$ *in* Smf/τ^{12}.
- (51,1) $\nu_2 = \langle \bar{\kappa}, \tau^4\nu_1, \nu \rangle$ *in* $\pi_{51,1}\,\mathrm{Smf}/\tau^{12}$.
- (123,1) $\nu_5 = \langle \bar{\kappa}, \tau^4\nu, 2\nu_4 \rangle$ *in* Smf/τ^{12}.
- (147,1) $\nu_6 = \langle \bar{\kappa}, \tau^4\nu_1, \nu_4 \rangle$ *in* Smf/τ^{12}.
- (39,3) $\eta_1\kappa = \langle \nu_1, \nu, \varepsilon \rangle$ *in* Smf/τ^{10}.
- (111,3) $\eta_4\kappa = \langle \nu, 2\nu_4, \varepsilon \rangle$ *in* Smf/τ^{12}.
- (135,3) $\eta_5\kappa = \langle \nu_1, \nu_4, \varepsilon \rangle$ *in* Smf/τ^{10}.
- (117,3) $\tau^2\eta_4\bar{\kappa} = \langle \nu, 2\nu_4, \kappa \rangle$ *in* Smf/τ^{12}.
- (20,4) $\pm 2\bar{\kappa} \in \langle \nu, \eta, \eta\kappa \rangle$ *in* Smf/τ^{10}.

Proof. The brackets for ν_1 and ν_5 are nonempty by Lemma 9.33. These brackets also have zero indeterminacy as $\pi_{7,-3}\,\mathrm{Smf}/\tau^{12}$ and $\pi_{103,-3}\,\mathrm{Smf}/\tau^{12}$ both vanish and $\pi_{24,0}\,\mathrm{Smf}/\tau^{12}$ and $\pi_{120,0}\,\mathrm{Smf}/\tau^{12}$ are both 2ν-torsion. The associated Massey product on E_5 contains $2\nu\Delta$, so from Theorem B.15 it follows that $\langle\bar{\kappa}, \tau^4\nu, 2\nu\rangle = \nu_1$. Similarly, we see $\nu_5 = \langle\bar{\kappa}, \tau^4\nu, 2\nu_4\rangle$. The brackets for ν_2 and ν_6 follow similarly, except to see they are nonempty one refers to Corollary 9.35 and the proof of Lemma 9.36.

For the bracket expressions for $\eta_1\kappa$, $\eta_4\kappa$, and $\eta_5\kappa$, we use Proposition B.14. One uses Theorem 3.70 to see that the appropriate reduction map to Smf/τ is injective in these degrees, so it suffices to work in Smf/τ. In this case, these brackets are easily seen to have no indeterminacy, and therefore follow from Proposition 8.36 by multiplication by a power of Δ. For $\tau^2\eta_4\bar{\kappa}$, we make the same arguments but only reduce to Smf/τ^4.

For the last Toda bracket for $\pm 2\bar{\kappa}$, we again use Proposition B.14 and the fact that the reduction map

$$\pi_{20,4}\,\mathrm{Smf}/\tau^{10} \longrightarrow \mathrm{Smf}/\tau$$

is injective by Theorem 3.70, and surjective as the generator g is hit by $\bar{\kappa}$ by Notation 8.34. We can now use Bauer's computation of $2g \in \langle h_2, h_1, h_1 d\rangle$ from [Bau08, Appendix A] together with Proposition B.14 and Corollary 8.15 which validates this relation in Smf/τ. ∎

9.4.7 Hidden extensions

Lemma 9.39.

- (28,2) $\nu_1\eta = \tau^4\varepsilon\bar{\kappa}$ *in* Smf/τ^{12}.
- (35,3) $\nu_1\varepsilon = \tau^4\eta\kappa\bar{\kappa}$ *in* Smf/τ^{12}.
- (41,3) $\nu_1\kappa = \tau^6\eta\bar{\kappa}^2$ *in* Smf/τ^{12}.
- (52,2) $\nu_2\eta = \tau^4\varepsilon_1\bar{\kappa}$ *in* Smf/τ^{12}.
- (59,3) $\nu_2\varepsilon = \tau^4\eta_1\kappa\bar{\kappa}$ *in* Smf/τ^{10}.
- (124,2) $\nu_5\eta = \tau^4\varepsilon_4\bar{\kappa}$ *in* Smf/τ^{12}.
- (131,3) $\nu_5\varepsilon = \tau^4\eta_4\kappa\bar{\kappa}$ *in* Smf/τ^{12}.
- (137,3) $\nu_5\kappa = \tau^6\eta_4\bar{\kappa}^2$ *in* Smf/τ^{12}.
- (148,2) $\nu_6\eta = \tau^4\varepsilon_5\bar{\kappa}$ *in* Smf/τ^{12}.
- (155,3) $\nu_6\varepsilon = \tau^4\eta_5\kappa\bar{\kappa}$ *in* Smf/τ^{10}.

Proof. The Toda bracket $\nu_1 = \langle\bar{\kappa}, \tau^4\nu, 2\nu\rangle$ of Lemma 9.38, the shuffling formulas of Proposition B.12, the naturality of Toda brackets of Lemma B.9, and the Toda bracket

$\varepsilon \in \langle \nu, 2\nu, \eta \rangle$ of Proposition 8.36 yield the relations

$$\nu_1 \eta = \langle \bar{\kappa}, \tau^4 \nu, 2\nu \rangle \eta = \bar{\kappa} \langle \tau^4 \nu, 2\nu, \eta \rangle \supseteq \bar{\kappa} \tau^4 \langle \nu, 2\nu, \eta \rangle \ni \tau^4 \varepsilon \bar{\kappa}$$

in $\pi_{28,2} \operatorname{Smf}/\tau^{12}$. Similarly, using the brackets $\nu_5 = \langle \bar{\kappa}, \tau^4 \nu_4, 2\nu \rangle$ and $\varepsilon_4 = \langle \nu, 2\nu_4, \eta \rangle$ of Lemma 9.38 and Notation 9.37, we obtain $\nu_5 \eta = \tau^4 \varepsilon_4 \bar{\kappa}$.

The other extensions on ν_1 and ν_2 follow similarly, referring to Proposition 8.36 and Lemma 9.38 when necessary. The same goes for all of the extensions on ν_2 and ν_6, referring to Lemma 9.38 and the definition of ε_1 and ε_5 from Notation 9.37. ■

9.5 Page 9

9.5.1 Atomic differentials

Proposition 9.40.

- (49,1) $d_9(h_1 \Delta^2) = cg^2$.[1]
- (56,2) $d_9(c\Delta^2) = h_1 dg^2$.
- (73,1) $d_9(h_1 \Delta^3) = cg^2 \Delta$.
- (81,2) $d_9(c\Delta^3) = h_1 dg^2 \Delta$.
- (145,1) $d_9(h_1 \Delta^6) = cg^2 \Delta^4$.
- (152,2) $d_9(c\Delta^6) = h_1 dg^2 \Delta^4$.
- (169,1) $d_9(h_1 \Delta^7) = cg^2 \Delta^5$.
- (176,2) $d_9(c\Delta^7) = h_1 dg^2 \Delta^6$.

Proof. From Proposition 9.32 it follows that also $\partial_4^{10}(\Delta^2) = u \cdot \nu_1 \bar{\kappa}$ for some $u \in (\mathbf{Z}/8)^\times$. Using Lemma 9.39, we learn that

$$\partial_4^{10}(\eta \Delta^2) = \eta \cdot \partial_4^{10}(\Delta^2) = u \cdot \eta \cdot \nu_1 \bar{\kappa} = u \cdot \tau^4 \varepsilon \bar{\kappa}^2$$

and

$$\partial_4^{10}(\varepsilon \Delta^2) = u \cdot \varepsilon \nu_1 \bar{\kappa} = u \cdot \tau^4 \eta \kappa \bar{\kappa}^2.$$

We obtain our desired d_9 on $h_1 \Delta^2$ courtesy of Proposition 3.37. The other differentials follow similarly using the hidden extensions of Lemma 9.39. ■

[1] As an alternative to the proof of this differential provided below, one can observe that

$$\tau^8 \varepsilon \bar{\kappa}^2 = \tau^8 \bar{\kappa}^2 \langle \nu, \eta, \nu \rangle = \tau^4 \bar{\kappa} \langle \tau^4 \bar{\kappa}, \nu, \eta \rangle \nu = 0,$$

and then apply Theorem 3.67.

9.5.2 Meta-arguments

Proposition 9.41. *The condition of Proposition 9.4 holds for d_9. Moreover, Δ^8 is a d_9-cycle.*

Proof. The condition of Proposition 9.4 is checked directly as before. The class Δ^8 is a d_9-cycle for degree reasons. ■

Proposition 9.42. *There are no line-crossing d_9-differentials.*

Proof. By Proposition 9.30, we may invoke the meta-argument of Proposition 9.5, which implies we only need to check for line crossing differentials through the 32-stem. The only possible atomic d_9's crossing the line in this range have source h_1, c, $h_1\Delta$, or $c\Delta$. The first two are permanent cycles. The third is a d_9-cycle because the only possible target supports a d_9 differential. Lastly, the d_9 on $c\Delta$ is excluded by the earlier Lemma 9.36 (combined with Theorem 3.62 (1)). ■

9.5.3 Lifts

In Lemma 9.21, we provided lifts η_1, η_4 and η_5 to Smf/τ^8. We will need further lifts of two of these elements.

Lemma 9.43 (25,1), (121,1). *The reduction maps $\mathrm{Smf}/\tau^{10} \to \mathrm{Smf}/\tau^8$ induces an isomorphism on bigraded homotopy groups in degrees $(25,1)$ and $(121,1)$.*

Proof. Our computation of the d_9-differentials shows that $h_1\Delta$ is a $d_{\leqslant 10}$-cycle. By Theorem 3.67, it therefore lifts to Smf/τ^{10}; since η_1 is by Lemma 9.21 the unique lift of $h_1\Delta$ to Smf/τ^8, it follows that η_1 lifts to Smf/τ^{10}. It follows from Theorem 3.70 that the reduction map $\mathrm{Smf}/\tau^{10} \to \mathrm{Smf}/\tau^8$ is injective on bidegree $(25,1)$, proving the claim. The case for bidegree $(121,1)$ and lifting $h_1\Delta^5$ is the same. ■

Notation 9.44.

- We write $\eta_1 \in \pi_{25,1}\,\mathrm{Smf}/\tau^{10}$ for the unique lift of $\eta_1 \in \pi_{25,1}\,\mathrm{Smf}/\tau^8$. In particular, it is also the unique lift of $h_1\Delta \in \pi_{25,1}\,\mathrm{Smf}/\tau$.
- We write $\eta_5 \in \pi_{121,1}\,\mathrm{Smf}/\tau^{10}$ for the unique lift of $\eta_5 \in \pi_{121,1}\,\mathrm{Smf}/\tau^8$. In particular, it is also the unique lift of $h_1\Delta^5 \in \pi_{121,1}\,\mathrm{Smf}/\tau$.

Lemma 9.45 (116,4). *The reduction maps $\mathrm{Smf}/\tau^{20} \to \mathrm{Smf}/\tau^{10}$ and $\mathrm{Smf}/\tau^{10} \to \mathrm{Smf}/\tau$ are injective on homotopy groups in degree $(116,4)$. Moreover, the map $\pi_{116,4}\,\mathrm{Smf}/\tau^{20} \to \pi_{116,4}\,\mathrm{Smf}/\tau$ can be identified with*

$$\mathbf{Z}/4\langle x\rangle \hookrightarrow \mathbf{Z}/8\langle g\Delta^4\rangle, \quad x \longmapsto 2g\Delta^4.$$

Proof. It follows from Theorem 3.70 that the maps are injective, so it remains to be shown that $2g\Delta^4$ lifts to Smf/τ^{20}, while $g\Delta^4$ does not lift to Smf/τ^{10}. The latter claim is clear, as it supports a d_7. To see that $2g\Delta^4$ lifts to Smf/τ^{20}, it suffices to show that $\partial_4^{20}(2\bar{\kappa}\Delta^4) = 0$, for which it in turn suffices to show that $\pi_{115,9}\,\mathrm{Smf}/\tau^{16}$ vanishes. This follows from another application of Theorem 3.70. ■

Notation 9.46. We write $\bar{\kappa}_4 \in \pi_{116,4}\,\mathrm{Smf}/\tau^{20}$ for the unique lift of the element $2g\Delta^4 \in \pi_{116,4}\,\mathrm{Smf}/\tau$. Note that it is also the unique lift to Smf/τ^{10} of $2g\Delta^4$.

9.5.4 Toda brackets

Lemma 9.47.

- (25,1) $\eta_1 \in \langle \bar{\kappa}, \tau^4\nu, \eta \rangle$ *in* Smf/τ^{10}. *This Toda bracket has indeterminacy given by the subgroup of* $\pi_{25,1}\,\mathrm{Smf}/\tau^{10}$ *spanned by the* κ*-torsion classes.*
- (121,1) $\eta_5 \in \langle \bar{\kappa}, \tau^4\nu_4, \eta \rangle$ *in* Smf/τ^{10}. *This Toda bracket has indeterminacy given by the subgroup of* $\pi_{121,1}\,\mathrm{Smf}/\tau^{10}$ *spanned by the* κ*-torsion classes.*
- (116,4) $\pm\bar{\kappa}_4 \in \langle \nu_4, \eta, \eta\kappa \rangle$ *in* Smf/τ^{10}.

Proof. The first two brackets are nonempty by Lemma 9.33, and it straightforward to compute the indeterminacies. The Massey products on the E_5-page associated with these Toda brackets contains $h_1\Delta$ and $h_1\Delta^5$, respectively, hence the synthetic Moss's Theorem B.15 shows that $\langle \bar{\kappa}, \tau^4\nu, \eta \rangle$ contains η_1 and $\langle \bar{\kappa}, \tau^4\nu_4, \eta \rangle$ contains η_5.

For the Toda bracket expression of $\pm\bar{\kappa}_4$, we note that by Notation 9.46, it suffices to show that this Toda bracket contains a lift of $\pm 2g\Delta^4$. We will show that $\pm 2g\Delta^4 \subseteq \langle h_2\Delta^4, h_1, h_1 d \rangle$. Using Proposition B.14 and the bracket $\pm 2\bar{\kappa} \in \langle \nu, \eta, \eta\kappa \rangle$ of Proposition 8.36, it follows that $\pm 2g \in \langle h_2, h_1, h_1 d \rangle$, so we learn that

$$\pm 2g\Delta^4 \in \pm\Delta^4 \langle h_2, h_1, h_1 d \rangle \subseteq \langle h_2\Delta^4, h_1, h_1 d \rangle.$$

■

9.5.5 Hidden extensions

Lemma 9.48.

- (40,4) $\eta\eta_1\kappa = \pm\tau^4 2\bar{\kappa}^2$ *in* Smf/τ^{10}.
- (136,4) $\eta\eta_5\kappa = \pm\tau^4\bar{\kappa}\bar{\kappa}_4$ *in* Smf/τ^{10}.

Proof. The formula for $\eta\eta_1\kappa$ follows from the Toda bracket expressions of Proposition 8.36 and Lemma 9.38:

$$\eta\eta_1\kappa = \langle \bar{\kappa}, \tau^4\nu, \eta \rangle \eta\kappa = \bar{\kappa}\langle \tau^4\nu, \eta, \eta\kappa \rangle \supseteq \tau^4\bar{\kappa}\langle \nu, \eta, \eta\kappa \rangle \ni \pm\tau^4 2\bar{\kappa}^2.$$

Here the second equality follows from the fact that the indeterminacy of $\langle \bar{\kappa}, \tau^4\nu, \eta \rangle$ is all κ-torsion. The equality for $\eta\eta_5\kappa$ is the same, except we refer only to Lemma 9.38. ■

9.6 Page 11

9.6.1 Atomic differentials

Proposition 9.49.

- (62,2) $d_{11}(d\Delta^2) = h_1 g^3$.
- (158,2) $d_{11}(d\Delta^6) = h_1 g^3 \Delta^4$.

Proof. This follows from the total differentials in Proposition 9.32 and the κ-extensions of Lemma 9.39. ∎

9.6.2 Meta-arguments

Proposition 9.50. *The conditions of Proposition* 9.4 *hold for d_{11}. Moreover, Δ^8 is a d_{11}-cycle.*

Proof. The condition of Proposition 9.4 is checked directly as before. The class Δ^8 is a d_{11}-cycle for degree reasons. ∎

Proposition 9.51. *There are no line-crossing d_{11}-differentials.*

Proof. By Proposition 9.41, we may invoke the meta-argument of Proposition 9.5, which implies we only need to check for line crossing differentials through the 40-stem. The only possible atomic d_{11}'s crossing the line in this range have sources c_4 or $h_1^3\Delta$, but the only possible target supports a d_{11} differential in both cases. ∎

9.6.3 Hidden extensions

To establish the d_{13}-differentials, we need some hidden extensions. These turn out to require four-fold Toda brackets, and computing these is a delicate matter. We provide a very detailed and careful analysis of these in Section 9.9. Using these, we now apply various shuffling formulas to obtain the following.

Lemma 9.52.

- (54,2) $\nu_2 \cdot 2\nu = \tau^8 \bar{\kappa}^2 \widetilde{d}$ *in* Smf/τ^{10}, *where* $\widetilde{d} \in \pi_{14,2}\,\mathrm{Smf}/\tau^{10}$ *is an element that is sent to d under the map* $\mathrm{Smf}/\tau^{10} \to \mathrm{Smf}/\tau$.
- (150,2) $\nu_6 \cdot 2\nu = \tau^8 \bar{\kappa}^2 \widetilde{d}_4$ *in* Smf/τ^{10}, *where* $\widetilde{d}_4 \in \pi_{110,2}\,\mathrm{Smf}/\tau^{10}$ *is an element that is sent to $d\Delta^4$ under the map* $\mathrm{Smf}/\tau^{10} \to \mathrm{Smf}/\tau$.

Proof. In Proposition 9.61, we show that

$$\nu_2 = \langle \nu,\, 2\nu\tau^4,\, \nu\tau^4,\, \bar{\kappa}^2 \rangle \quad \text{in } \pi_{51,1}\,\mathrm{Smf}/\tau^{10}.$$

The Toda brackets $\langle 2\nu, \nu, 2\nu\tau^4\rangle$ and $\langle \nu, 2\nu, \nu\tau^4\rangle$ are strictly zero, so that the four-fold Toda bracket $\langle 2\nu, \nu, 2\nu\tau^4, \nu\tau^4\rangle$ is nonempty and we may apply the shuffling formula of Proposition B.12 to see that

$$2\nu \cdot \nu_2 = 2\nu\langle \nu,\ 2\nu\tau^4,\ \nu\tau^4,\ \bar{\kappa}^2\rangle = \langle 2\nu,\ \nu,\ 2\nu\tau^4,\ \nu\tau^4\rangle\bar{\kappa}^2.$$

In particular, we also see that the set $\langle 2\nu, \nu, 2\nu\tau^4, \nu\tau^4\rangle\bar{\kappa}^2$ is a singleton. Applying the shuffling formula of Proposition B.12, one has

$$\langle 2\nu, \nu, 2\nu, \nu\rangle\tau^8 \subseteq \langle 2\nu, \nu, 2\nu, \nu\tau^8\rangle \subseteq \langle 2\nu, \nu, 2\nu\tau^4, \nu\tau^4\rangle.$$

Since $\langle 2\nu, \nu, 2\nu\tau^4, \nu\tau^4\rangle\bar{\kappa}^2$ is a singleton, we are reduced to showing that every element of $\langle 2\nu, \nu, 2\nu, \nu\rangle$ projects to d mod τ. However, using Proposition B.14, any class in $\langle 2\nu, \nu, 2\nu, \nu\rangle$ projects to $d \in \pi_{14,2}\,\mathrm{Smf}/\tau$, as the latter has zero indeterminacy and contains d by [Bau08, Equation (7.14)].

The expression for $2\nu \cdot \nu_6$ follows by similar arguments. First, we use the Toda bracket expression

$$\nu_6 = \langle \nu_4, 2\nu\tau^4, \nu\tau^4, \bar{\kappa}^2\rangle$$

of Proposition 9.61 and the shuffling formula of Proposition B.12

$$2\nu \cdot \nu_6 = 2\nu\langle \nu_4, 2\nu\tau^4, \nu\tau^4, \bar{\kappa}^2\rangle = \langle 2\nu, \nu_4, 2\nu\tau^4, \nu\tau^4\rangle\bar{\kappa}^2.$$

Combining this with the the containments

$$d\Delta^4 \in \Delta^4\langle 2h_2, h_2, 2h_2, h_2\rangle \subseteq \langle 2h_2\Delta^4, h_2, 2h_2, h_2\rangle \subseteq \langle 2h_2, h_2\Delta^4, 2h_2, h_2\rangle$$

in Smf/τ yields the result. ■

9.7 Page 13

9.7.1 Atomic differentials

Proposition 9.53.

- (75,1) $d_{13}(2h_2\Delta^3) = dg^3$.
- (81,3) $d_{13}(h_2^3\Delta^3) = \pm 2g^4$.
- (171,1) $d_{13}(2h_2\Delta^7) = dg^3\Delta^4$.
- (177,3) $d_{13}(h_2^3\Delta^7) = \pm 2g^4\Delta^4$.

Proof. Recall the total differentials of Proposition 9.32

$$\partial_4^{14}(\Delta^3) = u \cdot 3\nu_2\bar{\kappa} \qquad \text{and} \qquad \partial_4^{14}(\Delta^7) = u \cdot 7\nu_6\bar{\kappa}$$

and the extensions of Lemma 9.52

$$\nu_2 \cdot 2\nu = \tau^8\,\bar{\kappa}^2\,\widetilde{d} \qquad \text{and} \qquad \nu_6 \cdot 2\nu = \tau^8\,\bar{\kappa}\,\widetilde{d}_4,$$

where $\widetilde{d}$ is some element whose mod τ reduction is d, and where $\widetilde{d}_4$ is some element whose mod τ reduction is $d\Delta^4$. Together these give

$$\partial_4^{14}(2\nu\Delta^3) = 3u \cdot 2\nu \cdot \nu_2\bar{\kappa} = 3u \cdot \tau^8 \, \widetilde{d}\, \bar{\kappa}^3,$$

and similarly $\partial_4^{14}(2\nu\Delta^7) = 7u \cdot \tau^8 \, \widetilde{d}_4 \, \bar{\kappa}^3$. Combining these equalities with Proposition 3.37 gives the differentials supported in degrees $(75, 1)$ and $(171, 1)$.

For the other two, first recall the classical relation $\nu^3 = \eta\varepsilon$ in the non-synthetic sphere; see [Koc, Theorem 3.3.15 (a)]. This relation immediately lifts to the synthetic sphere $\mathbf{S}$ as there is no τ-power torsion in $\pi_{8,2}\,\mathbf{S}$, so we also have $\eta\varepsilon = \nu^3$ in Smf as well as $h_2^3 = h_1 c$ in Smf/τ. Using the extensions

$$\varepsilon\nu_2 = \tau^4\eta_1\kappa\bar{\kappa} \qquad \text{and} \qquad \eta\eta_1\kappa = \pm\tau^4 2\bar{\kappa}$$

of Lemma 9.39 and Lemma 9.48, respectively, we obtain the total differential

$$\partial_4^{14}(\nu^3\Delta^3) = 3u \cdot \eta\varepsilon \cdot \nu_2\bar{\kappa} = 3u \cdot (\tau^4\eta\eta_1\kappa\bar{\kappa})\bar{\kappa} = 3u \cdot \tau^4\bar{\kappa}^2(\eta\eta_1\kappa) = \pm 2\tau^8\bar{\kappa}^4.$$

Similarly, we have

$$\partial_4^{14}(\nu^3\Delta^7) = 7u \cdot \eta\varepsilon \cdot \nu_6\bar{\kappa} = 7u \cdot (\tau^4\eta\eta_5\kappa\bar{\kappa})\bar{\kappa} = 7u \cdot \tau^4\bar{\kappa}^2(\eta\eta_5\kappa) = \pm\tau^8\bar{\kappa}_4\bar{\kappa}^3$$

using Lemma 9.39 and Lemma 9.48. These two total differentials combined with Proposition 3.37 yield the remaining two atomic d_{13}'s. ∎

9.7.2 Meta-arguments

Proposition 9.54. *The condition of Proposition 9.4 holds for d_r for $13 \leqslant r \leqslant 21$. Moreover, Δ^8 is a d_{21}-cycle.*

Proof. The condition of Proposition 9.4 is checked directly as before. The class Δ^8 is a d_{21}-cycle for degree reasons. ∎

Proposition 9.55. *There are no line-crossing d_r-differentials for $13 \leqslant r \leqslant 21$.*

Proof. By Proposition 9.50, we may invoke the meta-argument of Proposition 9.5 for the case $r = 13$, which implies we only need to check for line-crossing differentials through the 46-stem, and there are no possibilities in this range.

For $15 \leqslant r \leqslant 21$, Proposition 9.54 implies we only need to check for line-crossing differentials through the 96-stem, and again there are no possibilities in this range. ∎

9.7.3 Lifts

Before we start the computation of the E_{23}-page, we need to lift a relation from Smf/τ^{12} to Smf/τ^{24}.

Lemma 9.56 (23,1). $\tau^4\nu\bar{\kappa} = 0$ *in* Smf/τ^{24}.

Proof. We know this relation holds in Smf/τ^{12} by Lemma 9.33, so it suffices to show that the reduction map

$$\pi_{23,1}\,\mathrm{Smf}/\tau^{24} \longrightarrow \pi_{23,1}\,\mathrm{Smf}/\tau^{12}$$

is injective. This now follows from Theorem 3.70; the nontrivial item to check is that there is a d_9 hitting the class in filtration 19, and this is a consequence of Section 9.5.2. ■

9.8 Page 23

9.8.1 Atomic differentials

Proposition 9.57 (121,1). $d_{23}(h_1\Delta^5) = g^6$.

Proof. First, we claim that by Theorem 3.62, it suffices to show that $\tau^{22}\bar{\kappa}^6 = 0$ in Smf/τ^{24}. Indeed, the element g^6 is a $d_{\leqslant 22}$-cycle and the only potential source of a differential to hit g^6 is $h_1\Delta^5$. Next, we claim that all of the proper sub-brackets of the four-fold Toda brackets

$$\langle \kappa,\, 2,\, \eta,\, \nu\rangle \qquad \text{and} \qquad \langle \tau^{16}\bar{\kappa}^4,\, \kappa,\, 2,\, \eta\rangle$$

are equal to zero in Smf/τ^{24}. This computation for the first bracket is straightforward. For the second bracket, we need to use the key fact that $\pi_{95,1}\,\mathrm{Smf}/\tau^{24} = 0$. This follows from Theorem 3.70, using the d_7 on Δ^4 of Proposition 9.29. Lastly, we have that $\tau^2\bar{\kappa} \in \langle \kappa, 2, \eta, \nu\rangle$ in Smf/τ^{24}, which follows from Corollary 8.35.

The fact that these sub-brackets are strictly zero allows us to apply the shuffling formula of Proposition B.12 (1), and we obtain

$$\tau^{22}\bar{\kappa}^6 = \tau^{16}\bar{\kappa}^4 \cdot (\tau^2\bar{\kappa}) \cdot \tau^4\bar{\kappa} \in \tau^{16}\bar{\kappa}^4 \cdot \langle \kappa, 2, \eta, \nu\rangle \cdot \tau^4\bar{\kappa} = \langle \tau^{16}\bar{\kappa}^4, \kappa, 2, \eta\rangle \cdot \tau^4\nu\bar{\kappa} = 0,$$

where for the last equality we used the relation $\tau^4\nu\bar{\kappa} = 0$ from Lemma 9.56. ■

All other d_{23}'s in the connective region follow from the Leibniz rule as g, $[\Delta^8]$, and $[h_1\Delta]$ are $d_{\leqslant 23}$-cycles.

9.8.2 Meta-arguments

Proposition 9.58. *The conditions of Proposition* 9.4 *hold for* d_{23}.

Proof. The condition is checked directly. ■

Proposition 9.59. *There are no line-crossing d_{23}-differentials.*

Proof. By Proposition 9.54, we may invoke the meta-argument of Proposition 9.5, which implies we only need to check for line crossing differentials through the 110-stem. There are no possible line-crossing d_{23}'s in this range. ■

Proposition 9.60. *The groups $\mathrm{E}_{24}^{n,s}$ vanish for $-21 < n < 0$ and all s, and also for all (n,s) with $n \geqslant 0$ and $s > 23$. In particular, there is no nontrivial differential of length > 23 in the DSS whose source lives in bidegree (n,s) for $n > -21$, and Δ^8 is a permanent cycle.*

Proof. The region S of Definition 9.7 on E_{24} is Δ^8-torsion free by Propositions 9.8 and 9.58. It therefore suffices to check the claim in the connective region, by multiplying any class with a power of Δ^8. However, by inductively applying Proposition 9.4, every element in the connective region on E_{24} of filtration $\geqslant 24$ is divisible by g^6, which is zero by Proposition 9.57. ■

9.9 Rainchecked four-fold Toda brackets in Smf

Logically speaking, this section appears just after we have finished with the E_{11}-page computations above and just before our discussion of the E_{13}-page. To establish our earlier d_{13}'s, we needed to establish some hidden extensions. This used that the element ν_2 (Notation 9.26) is contained in the bracket

$$\langle \nu,\, 2\nu\tau^4,\, \nu\tau^4,\, \bar{\kappa}^2 \rangle \subseteq \pi_{51,1}\,\mathrm{Smf}/\tau^{10},$$

which allowed us to determine the product $2\nu \cdot \nu_2$. We now compute this four-fold bracket. In the following discussion we work in the symmetric monoidal ∞-category of modules in Syn over Smf/τ^{10}. In particular, we write $\mathbf{1}$ for Smf/τ^{10}.

A simple long exact sequence argument shows there is a diagram as follows

$$\begin{tikzcd}
\mathbf{1}^{7,-7} \arrow[r, dashed, "\exists!\overline{\nu\tau^4}"] \arrow[dr, "\nu\tau^4"'] & C(2\nu\tau^4) \arrow[d] \\
& \mathbf{1}^{4,-4},
\end{tikzcd}$$

where the vertical map is projection onto the top cell, and similarly with the roles of ν and 2ν replaced. Taking the cofibers with respect to the dashed maps gives the unique forms of the 3-cell complexes $C(\nu\tau^4, 2\nu\tau^4)$ and $C(2\nu\tau^4, \nu\tau^4)$, in the sense of Definition B.3.

Proposition 9.61.

- (51,1) *The set*
$$\langle \nu,\, 2\nu\tau^4,\, \nu\tau^4,\, \bar{\kappa}^2 \rangle \subseteq \pi_{51,1}\,\mathrm{Smf}/\tau^{10}$$
is a singleton consisting of (*the mod* τ^{10} *reduction of*) *the element* ν_2 *from Notation* 9.26.

- (147,1) *The set*
$$\langle \nu_4,\, 2\nu\tau^4,\, \nu\tau^4,\, \bar{\kappa}^2 \rangle \subseteq \pi_{51,1}\,\mathrm{Smf}/\tau^{10}$$
is a singleton consisting of (*the mod* τ^{10} *reduction of*) *the element* ν_6 *from Notation* 9.26.

In a diagram, this proposition says the composition

$$\begin{tikzcd}
\boxed{48,0} \arrow[r, "\bar{\kappa}^2"] & \boxed{8,-8} \arrow[d, dash, "\nu\tau^4"] & \\
& \boxed{4,-4} \arrow[d, dash, "2\nu\tau^4"] & \\
& \boxed{0,0} \arrow[r, "\nu"] & \boxed{-3,-1}
\end{tikzcd}$$

can be constructed uniquely and that it is sent to $\nu\Delta^2$ in Smf/τ^4, which uniquely specifies $\nu_2 \in \pi_{51,1}\,\mathrm{Smf}/\tau^{10}$; see Lemma 9.25.

Proof. To construct $\underline{\nu}$, we must extend the map ν over $C(2\nu\tau^4, \nu\tau^4)$, which by Proposition B.11 exists if
$$0 \in \langle \nu, 2\nu\tau^4, \nu\tau^4 \rangle.$$
This bracket is nonempty since $2\nu \cdot \nu = 0$ in $\pi_{*,*}\,\mathbf{1}$, and a degree check shows it has zero indeterminacy. Shuffling shows it contains $\tau^8\langle \nu, 2\nu, \nu\rangle$, hence it suffices to show $0 \in \langle \nu, 2\nu, \nu \rangle$. This also has zero indeterminacy, and for degree reasons is either 0 or $\eta^2 c_4$. However, in the latter case, the shuffling formula
$$\langle \nu, 2\nu, \nu \rangle \eta = \nu \langle 2\nu, \nu, \eta \rangle = \nu\epsilon = 0$$
would give a contradiction, as $\eta^3 c_4 \neq 0$. Two extensions $\underline{\nu}$ so constructed differ by an element in $\pi_{11,-7}\,\mathbf{1} = 0$, so this constructs $\underline{\nu}$ uniquely.

We construct $\overline{\bar{\kappa}^2}$ analogously by constructing its dual. As before, we must show that
$$0 \in \langle 2\nu\tau^2, \nu\tau^4, \bar{\kappa}^2 \rangle.$$
This bracket is nonempty since $2\nu \cdot \nu = 0$ and $\nu\tau^4\bar{\kappa}^2 = 0$ by Corollary 8.35, and it lives in $\pi_{47,1}\,\mathbf{1}$, which is zero. A degree argument shows there is a unique extension

of $\bar{\kappa}^2$ over $C(\nu\tau^4)$, and two further extensions over $C(\nu\tau^4, 2\nu\tau^4)$ differ by an element in $\pi_{48,0}\,\mathbf{1} \neq 0$, via the exact sequence

$$\pi_{48,0}\,\mathbf{1} \xrightarrow{\iota} \pi_{48,0}\,C(\nu\tau^4, 2\nu\tau^4) \longrightarrow \pi_{48,0}\,\Sigma^{4,-4}C(2\nu\tau^4)$$

so $\overline{\bar{\kappa}^2}$ is not unique. However, the composite $\underline{\nu} \circ \iota = \nu$, and ν kills $\pi_{48,0}\,\mathbf{1}$, hence $\underline{\nu} \circ \overline{\bar{\kappa}^2}$ is unique, which proves the set

$$\langle \nu, 2\nu\tau^4, \nu\tau^4, \bar{\kappa}^2 \rangle$$

is a singleton.

To determine the image of $\langle \nu, 2\nu\tau^4, \nu\tau^4, \bar{\kappa}^2 \rangle$ under the map $\mathrm{Smf}/\tau^{10} \to \mathrm{Smf}/\tau^4$, we note that the primary attaching maps $2\nu\tau^4$ and $\nu\tau^4$ in $C(2\nu\tau^4, \nu\tau^4)$ are sent to zero in Smf/τ^4. Therefore the attaching map for the middle cell in $C(2\nu\tau^4, \nu\tau^4)$ is sent to zero, and the attaching map for the top cell in $C(2\nu\tau^4, \nu\tau^4)$ factors through a map in $\pi_{7,-7}\,\mathrm{Smf}/\tau^4 = 0$. Applying $- \otimes_{\mathbf{1}} \mathrm{Smf}/\tau^4$ to the composite $\underline{\nu} \circ \overline{\bar{\kappa}^2}$ therefore gives a diagram

$$\begin{tikzcd}
\boxed{48,0} \arrow[r, "\bar{\kappa}^2"] \arrow[dr, dashed, "f_2"] \arrow[ddr, dashed, "f_1"'] & \boxed{8,-8} & \\
& \boxed{4,-4} & \\
& \boxed{0,0} \arrow[r, "\nu"] & \boxed{-3,-1}
\end{tikzcd}$$

in modules over Smf/τ^4, using that the components of $\underline{\nu}$ on the middle and top cell are zero for degree reasons. The composite is therefore given by $\nu \cdot f_1$, and it remains to identify f_1 with $\Delta^2 \in \pi_{48,0}\,\mathrm{Smf}/\tau^4$ up to terms in the kernel of multiplication by ν.

We first construct a commutative square

$$\begin{tikzcd}
\Sigma^{4,-4}\,C(\nu\tau^4) \arrow[r] \arrow[d, "\underline{2\nu}"] & \mathbf{1}^{1,-1} \arrow[d, equal] \\
\mathbf{1}^{1,-5} \arrow[r, "\tau^4"] & \mathbf{1}^{1,-1}.
\end{tikzcd}$$

Here we choose the left-hand vertical map to be the unique extension of $2\nu\colon \mathbf{1}^{4,-4} \to$

$\mathbf{1}^{1,-5}$ over $C(\nu\tau^4)$. Taking fibers produces a commutative square

$$\begin{tikzcd}
C(2\nu\tau^4, \nu\tau^4) \arrow[r] \arrow[d] & \Sigma^{4,-4} C(\nu\tau^4) \arrow[d] \\
C(\tau^4) \arrow[r] & \mathbf{1}^{1,-5}
\end{tikzcd} \tag{9.62}$$

where each of the horizontal maps crush the bottom cell. Applying the counterclockwise composite in this diagram to $\overline{\kappa}^2$ gives a composition that we may represent with the following diagram.

$$\begin{tikzcd}
\boxed{48,0} \arrow[r, "\overline{\kappa}^2"] & \boxed{8,-8} \arrow[d, dash, "\nu\tau^4"] & & \\
& \boxed{4,-4} \arrow[r, "2\nu"] \arrow[d, dash, "2\nu\tau^4"] & \boxed{1,-5} \arrow[r, equal] \arrow[d, dash, "\tau^4"] & \boxed{1,-5} \\
& \boxed{0,0} \arrow[r, equal] & \boxed{0,0} &
\end{tikzcd}$$

After tensoring down to Smf/τ^4, the composite of the first two morphisms above has bottom component equal to f_1 and the diagram shows that $\partial_4^8(f_1) = 2\nu \cdot f_2$.

To determine $2\nu \cdot f_2$, we apply the clockwise composite in (9.62) to $\overline{\kappa}^2$, which we may represent with the following diagram.

$$\begin{tikzcd}
\boxed{48,0} \arrow[r, "\overline{\kappa}^2"] & \boxed{8,-8} \arrow[r, equal] \arrow[d, dash, "\nu\tau^4"] & \boxed{8,-8} \arrow[d, dash, "\nu\tau^4"] & \\
& \boxed{4,-4} \arrow[r, equal] \arrow[d, dash, "2\nu\tau^4"] & \boxed{4,-4} \arrow[r, "2\nu"] & \boxed{1,-5} \\
& \boxed{0,0} & &
\end{tikzcd}$$

After tensoring down to Smf/τ^4, the composite φ of the first two morphisms above has bottom component equal to f_2, so we analyze φ carefully. One may construct a diagram

$$\begin{tikzcd}
\boxed{48,0} \arrow[r, "\bar{\kappa}^2"] & \boxed{8,-8} \arrow[r, equal] \arrow[d, dash, "\nu\tau^4"] & \boxed{8,-8} \arrow[r, "\nu"] \arrow[d, dash, "\nu\tau^4"] & \boxed{5,-9} \arrow[d, dash, "\tau^4"] \\
& \boxed{4,-4} \arrow[r, equal] \arrow[d, dash, "2\nu\tau^4"] & \boxed{4,-4} \arrow[r, equal] & \boxed{4,-4} \\
& \boxed{0,0} & &
\end{tikzcd}$$

where the composite of the first two maps is φ. After tensoring down to Smf/τ^4, the component of this composite onto the cell of dimension $(4,-4)$ remains f_2, so this diagram shows that $\partial_4^8(f_2) = \bar{\kappa}^2\nu$, which guarantees that $f_2 = \bar{\kappa}\Delta$.

Putting these facts together we conclude that $\partial_4^8(f_1) = 2\nu\bar{\kappa}\Delta$, which guarantees that $f_1 = \Delta^2 \in \pi_{*,*}\,\mathrm{Smf}/\tau^4$ modulo ν-torsion, and we conclude that $\nu \cdot f_1 = \nu_2$.

The bracket for ν_6 follows from analogous arguments. ■

9.10 Computations away from the prime 2

Above we computed (most of) the signature spectral sequence of $\mathrm{Smf}_{(2)}$, and here we would like to do the same for $\mathrm{Smf}_{(3)}$ and $\mathrm{Smf}[\frac{1}{6}]$. All together, these three results yield the signature spectral sequence for Smf.

For any collection of primes J, the signature spectral sequence associated to $\mathrm{Smf}[J^{-1}]$ is naturally identified with the DSS for $\mathrm{Tmf}[J^{-1}]$. Indeed, by Proposition 7.56, the natural map of synthetic $\mathbf{E}_\infty$-rings

$$\mathrm{Smf}[J^{-1}] = \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}})[J^{-1}] \xrightarrow{\cong} \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}[J^{-1}])$$

is an isomorphism. This is implicitly used below.

9.10.1 Computations at the prime 3

As is often the case for Tmf, the 3-local DSS computation is a vast simplification of the 2-local analogue discussed previously. We implicitly work 3-locally in this subsection.

Theorem 9.63. *The signature spectral sequence of* Smf*, i.e., the DSS for* Tmf*, is determined below, and has precisely the form depicted in Figure* A.1*.*

The E_2-page is computed using sheaf cohomology; see Figure A.1 or [Kon12, Figure 10]. As before, we recommend that the reader keeps these charts nearby throughout the following arguments.

There are only two atomic differentials.

Proposition 9.64.

- (24,0) $d_5(\Delta) = \pm\alpha\beta^2$.
- (51,1) $d_9(\alpha\Delta^2) = \pm\beta^5$.

Proof. Recall from Proposition 8.27 that $\alpha \in \pi_{3,1}\,\mathbf{S}$ has nonzero image in $\pi_{3,1}\,\mathrm{Smf}$. The class in the sphere $\beta \in \pi_{10,2}\,\mathbf{S}$ is defined as the Toda bracket $\langle\alpha,\alpha,\alpha\rangle$. From the Massey product structure of the 3-local cubic Hopf algebroid, combined with Proposition B.14 and Corollary 8.15, we see that β hits the generator of $\pi_{10,2}\,\mathrm{Smf}/\tau \cong \mathbf{F}_3$, which we also call β; see [Bau08, Equation (5.1)].

In the ANSS for $\mathbf{S}$, there is the classical Toda differential $d_5(\beta_{3/3}) = \pm\alpha\beta^3$; see [Rav04, Theorem 4.4.22]. As a result, the element $\alpha\beta^3$ in Smf/τ must also be hit by a differential. For degree reasons, the only possibility is $d_5(\pm\beta\Delta) = \alpha\beta^3$. From the Leibniz rule and the fact that β is a permanent cycle, we obtain $d_5(\Delta) = \pm\alpha\beta^2$. The Leibniz rule gives all other d_5's in this spectral sequence.

For degree reasons, the next possible differential is a d_9. To compute this atomic d_9, we show that the relation

$$\tau^4\alpha\beta^2 = 0 \tag{9.65}$$

holds in Smf/τ^{14}. As the mod τ reduction $\alpha\beta^2 \in \pi_{23,5}\,\mathrm{Smf}/\tau$ is the target of a d_5, Theorem 4.77 tells us that there exists a τ^4-torsion lift to $\pi_{23,5}\,\mathrm{Smf}/\tau^{14}$. Using Corollary 3.72, we see that this lift is unique, proving that (9.65) holds.

Next, recall that the defining Toda bracket expression $\beta = \langle\alpha,\alpha,\alpha\rangle$ also holds in Smf/τ^{14}. Applying the usual juggling formula Proposition B.12 and the relation (9.65), we find that

$$\tau^8\beta^5 = \tau^4\beta^2\langle\alpha,\alpha,\alpha\rangle\tau^4\beta^2 = \langle\tau^4\beta^2,\alpha,\alpha\rangle\tau^4\alpha\beta^2 = 0.$$

(Note that (9.65) also justifies that the second bracket is nonempty.) In other words, we have learned that β^5 is τ^8-torsion, which by Theorem 3.67 means that its mod τ reduction $\beta^5 \in \pi_{50,10}\,\mathrm{Smf}/\tau$ is hit by a $d_{\leqslant 9}$-differential. As this class is a $d_{\leqslant 8}$-cycle for degree reasons, it must be hit by d_9. The only possibility is the desired $d_9(\alpha\Delta^2) = \pm\beta^5$. All other d_9's follow from the Leibniz rule. ∎

This spectral sequence then collapses with a horizontal vanishing line at $s = 8$. This yields the homotopy groups of Tmf, which can be read off from Figure A.1. In other words, we have proved Theorem 9.63.

9.10.2 Computations away from 6

Theorem 9.66. *There is an isomorphism of bigraded $\mathbf{Z}[\frac{1}{6},\tau]$-modules*

$$\pi_{*,*}\operatorname{Smf}[\tfrac{1}{6}] \cong \mathbf{Z}[\tfrac{1}{6},\tau][c_4,\Delta]\otimes E(c_6)\oplus \mathbf{Z}[\tfrac{1}{6},\tau]\Big\{\, c_4^i c_6^j \Delta^k \,\Big\}_{\substack{i,k\leqslant -1\\ 0\leqslant j\leqslant 1}}$$

where $E(-)$ denotes exterior algebra and

$$|c_4^i c_6^j \Delta^k| = \begin{cases} (8i+12j+24k,\ 0) & \text{if } i\geqslant 0,\\ (8i+12j+24k-1,\ 1) & \text{if } i\leqslant -1.\end{cases}$$

Moreover, in nonnegative degrees, this is an isomorphism of rings.

This $\mathbf{Z}[\tau]$-module structure on the synthetic homotopy groups shows that the signature spectral sequence of $\operatorname{Smf}[\frac{1}{6}]$ converging to $\operatorname{Tmf}[\frac{1}{6}]$ collapses on the E_2-page.

Proof. The signature of $\operatorname{Smf}[\frac{1}{6}]$ is given by the DSS for $\mathcal{O}^{\mathrm{top}}$ on $\overline{\mathfrak{M}}_{\mathrm{ell}}[\frac{1}{6}]$. This spectral sequence takes the form

$$\mathrm{E}_2^{n,s} \cong \mathrm{H}^s(\overline{\mathfrak{M}}_{\mathrm{ell}}[\tfrac{1}{6}],\, \omega^{\otimes(n+s)/2}) \implies \pi_n \operatorname{Tmf}[\tfrac{1}{6}].$$

As $\overline{\mathfrak{M}}_{\mathrm{ell}}[\frac{1}{6}]$ is the weighted projective stack $\mathbf{P}(4,6)$, we see that this spectral sequence is concentrated in filtrations $s=0,1$; see [Kon12, Section 6]. In particular, this spectral sequence collapses for degree reasons. By the Omnibus Theorem 4.77, it follows that $\pi_{*,*}\operatorname{Smf}[\frac{1}{6}]$ is τ-power torsion free, so the claim follows from the computation of the second page of the DSS. ∎

Chapter 10

Main results

With all of the computations out of the way, we can now prove our main theorems and corollaries, including the Gap Theorem (Theorem A) and a description of the DSS for Tmf (Theorem B). This fills the gaps in the literature discussed in Section 6.1. As corollaries, we also compute the homotopy groups of Tmf (Corollary C), the ANSS for tmf (Corollary D), and the ANSS for TMF (Corollary E). For the convenience of the reader, we repeat the statements of these results below.

Of particular interest is our method of solving extension problems (see also Remark 10.4), using the technique explained in Section 3.3.1, using our previous computations of total differentials. This in particular avoids the use of large Toda brackets.

We start with the Gap Theorem.

Theorem A. *The homotopy groups* $\pi_n \operatorname{Tmf}$ *vanish for* $-21 < n < 0$.

Proof. Using the usual fracture square for Tmf, it suffices to prove the Gap Theorem for $\operatorname{Tmf}_{(p)}$ for each prime p. By Theorem 9.66, this holds for all primes $p \neq 2, 3$. By Theorem 9.63 (see also Figure A.1), it holds at the prime 3. Finally, by Proposition 9.60, it holds at the prime 2. ∎

Our proof of the Gap Theorem only requires the DSS for $\operatorname{Tmf}_{(2)}$ only in stems $n \geqslant -20$. The additional computations of [CDvN25, Section 6.9] yield the entire DSS for Tmf.

Theorem B. *The descent spectral sequence for* Tmf *takes the form depicted in Figures* A.2 *to* A.6 *at the prime* 2, *depicted in Figure* A.1 *at the prime* 3, *and collapses otherwise as detailed in Theorem* 9.66.

Proof. By Proposition 7.56, the natural map of synthetic $\mathbf{E}_\infty$-rings

$$\operatorname{Smf}[J^{-1}] = \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}})[J^{-1}] \xrightarrow{\simeq} \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}} \times \operatorname{Spec} \mathbf{Z}[J^{-1}])$$

is an isomorphism for any set of primes J. In particular, we obtain the DSS for Tmf from the DSS for $\mathrm{Tmf}[\frac{1}{6}]$ of Theorem 9.66, which collapses, the DSS for $\mathrm{Tmf}_{(3)}$ of Theorem 9.63, and the DSS for $\mathrm{Tmf}_{(2)}$ of Chapter 9. Our computations show that these spectral sequences collapse on a finite page, so that the conditional convergence is in fact strong. ■

We deduce the following, which will be helpful for later results.

Lemma 10.1.

- *The element* $\Delta \in \pi_{24,0}\,\mathrm{Smf}[\frac{1}{6}]/\tau$ *lifts uniquely to* $\pi_{24,0}\,\mathrm{Smf}[\frac{1}{6}]$.
- *The element* $\Delta^3 \in \pi_{72,0}\,\mathrm{Smf}_{(3)}/\tau$ *lifts uniquely to* $\pi_{72,0}\,\mathrm{Smf}_{(3)}$.
- *The element* $\Delta^8 \in \pi_{192,0}\,\mathrm{Smf}_{(2)}/\tau$ *lifts uniquely to* $\pi_{192,0}\,\mathrm{Smf}_{(2)}$.
- *The element* $\Delta^{24} \in \pi_{576,0}\,\mathrm{Smf}/\tau$ *lifts uniquely to* $\pi_{576,0}\,\mathrm{Smf}$.

Proof. Our computation shows that, at every prime, the DSS collapses at a finite page. As a result, the conditional convergence of the DSS is strong. In particular, the Omnibus Theorem 4.77 applies. Using this, it suffices to show that the indicated powers of Δ are permanent cycles in the localised DSS for Tmf, and that there are no permanent cycles in higher filtrations. This now follows from the computations of the previous chapter. Indeed, when 6 is inverted, the class Δ is a permanent cycle; 3-locally, the power Δ^3 is a permanent cycle; 2-locally, the power Δ^8 is a permanent cycle. In all of these cases, there are indeed no elements in higher filtrations in the respective stems. The lowest common multiple of these powers is Δ^{24}, proving the final claim. ■

With the whole of the DSS for Tmf at hand, we have almost computed the homotopy groups of Tmf; all that is left is to solve some extension problems. Most of these extension problems follow from rudimentary algebraic arguments given our computations so far. The lone exception is a 2-extension in degree 110, which instead follows from our knowledge of total differentials using the method explained in Remark 3.41. The application of this method to this hidden extension is due to [Isa+24, Proposition 4.5].

Corollary C. *The homotopy groups of* Tmf, *and hence also those of* $\mathrm{tmf} = \tau_{\geqslant 0}\mathrm{Tmf}$, *are computed; see Theorem* 9.66 *away from* 6, *Figure* A.1 *at the prime* 3, *and Figures* A.3 *to* A.6 *at the prime* 2.

Proof. Away from 6, these homotopy groups are obtained from Theorem 9.66 by inverting τ. Localised at the prime 3, they follow immediately from Theorem 9.63 as there are no extension problems. At the prime 2, we use the DSS of Appendix A computed in Chapter 9, but there are some extension problems to solve.

First, let us deal with the positive stems. The 2-extensions which follow from the hidden extension $4\nu = \tau^2\eta^3$ from Lemma 9.15 in $\mathrm{Smf}_{(2)}/\tau^4$ are clear, so we ignore these. There can be no extensions between the ko-patterns, indicated by solid diamonds on the 0-line of Figures A.3 to A.6; this can be checked on a case by case basis. Indeed, there cannot be any 2-extensions with source in filtration 0 for algebraic reasons, as these sources are torsion-free. All other parts of the ko-patterns are divisible by η, which cannot support multiplication by 2. Similar arguments discount many other potential extensions: there cannot be a 2-extension in stem 65 from filtration 3 to 9 as the source is divisible by the 2-torsion class κ. Using the lift of Δ^8 from Lemma 10.1, the above arguments reduce us to verifying 2-extensions in the following stems:

$$54, \quad 110, \quad 130, \quad 150.$$

The extension in the 110-stem implies the one in the 130-stem by $\bar{\kappa}$-multiplication. In the 54-stem, this extension is precisely Lemma 9.52. This also gives the lower half of the extensions in the 150-stem, and the upper half follows from the ones in the 130-stem by $\bar{\kappa}$-multiplication. We are reduced to the extension in degree 110, where we want to show that

$$2 \cdot \kappa_4 = \eta_1^2\bar{\kappa}^3 \qquad \text{in } \pi_{110}\,\mathrm{Tmf}_{(2)}. \tag{10.2}$$

Here we write κ_4 for a choice of element denoted by $\widetilde{d_4}$ in Lemma 9.52. In this case, consider the exact sequence

$$\pi_{*,*}\,\mathrm{Smf}_{(2)}/\tau^{28} \longrightarrow \pi_{*,*}\,\mathrm{Smf}_{(2)}/\tau^4 \xrightarrow{\partial_4^{28}} \pi_{*-1,\,*+4}\,\mathrm{Smf}_{(2)}/\tau^{28} \xrightarrow{\tau^4} \pi_{*-1,\,*}\,\mathrm{Smf}_{(2)}/\tau^{28}.$$

We will show that there is an equality

$$2\,\tau^8\,\kappa_4\,\bar{\kappa}^3 = \tau^{20}\,\eta_1^2\,\bar{\kappa}^6 \qquad \text{in } \pi_{170,6}\,\mathrm{Smf}_{(2)}/\tau^{28}, \tag{10.3}$$

where all of the elements displayed are their unique lifts to Smf/τ^{28}; from this our desired extension in $\mathrm{Tmf}_{(2)}$ in the 110-stem follows as the DSS for $\mathrm{Tmf}_{(2)}$ collapses on the E_{26}-page. Note that this is indeed enough to deduce the hidden extension (10.2). The differentials

$$d_{13}(2h_2\Delta^7) = dg^3\Delta^4 \qquad \text{and} \qquad d_{23}(h_1^3\Delta^7) = h_1^2g^6\Delta^2$$

of Sections 9.7 and 9.8, respectively, show by Theorem 3.62 (3) that $\tau^8\,\kappa_4\,\bar{\kappa}$ and $\tau^{20}\,\eta_1^2\,\bar{\kappa}^6$ are both τ^4-torsion (as they are the unique lift of the respective E_2-elements). One then uses the exact sequence above to compute

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

In Smf/τ^4 we have $\tau^2\eta^3 = 4\nu$ courtesy of the mod τ^4 reduction of Lemma 9.15, which also yields $4\nu\Delta^7 = \tau^2\eta^3\Delta^7$. Combining what we have so far yields the desired equality (10.3):

$$2\tau^8\,\kappa_4\,\bar{\kappa}^3 = \partial_4^{28}(4\nu\Delta^7) = \partial_4^{28}(\tau^2\eta^3\Delta^7) = \tau^{20}\,\eta_1^2\,\bar{\kappa}^5.$$

We are reduced to extension problems in negative degrees. Again, there are no extension problems away from 2, so we are reduced to $\mathrm{Tmf}_{(2)}$. In this case, all of the extension problems in negative degrees follow from their counterparts in positive degrees by Δ^{8t}-multiplication for large enough t. ∎

Remark 10.4 (Further hidden extensions). The method from Remark 3.41 used to deduce the 2-extension in stem 110 generalises to capture many hidden extensions in the DSS for Tmf. For example, consider the total differentials

$$\partial_4^{28}(\eta\Delta^5) = \tau^{18}\bar{\kappa}^6 \qquad \text{and} \qquad \partial_4^{28}(\Delta^5) = u\nu_4\bar{\kappa}$$

in $\mathrm{Smf}_{(2)}/\tau^{28}$, where u is a unit of $\mathbf{Z}/8$; these can be computed from the computations of the previous chapter. From this, one obtains

$$\eta\nu_4\bar{\kappa} = \eta \cdot \partial_4^{28}(\Delta^5) = \partial_4^{28}(\eta\Delta^5) = \tau^{18}\bar{\kappa}^6,$$

and hence the hidden η-extension

$$\eta\nu_4 = \bar{\kappa}^5 \qquad \text{in } \pi_{100}\,\mathrm{Tmf}_{(2)}$$

of [Bau08, Corollary 8.7 (2)]. These arguments completely avoid the use of six-fold Toda brackets seen in [Bau08].

Using the Gap Theorem, Mathew computed the Hopf algebroid computing the E_2-page of the ANSS of tmf, which Bauer used to compute the E_2-page. We will now show how this, combined with our computation of the DSS for Tmf, computes the ANSS of tmf. This recovers Bauer's differentials of [Bau08] without any circularity issues.

Corollary D. *There is an inclusion of the ANSS for* tmf *into the DSS for* Tmf *as a retract of spectral sequences. In particular, the ANSS for* tmf *is the region under the blue line of Figures* A.3 *to* A.6 *at the prime* 2 *from the* E_5*-page, the region under the blue line of Figure* A.1 *and* 3*, and the connective part of Theorem* 9.66 *away from* 6*.*

Proof. By [Mat16, Corollary 5.3], which in turn relies on the Gap Theorem, the E_2-page of ANSS for tmf is isomorphic to the cohomology of the cubic Hopf algebroid of Section 8.3. By Corollary 8.15, this means that the composition of natural maps of synthetic $\mathbf{E}_\infty$-rings

$$\nu\,\mathrm{tmf} \longrightarrow \nu\,\mathrm{Tmf} \longrightarrow \mathrm{Smf}$$

induces an isomorphism on $\pi_{n,s}(C\tau \otimes -)$ for $5s \leqslant n + 12$, and in general is a retract of bigraded abelian groups. At $p = 2$, the atomic d_3 in the signature spectral sequence for Smf lifts uniquely to the signature spectral sequence for $\nu\,\mathrm{tmf}$. This propagates to all other d_3's using the Leibniz rule. From the E_5-page on, there are no differentials whose source is in the connective region and whose target is outside the connective region (i.e., there are no line-crossing differentials à la

Section 9.1.1). Therefore, we can safely import the differentials in the connective region of the signature spectral sequence of $\mathrm{Smf}_{(2)}$ to differentials in the signature spectral sequence of $\nu\,\mathrm{tmf}_{(2)}$.

The same is true at $p = 3$, i.e., there are no line-crossing differentials. Away from 6 this is tautological as $\mathrm{tmf}[\frac{1}{6}]$ is complex-oriented. It follows that the map $\nu\,\mathrm{tmf} \to \mathrm{Smf}$ induces a retract of spectral sequences. ■

The DSS for TMF follows easily from Theorem B by inverting an appropriate power of Δ. In fact, this inversion happens in a very structured way, at the level of synthetic modular forms. We promised this result back in Section 8.2.

Proposition 10.5. *The natural map of synthetic* $\mathbf{E}_\infty$-Smf-*algebras* $\mathrm{Smf}[\Delta^{-24}] \xrightarrow{\cong} \mathrm{SMF}$ *is an isomorphism.*

Proof. Using Lemma 10.1, this follows from Proposition 7.56 and the fact that $\mathfrak{M}_{\mathrm{ell}}$ is exactly the nonvanishing locus of $\overline{\Delta}^{24}$ inside $\overline{\mathfrak{M}}_{\mathrm{ell}}$. ■

Remark 10.6. Note that the proof of the above does not depend on the isomorphism $\mathrm{SMF} \cong \nu\,\mathrm{TMF}$, but only uses that SMF is a synthetic lift of TMF, and that its signature is the DSS for TMF.

Using this, the computation of the DSS for TMF now follows.

Corollary E. *The ANSS for* TMF *is obtained from the DSS for* Tmf *by inverting* Δ^{24}*. Specifically, at the prime* 2 *it is obtained by inverting* Δ^8 *in Figures* A.2 *to* A.6*, at the prime* 3 *by inverting* Δ^3 *in Figure* A.1*, and away from* 6 *by inverting* Δ *in Theorem* 9.66*.*

Proof. The signature spectral sequence of SMF is the DSS for TMF. Proposition 10.5 therefore identifies the E_2-page of the DSS for TMF with the Δ-inversion of the E_2-page of the DSS for Tmf. The $\mathbf{Z}[\Delta^{\pm}]$-module $\pi_{*,*}\,\mathrm{SMF}/\tau$ is generated by the image of the connective region in $\pi_{*,*}\,\mathrm{Smf}/\tau$. This means that by inverting Δ on the DSS for Tmf, we obtain all of the differentials in the DSS for TMF. The ANSS for TMF then follows from the isomorphism of synthetic $\mathbf{E}_\infty$-rings

$$\nu\,\mathrm{TMF} \cong \mathrm{SMF}$$

from Proposition 8.6, as $\nu\,\mathrm{TMF}$ implements the ANSS for TMF; see Theorem 4.71. ■

As a final corollary of the computation, we have the following alternate expression for the synthetic $\mathbf{E}_\infty$-ring Smf. This definition was originally proposed by Jack Davies. Recall that $\mathrm{tmf} = \tau_{\geqslant 0}\mathrm{Tmf}$. From our computation of Smf, it follows that we have an element $c_4 \in \pi_{8,0}\,\mathrm{Smf}$ defined uniquely by being the $\overline{\kappa}$-torsion class detecting the normalised Eisenstein series of weight 4. As a consequence of Corollary D, it follows that this defines an element $c_4 \in \pi_{8,0}(\nu\,\mathrm{tmf})$, and in particular an element in $\pi_8\,\mathrm{tmf}$.

Proposition 10.7. *There is a pullback diagram of synthetic $\mathbf{E}_\infty$-rings*

$$\begin{tikzcd} \mathrm{Smf} \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & \nu\,\mathrm{TMF} \arrow[d] \\ \nu\,\mathrm{tmf}[c_4^{-1}] \arrow[r] & \nu\,\mathrm{tmf}[c_4^{-1}, \Delta^{-24}]. \end{tikzcd}$$

Proof. The diagram of synthetic $\mathbf{E}_\infty$-rings

$$\begin{tikzcd} \mathrm{Smf} \arrow[r] \arrow[d] & \mathrm{Smf}[\Delta^{-24}] \arrow[d] \\ \mathrm{Smf}[c_4^{-1}] \arrow[r] & \mathrm{Smf}[c_4^{-1},\, \Delta^{-24}] \end{tikzcd}$$

is a pullback square. Indeed, by Proposition 7.56, we can identify this square with

$$\begin{tikzcd} \mathcal{O}^{\mathrm{syn}}(D(1)) \arrow[r] \arrow[d] & \mathcal{O}^{\mathrm{syn}}(D(\Delta)) \arrow[d] \\ \mathcal{O}^{\mathrm{syn}}(D(c_4)) \arrow[r] & \mathcal{O}^{\mathrm{syn}}(D(c_4\Delta)). \end{tikzcd}$$

The fact that $\overline{\mathfrak{M}}_{\mathrm{ell}} = D(\Delta) \cup D(c_4)$ and the fact that $\mathcal{O}^{\mathrm{syn}}$ is a Zariski sheaf then yields the above pullback of synthetic $\mathbf{E}_\infty$-rings. We are left to identify this square with the square in the statement of the proposition. Combining Propositions 8.6 and 10.5, we obtain isomorphisms $\mathrm{Smf}[\Delta^{-24}] \cong \mathrm{SMF} \cong \nu\,\mathrm{TMF}$. It suffices to show that the composite morphism of synthetic $\mathbf{E}_\infty$-rings

$$\nu\,\mathrm{tmf} \longrightarrow \nu\,\mathrm{Tmf} \longrightarrow \mathrm{Smf}$$

becomes an isomorphism after inverting c_4. As both $\nu\,\mathrm{tmf}$ and Smf are τ-complete (by Theorem 4.71 and Proposition 8.2, respectively), it suffices to show that it becomes an isomorphism after tensoring with $C\tau$. After tensoring with $C\tau$, we note that by [Mat16, Corollary 5.3] and Corollary 7.33, we can identify the bigraded homotopy groups of $C\tau \otimes \nu\,\mathrm{tmf}$ and $C\tau \otimes \mathrm{Smf}$ as the E_2-page of a DSS for $\mathfrak{M}_{\mathrm{cub}}$ and $\overline{\mathfrak{M}}_{\mathrm{ell}}$, respectively. As $\overline{\mathfrak{M}}_{\mathrm{ell}}$ can be written inside $\mathfrak{M}_{\mathrm{cub}}$ as the union of the open substacks $D(c_4) \cup D(\Delta)$, and filtered colimits commute with cohomology (see [Stacks, Tag 0GQV]), we obtain the desired isomorphism. ∎

Remark 10.8. Note that the previous result requires a computational input, namely that the classes c_4 and Δ^{24} lift to $\nu\,\mathrm{tmf}$. Both of these results are nontrivial: showing that c_4 is permanent involves computing d_{11}-differentials (see the proof of Proposition 9.51), and the case of Δ is one of our meta-arguments we had to check at every page (Proposition 9.6). As such, defining Smf in this way cannot be used as an alternative definition to do the computation of Chapter 9 with.

Chapter 11

Epilogue

In this chapter, we describe a linear path to the computation of topological modular forms. Our aim is not to give a historical account, but rather to explain the logical dependencies of the various parts of the literature.

The construction of the sheaf $\mathcal{O}^{\mathrm{top}}$ of $\mathbf{E}_\infty$-rings on $\overline{\mathfrak{M}}_{\mathrm{ell}}$ is the starting point of the theory. This was done by Goers–Hopkins–Miller, and later Lurie for its restriction to $\mathfrak{M}_{\mathrm{ell}}$; see [DFHH, Chapter 12], [Goe10], and [Dav24b]. With this sheaf, one can define the following incarnations of topological modular forms:

- the **projective topological modular forms**

$$\mathrm{Tmf} = \mathcal{O}^{\mathrm{top}}(\overline{\mathfrak{M}}_{\mathrm{ell}});$$

- the **periodic topological modular forms**

$$\mathrm{TMF} = \mathcal{O}^{\mathrm{top}}(\mathfrak{M}_{\mathrm{ell}});$$

- the **connective topological modular forms**

$$\mathrm{tmf} = \tau_{\geqslant 0}\,\mathrm{Tmf}.$$

However, the reason for defining connective tmf in this way, or for calling TMF periodic, are not clear at this point: knowing that Tmf is not connective, or that TMF is 576-periodic, requires a computation of their homotopy groups, which is a very nontrivial undertaking. Moreover, the ad-hoc definition of tmf means that tmf does not share the close connection to algebraic geometry that Tmf and TMF have.

The computation of the homotopy groups of these three variants of topological modular forms requires a spectral sequence. The descent spectral sequence (DSS) for Tmf is the fundamental one. As argued previously in this thesis, the key to computing this is to define the MU-synthetic lift $\mathcal{O}^{\mathrm{syn}}$, and define

$$\mathrm{Smf} = \mathcal{O}^{\mathrm{syn}}(\overline{\mathfrak{M}}_{\mathrm{ell}}).$$

This definition also does not require any computational input.

The computation of the second page of the DSS for Tmf was done, as detailed in Section 8.3, by Konter [Kon12], using Bauer's [Bau08] computation of the cohomology of $\mathfrak{M}_{\mathrm{cub}}$. Again note that this only uses the algebraic geometry of $\overline{\mathfrak{M}}_{\mathrm{ell}}$ and $\mathfrak{M}_{\mathrm{cub}}$, without their connection to topology. Next, the structure on Smf of a synthetic $\mathbf{E}_\infty$-ring results in a map of MU-synthetic $\mathbf{E}_\infty$-rings

$$\nu\mathbf{S} \longrightarrow \mathrm{Smf}.$$

Granting a certain amount of knowledge of the ANSS of the sphere in low dimensions (which does not require topological modular forms) allows one to then show that this map detects these elements, as done in Section 8.5.

At this point, the computation of the DSS of Tmf is essentially a matter of one's ability to deduce differentials. As demonstrated in detail in Chapter 9, the filtered techniques developed in Part I, together with synthetic Moss's Theorem of Theorem B.15, give a sufficient supply of this ability, allowing for the complete calculation of the homotopy groups of Tmf. In particular, this proves the Gap Theorem for Tmf, and computes the homotopy groups of tmf as well; see Theorem A and Corollary C in Chapter 10.

What this does not yet do is compute the second page of either the ANSS or the ASS for tmf. In the case of the ANSS, only the computation of the E_2-page is missing: once this is known, the differentials in this spectral sequence can be directly imported from the DSS for Tmf; see Corollary D in Chapter 10. As shown by Mathew [Mat16], to compute the MU- and $\mathbf{F}_p$-homologies of tmf, one needs the Gap Theorem for Tmf.

In some more detail, to compute $\mathrm{MU}_*\mathrm{tmf}$, one starts by computing $\mathrm{MU}_*\mathrm{Tmf}$. For this, one needs to know that the natural map

$$\mathrm{MU} \otimes \mathrm{Tmf} = \mathrm{MU} \otimes \Gamma(\mathcal{O}^{\mathrm{top}}) \longrightarrow \Gamma(\mathrm{MU} \otimes \mathcal{O}^{\mathrm{top}})$$

is an isomorphism. This either follows using the finite complexes defined by Mathew in [Mat16], or using the affineness results of Mathew–Meier [MM15] (see Section 7.4.1 for a detailed discussion of this second approach). As a result, the homotopy of $\mathrm{MU} \otimes \mathrm{Tmf}$ can be computed by a DSS of its own, which turns out to collapse immediately; see [Mat16, Proposition 5.1]. Using the Gap Theorem, Mathew deduces $\mathrm{MU}_*\mathrm{tmf}$ as an $\mathrm{MU}_*\mathrm{MU}$-comodule from this, and proves that the E_2-page of the ANSS for tmf is the cohomology of the stack $\mathfrak{M}_{\mathrm{cub}}$; see Corollaries 5.2 and 5.3 of op. cit., respectively.

Remark 11.1. In Sections 7.4 and 7.5, we used the affineness results of Mathew–Meier to study $\mathcal{O}^{\mathrm{syn}}$, for instance to study the comparison map

$$\nu(\mathcal{O}^{\mathrm{top}}(\mathfrak{X})) \longrightarrow \mathcal{O}^{\mathrm{syn}}(\mathfrak{X}).$$

Strictly speaking, this is not necessary if one is only interested in a geodesic path to the computation of Smf. On the other hand, it is not circular to use these results before the computation of π_* Tmf is done, as Mathew–Meier do not use computational input for their affineness results. However, certain *applications* of their results, such as the isomorphisms

$$\mathrm{TMF} \cong \mathrm{Tmf}[\Delta^{-24}] \qquad \text{and} \qquad \mathrm{SMF} \cong \mathrm{Smf}[\Delta^{-24}]$$

do require computational knowledge, namely the fact that Δ^{24} is a permanent cycle.

Knowing the MU-homology of tmf, Mathew [Mat16, Section 5.3] computes the $\mathbf{F}_2$-homology, resulting in the famous expression

$$\mathrm{H}_*(\mathrm{tmf}; \mathbf{F}_2) \cong \mathcal{A}_* /\!\!/ \mathcal{A}_*(2).$$

This is taken as the defining property of the $\mathbf{E}_\infty$-ring tmf by Bruner–Rognes [BR]. From this defining feature, they compute the entirety of the ASS for tmf. (Note, by contrast, that the $\mathbf{F}_p$-homologies of both Tmf and TMF vanish, for all p.) Marek [Mar24] uses the outcome of the computation of Bruner–Rognes to describe the $\mathbf{F}_2$-synthetic analogue of tmf, computing $\pi_{*,*}(\nu_{\mathbf{F}_2}\mathrm{tmf})^\wedge_2$.

Recently, Isaksen–Kong–Li–Ruan–Zhu [Isa+24] use the second pages of both the $\mathbf{F}_2$-Adams and Adams–Novikov spectral sequences of tmf to compute ν_{MU} tmf (referred to as mmf in op. cit.; see Example 5.43 for a comparison). More specifically, they analyse both the $\nu_{\mathrm{MU}}\mathrm{MU}$-based and $\nu_{\mathrm{MU}}\mathbf{F}_2$-based Adams spectral sequence internal to $\mathrm{Syn}_{\mathrm{MU}}$ for ν_{MU} tmf. By playing these out against each other, they are able to determine many differentials and extensions through much more algebraic methods, in many places avoiding the use of Toda brackets. Combined with synthetic and filtered techniques (explained in this thesis in Part I), this proves to be a very effective method for computing the further structure of tmf.

Appendices

Appendix A

Tables and charts

A.1 Tables

Tables A.1 to A.6 collect the lifts, hidden extensions, Toda brackets, and total differentials proved in the 2-primary computation of Chapter 9. Every entry in the table is accompanied with the location where the element is defined or the relation is proved.

In these tables, the term *validity* refers to the number k for which the element or relation lives in Smf/τ^k, where validity ∞ means it lives in Smf. However, we only list the validity that we prove (and that we require); listing a finite validity does not necessarily mean that it does not lift further. For total differentials, the validity is a pair of numbers (n, N); this refers to the total differential ∂_n^N. We omit the unknown units in the formulas for the total differentials, and refer to the location where the differential is proved for the expression with units included. For the Toda brackets, if no indeterminacy is listed, this means it is zero.

Table A.1: Elements imported from the sphere to $\mathrm{Smf}_{(2)}$.

Name	Degree	Detected by	Location	Comment
η	$(1,1)$	h_1	8.27, 9.2	
ν	$(3,1)$	h_2	8.27, 9.2	
ε	$(8,2)$	c	8.27, 9.2	
κ	$(14,2)$	d	8.28, 9.2	
$\bar{\kappa}$	$(20,4)$	g	8.34, 9.2	Determined up to $\nu^2\kappa$-multiples

Table A.2: Lifts of elements from $\mathrm{Smf}_{(2)}/\tau$ to $\mathrm{Smf}_{(2)}/\tau^k$.

Name	Degree	Validity	Lift of	Location
Δ	$(24,0)$	4	Δ in Smf/τ	9.17
η_1	$(25,1)$	8	$h_1\Delta$ in Smf/τ	9.22
		10		9.44
η_4	$(97,1)$	8	$h_1\Delta^4$ in Smf/τ	9.22
η_5	$(121,1)$	8	$h_1\Delta^5$ in Smf/τ	9.22
		10		9.44
ν_1	$(27,1)$	14	$2\nu\Delta$ in Smf/τ^4	9.26
ν_2	$(51,1)$	14	$\nu\Delta^2$ in Smf/τ^4	9.26
ν_4	$(99,1)$	14	$\nu\Delta^4$ in Smf/τ^4	9.26
ν_5	$(123,1)$	14	$2\nu\Delta^5$ in Smf/τ^4	9.26
ν_6	$(147,1)$	14	$\nu\Delta^6$ in Smf/τ^4	9.26
ε_1	$(32,2)$	12	$c\Delta$ in Smf/τ	9.37
ε_4	$(104,2)$	12	$c\Delta^4$ in Smf/τ	9.37
ε_5	$(128,2)$	12	$c\Delta^5$ in Smf/τ	9.37
$\bar{\kappa}_4$	$(116,4)$	20	$2g\Delta^4$ in Smf/τ	9.46

Table A.3: A few τ-power torsion relations in $\pi_{*,*}\,\mathrm{Smf}_{(2)}/\tau^k$.

Relation	Degree	Validity	Location
$\tau^4\nu\bar{\kappa}=0$	$(23,1)$	12	9.33
		24	9.56
$\tau^4\nu_1\bar{\kappa}=0$	$(47,1)$	12	9.35
$\tau^4\nu_4\bar{\kappa}=0$	$(119,1)$	12	9.33
$\tau^4\nu_5\bar{\kappa}=0$	$(143,1)$	12	9.35

Table A.4: Relations and hidden relations in $\pi_{*,*}\,\mathrm{Smf}_{(2)}/\tau^k$.

Relation	Degree	Validity	Location
$4\nu = \tau^2\eta^3$	$(3,1)$	14	9.15
$2\nu_1 = \tau^2\eta^2\eta_1$	$(27,1)$	8	9.34
$2\nu_5 = \tau^2\eta^2\eta_5$	$(123,1)$	8	9.34
$\nu_1\eta = \tau^4\varepsilon\bar{\kappa}$	$(28,2)$	12	9.39
$\nu_1\varepsilon = \tau^4\eta\kappa\bar{\kappa}$	$(35,3)$	12	9.39
$\nu_1\kappa = \tau^6\eta\bar{\kappa}^2$	$(41,3)$	12	9.39
$\nu_2\eta = \tau^4\varepsilon_1\bar{\kappa}$	$(52,2)$	12	9.39
$\nu_2\varepsilon = \tau^4\eta_1\kappa\bar{\kappa}$	$(59,3)$	10	9.39
$\nu_5\eta = \tau^4\varepsilon_4\bar{\kappa}$	$(124,2)$	12	9.39
$\nu_5\varepsilon = \tau^4\eta_4\kappa\bar{\kappa}$	$(131,3)$	12	9.39
$\nu_5\kappa = \tau^6\eta_4\bar{\kappa}^2$	$(137,3)$	12	9.39
$\nu_6\eta = \tau^4\varepsilon_5\bar{\kappa}$	$(148,2)$	12	9.39
$\nu_6\varepsilon = \tau^4\eta_5\kappa\bar{\kappa}$	$(155,3)$	10	9.39
$\nu_2 \cdot 2\nu = \tau^8\bar{\kappa}^2\widetilde{d}$	$(54,2)$	10	9.52
$\nu_6 \cdot 2\nu = \tau^8\bar{\kappa}^2\widetilde{d}_4$	$(150,2)$	10	9.52

Table A.5: Total differentials on $\mathrm{Smf}_{(2)}/\tau$.

Source	Source degree	Target	Validity	Location
Δ	$(24,0)$	$\nu\bar{\kappa}$	$(4,14)$	9.28
Δ^2	$(48,0)$	$\nu_1\bar{\kappa}$	$(4,14)$	9.32
Δ^3	$(72,0)$	$3\nu_2\bar{\kappa}$	$(4,14)$	9.32
Δ^6	$(144,0)$	$3\nu_5\bar{\kappa}$	$(4,14)$	9.32
Δ^7	$(168,0)$	$7\nu_6\bar{\kappa}$	$(4,14)$	9.32

Table A.6: Toda brackets in $\mathrm{Smf}_{(2)}/\tau^k$.

Name	Degree	Toda bracket	Indeterminacy	Validity	Location
η_1	$(25,1)$	$\langle \bar{\kappa}, \tau^4\nu, \eta \rangle$	κ-torsion classes	10	9.47
η_5	$(121,1)$	$\langle \bar{\kappa}, \tau^4\nu_4, \eta \rangle$	κ-torsion classes	10	9.47
ν_1	$(27,1)$	$\langle \bar{\kappa}, \tau^4\nu, 2\nu \rangle$		12	9.38
ν_2	$(51,1)$	$\langle \bar{\kappa}, \tau^4\nu_1, \nu \rangle$		12	9.38
		$\langle \nu, \tau^4 2\nu, \tau^4\nu, \bar{\kappa}^2 \rangle$		10	9.61
ν_5	$(123,1)$	$\langle \bar{\kappa}, \tau^4\nu, 2\nu_4 \rangle$		12	9.38
ν_6	$(147,1)$	$\langle \bar{\kappa}, \tau^4\nu_1, \nu_4 \rangle$		12	9.38
		$\langle \nu_4, \tau^4 2\nu, \tau^4\nu, \bar{\kappa}^2 \rangle$		10	9.61
ε_1	$(32,2)$	$\langle \nu_1, \nu, \eta \rangle$		12	9.37
ε_4	$(104,2)$	$\langle \nu, 2\nu_4, \eta \rangle$		12	9.37
ε_4	$(128,2)$	$\langle \nu_1, \nu_4, \eta \rangle$		12	9.37
$\eta_1\kappa$	$(39,3)$	$\langle \nu_1, \nu, \varepsilon \rangle$		10	9.38
$\eta_4\kappa$	$(111,3)$	$\langle \nu, 2\nu_4, \varepsilon \rangle$		12	9.38
$\eta_5\kappa$	$(135,3)$	$\langle \nu_1, \nu_4, \varepsilon \rangle$		10	9.38
$\tau^2\eta_4\bar{\kappa}$	$(117,3)$	$\langle \nu, 2\nu_4, \kappa \rangle$		12	9.38
$\pm 2\bar{\kappa}$	$(20,4)$	$\langle \nu, \eta, \eta\kappa \rangle$	not discussed	10	9.38
$\tau^2\bar{\kappa}$	$(20,2)$	$\langle \kappa, 2, \eta, \nu \rangle$	not discussed	∞	8.35
$\pm\bar{\kappa}_4$	$(116,4)$	$\langle \nu_4, \eta, \eta\kappa \rangle$	not discussed	10	9.47

A.2 Descent spectral sequence charts

Here we display the descent spectral sequences for $\mathrm{Tmf}_{(3)}$ and $\mathrm{Tmf}_{(2)}$. We use the following conventions.

- Black arrows are differentials. Black lines are multiplication by either α or β for $\mathrm{Tmf}_{(3)}$, and either η or ν for $\mathrm{Tmf}_{(2)}$. Under the blue line is the *connective region* (Definition 9.3) and right of the orange line is the *S-region* (Definition 9.7).
- Red lines indicate hidden extensions by α for $\mathrm{Tmf}_{(3)}$ and by η, ν, ε, or κ for $\mathrm{Tmf}_{(2)}$. We only include those hidden extensions needed in this article.
- Hollow squares refer to $\mathbf{Z}_{(p)}$ and hollow circles to $\mathbf{F}_p$. For $\mathrm{Tmf}_{(2)}$, two enclosed circles represents $\mathbf{Z}/4$ and three enclosed circles represents $\mathbf{Z}/8$. Two symbols in the same bidegree represents their sum.

These charts correct some small oversights in those of [Kon12]. We would like to mention in particular the key d_{23}-differential of Proposition 9.57 is missing from Konter's chart of the DSS for $\mathrm{Tmf}_{(2)}$.

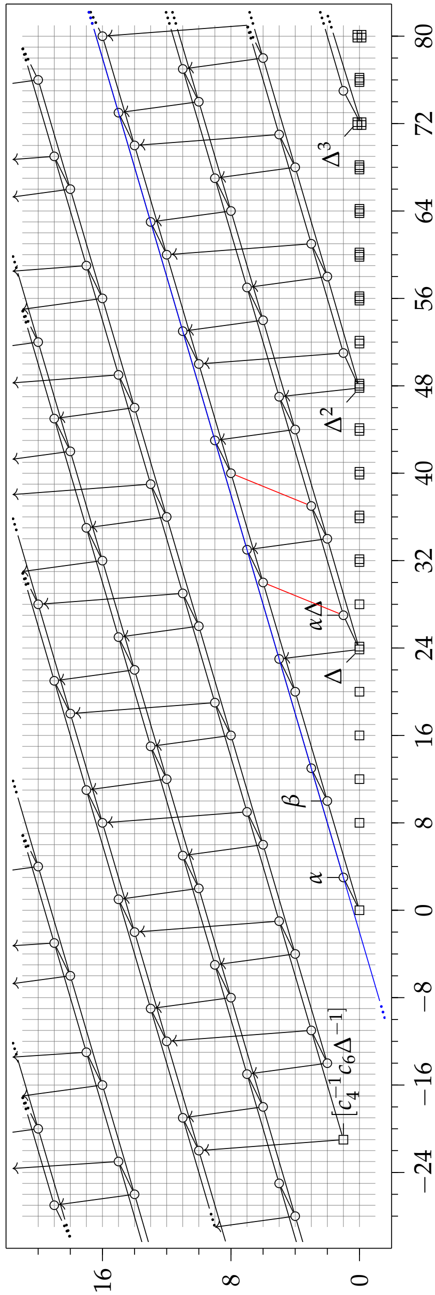

Figure A.1: The DSS for $\mathrm{Tmf}_{(3)}$. The lines represent multiplication by either α or β.

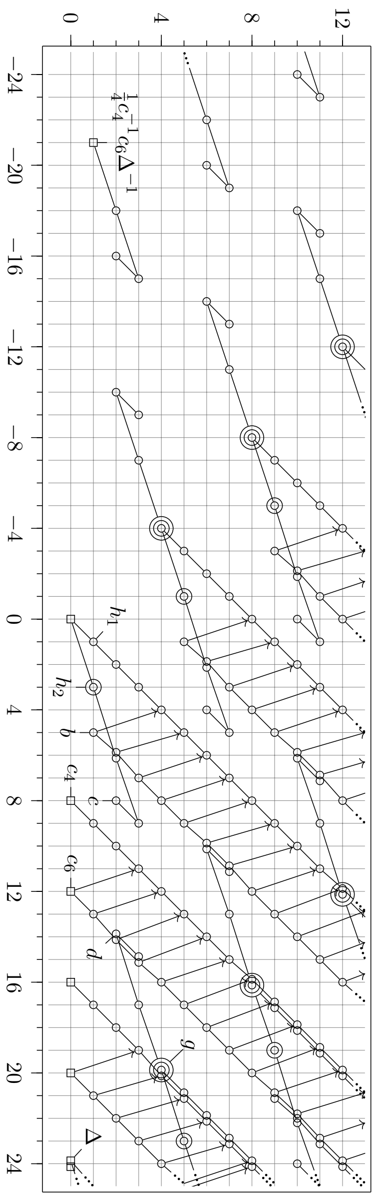

Figure A.2: The E_3-page of the DSS for $\mathrm{Tmf}_{(2)}$.

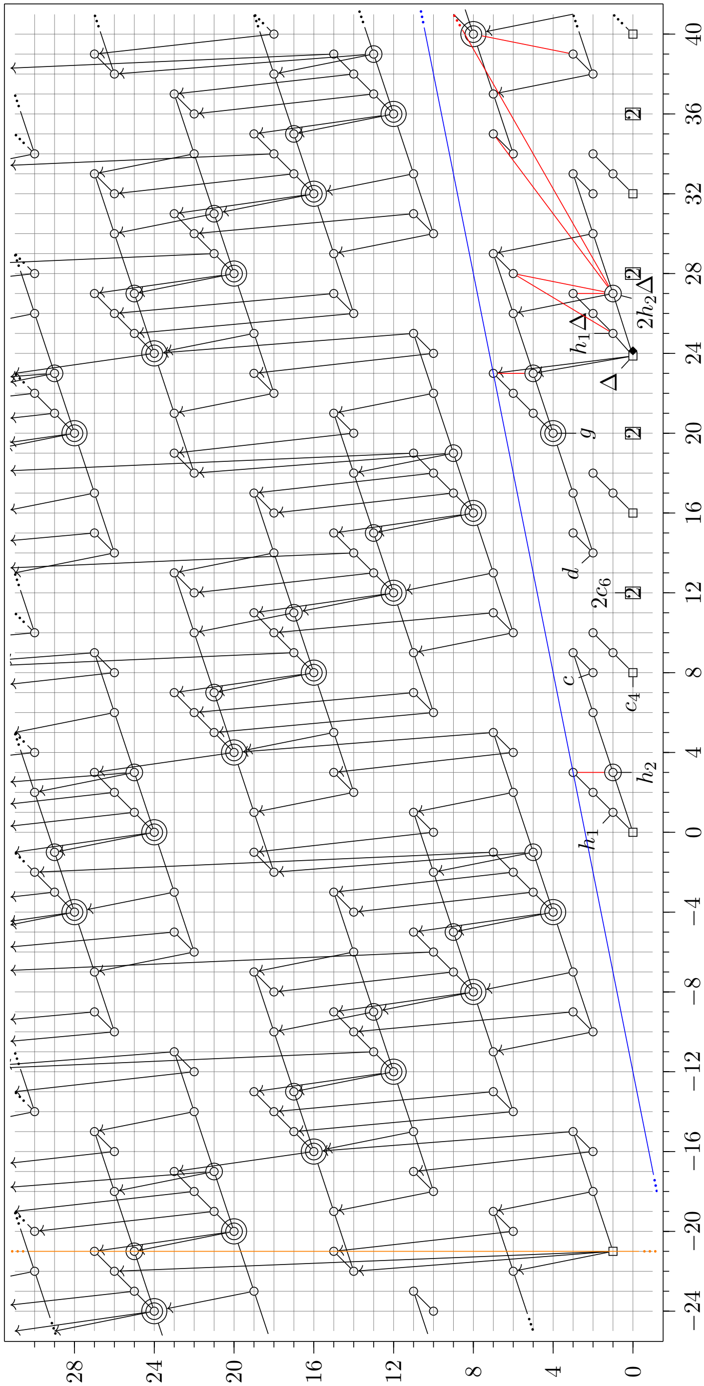

Figure A.3: The DSS for $\mathrm{Tmf}_{(2)}$ from the E_5-page onward in the range $-24 \leqslant n \leqslant 40$.

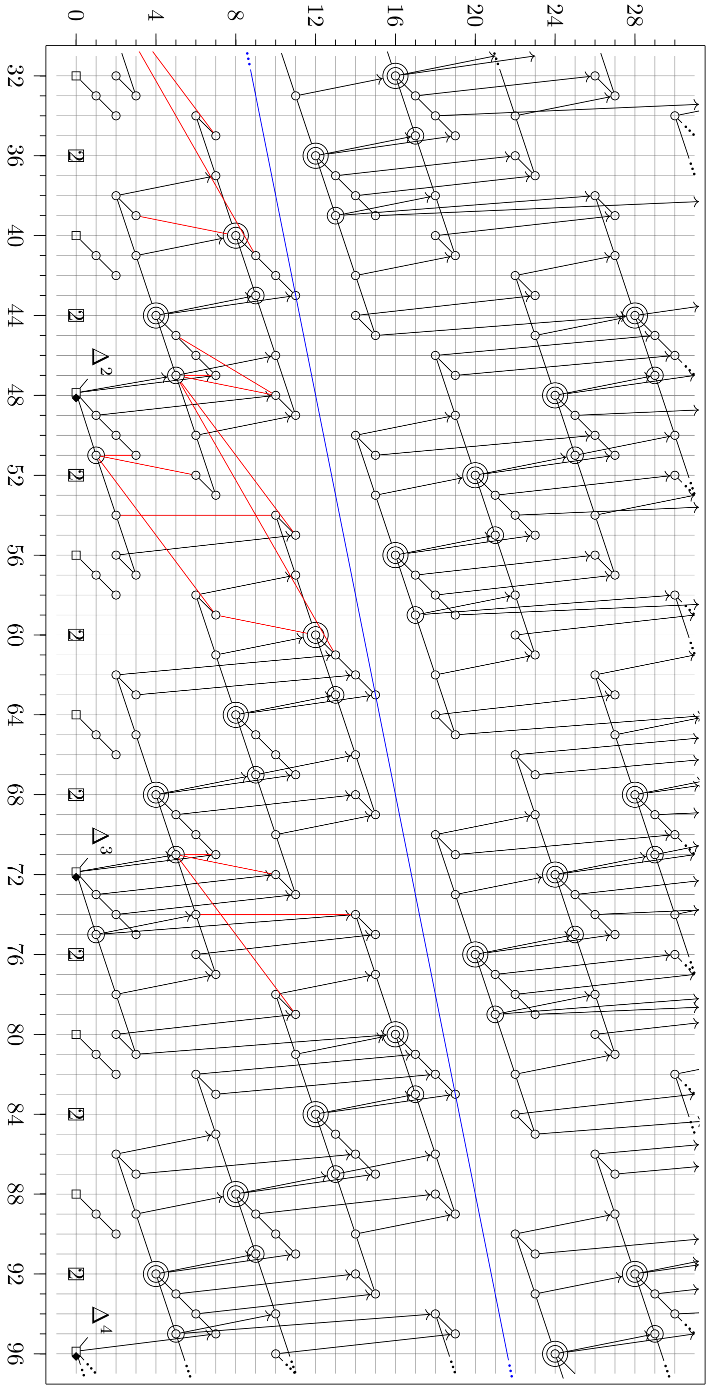

Figure A.4: The DSS for $\mathrm{Tmf}_{(2)}$ from the E_5-page onward in the range $32 \leqslant n \leqslant 96$.

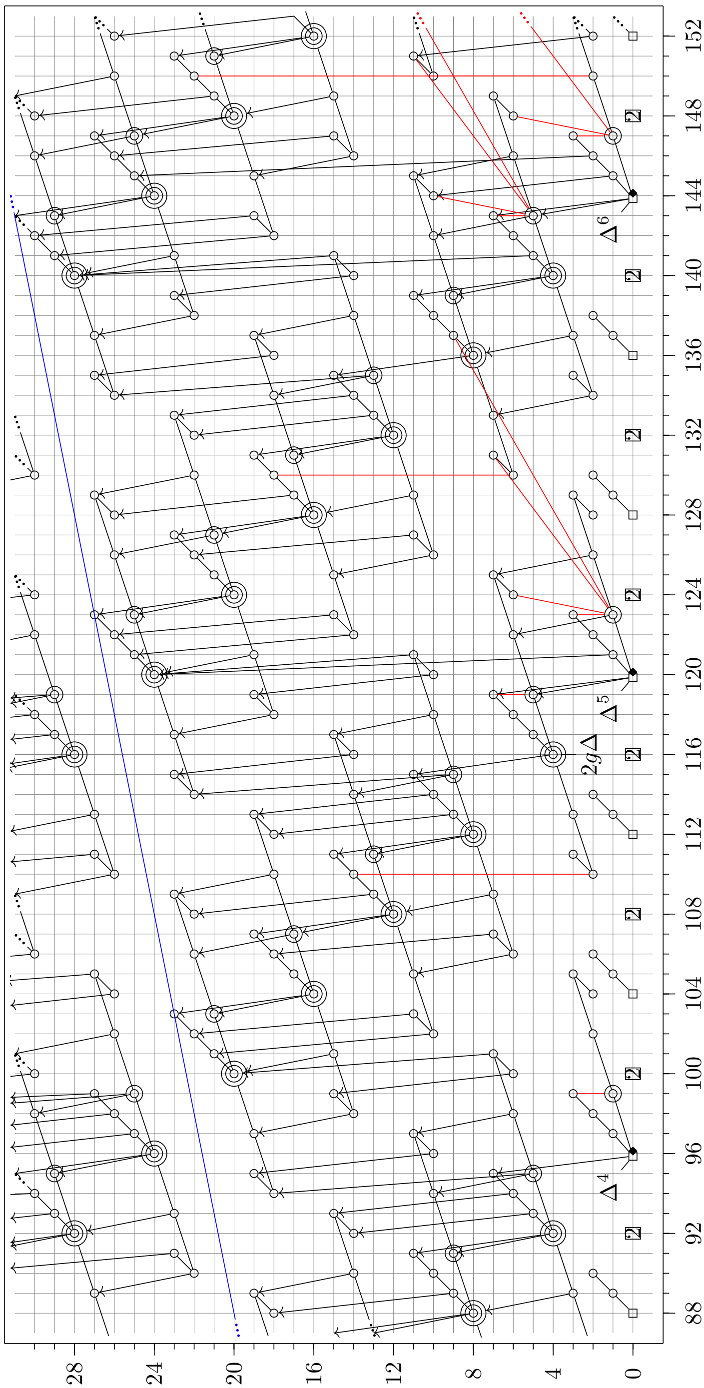

Figure A.5: The DSS for $\mathrm{Tmf}_{(2)}$ from the E_5-page onward in the range $88 \leqslant n \leqslant 152$.

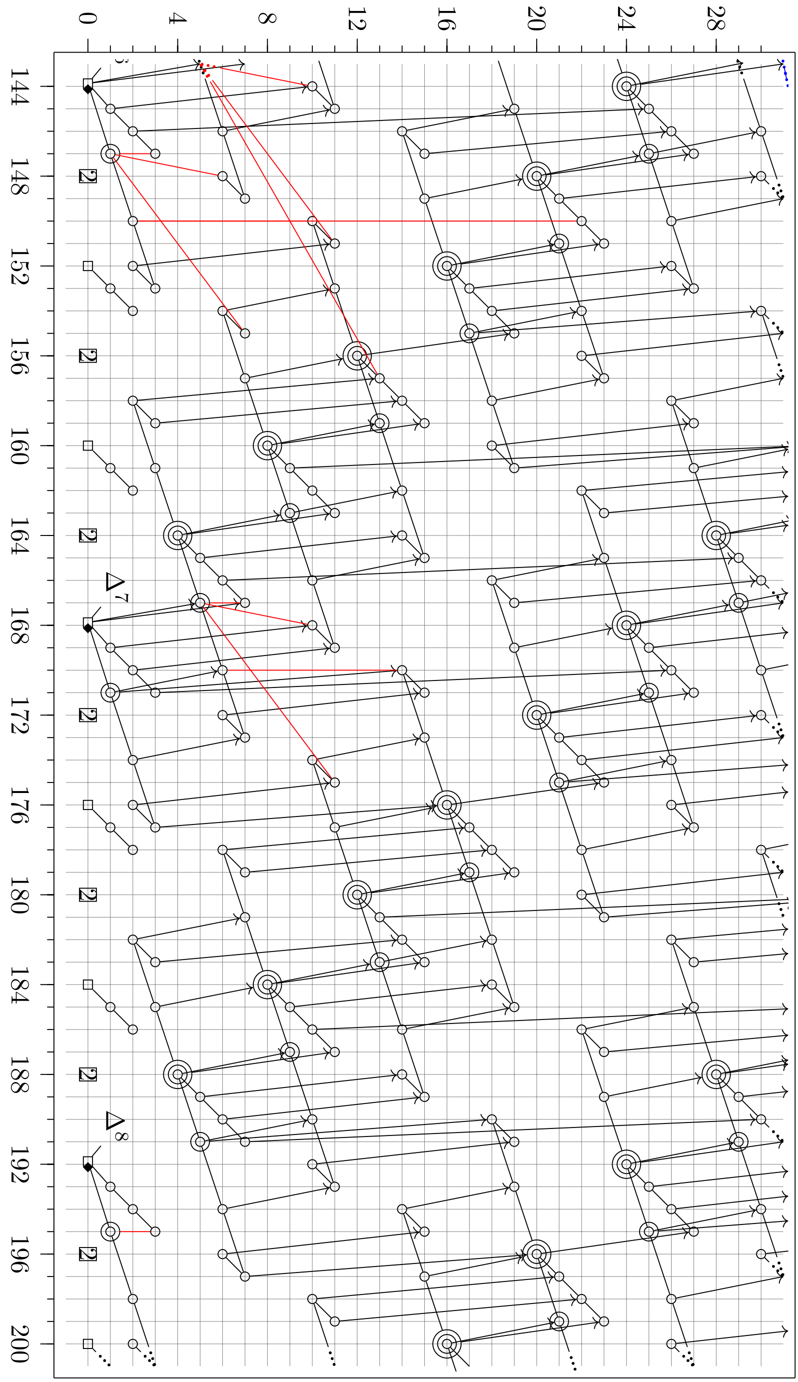

Figure A.6: The DSS for $\mathrm{Tmf}_{(2)}$ from the E_5-page onward in the range $144 \leqslant n \leqslant 200$.

Appendix B

Toda brackets

In this brief appendix, we collect a few well-known facts about Toda brackets to be used in the main text. This is merely a condensed form of the more extensive treatment we give in [CDvN24, Section 3]. For the part about Moss's Theorem, we closely follow ideas of Burklund from [Bur22].

It has long been known to experts that one can define Toda brackets in any sufficiently coherent homotopical context, for example, in an ∞-category. For simplicity and with our eyes towards applications to synthetic, equivariant, and motivic spectra, as well as the associated categories of modules, we restrict our attention to the Picard-graded homotopy groups of the unit in a monoidal stable ∞-category.

B.1 Toda brackets and iterated cones

Let $(\mathcal{C}, \otimes, \mathbf{1})$ be a stably monoidal ∞-category.

Definition B.1. We let a **shift** of an object or morphism in $\mathcal{C}$ refer to applying $X \otimes -$ or $- \otimes X$ for some $X \in \operatorname{Pic}(\mathcal{C})$.

We will abuse notation slightly and use the same name for a map $f\colon X \to Y$ for $X, Y \in \operatorname{Pic}(\mathcal{C})$ and its shift $f \otimes Z\colon X \otimes Z \to Y \otimes Z$ for $Z \in \operatorname{Pic}(\mathcal{C})$. One important example is the dual map $Y^\vee \to X^\vee$, as this can be identified with the shift of $Y^{-1} \otimes f \otimes X^{-1}$ of f up to a sign, in the following sense.

Definition B.2. We say that (sets of) morphisms agree up to a **sign** when they become homotopic after multiplying by a unit of the ring $[\mathbf{1}, \mathbf{1}]$ of endomorphisms of the unit of $\mathcal{C}$.

All of the formulas given below only hold up to a sign, and we do not make these signs explicit.

Definition B.3. Let $X_0, \ldots, X_n \in \mathrm{Pic}(\mathcal{C})$ and

$$X_n \xrightarrow{a_n} X_{n-1} \longrightarrow \cdots \longrightarrow X_1 \xrightarrow{a_1} X_0$$

a composable sequence of arrows in $\mathcal{C}$. We define the data of a **form of $\boldsymbol{C(a_1, \ldots, a_n)}$** in $\mathcal{C}$ inductively as follows:

- A form of $C(a_1)$ is a cofiber of a_1 together with the canonical projection $C(a_1) \to \Sigma X_1$. There is a contractible space of forms of $C(a_1)$.
- Suppose that forms of $C(a_1, \ldots, a_{n-1})$ have been defined so that each such form comes with a canonical "projection to the top cell"

  $$C(a_1, \ldots, a_{n-1}) \longrightarrow \Sigma^{n-1} X_{n-1}$$

  A form of $C(a_1, \ldots, a_n)$ is then a cofiber of a morphism

  $$\Sigma^{n-1} X_n \longrightarrow C(a_1, \ldots, a_{n-1})$$

  such that the composite

  $$\Sigma^{n-1} X_n \longrightarrow C(a_1, \ldots, a_{n-1}) \longrightarrow \Sigma^{n-1} X_{n-1}$$

  is homotopic to $\Sigma^{n-1} a_n$, for some form of $C(a_1, \ldots, a_{n-1})$. Rotating the cofiber sequence, one obtains a canonical map

  $$C(a_1, \ldots, a_n) \longrightarrow \Sigma^n X_n$$

Remark B.4. A form of $C(a_1, \ldots, a_n)$ is a cell complex built out of $X_0, \ldots, X_n$ which may be depicted as follows

$$\begin{tikzcd}
\boxed{\Sigma^n X_n} \arrow[d, dash, "a_n"] \\
\boxed{\Sigma^{n-1} X_{n-1}} \arrow[d, dash] \\
\vdots \arrow[d, dash] \\
\boxed{\Sigma X_1} \arrow[d, dash, "a_1"] \\
\boxed{X_0}
\end{tikzcd}$$

By definition such a form fits into a canonical cofiber sequence

$$C(a_1, \ldots, a_{n-1}) \longrightarrow C(a_1, \ldots, a_n) \longrightarrow \Sigma^n X_n$$

for some form of $C(a_1, \ldots, a_{n-1})$. Except in the case $n = 1$, these do not always exist and are not in general unique when they do.

Example B.5. The dual of a form of $C(a_1, \ldots, a_n)$ is a shift of a form of $C(a_n, \ldots, a_1)$, since taking duals preserves cofiber sequences. ▲

Definition B.6. Let K be a form of $C(a_1, \ldots, a_n)$. For maps $a\colon X \to X_n$ and $b\colon X_0 \to X$, we use the notation

$$\Sigma^n X \xrightarrow{\overline{a}} K$$
$$K \xrightarrow{\underline{b}} X$$

to denote maps that make the following diagrams commute

$$\begin{tikzcd} X_0 \arrow[r, "b"] \arrow[d, "\iota"'] & X \\ K \arrow[ur, "\underline{b}"'] & \end{tikzcd} \qquad \begin{tikzcd} \Sigma^n X \arrow[r, "a"] \arrow[d, "\overline{a}"'] & \Sigma^n X_n, \\ K \arrow[ur, "p"'] & \end{tikzcd}$$

where ι and p are the canonical inclusions and projections respectively.

Definition B.7. Let K be a form of $C(a_2, \ldots, a_{n-1})$, and suppose one has maps

$$\Sigma^{n-2} X_n \xrightarrow{\overline{a_n}} K$$
$$K \xrightarrow{\underline{a_1}} X_0$$

as in Definition B.6. The **Toda bracket**

$$\langle a_1, \ldots, a_n \rangle \subseteq [\Sigma^{n-2} X_n, X_0]$$

is the set of all composites $\underline{a_1} \circ \overline{a_n}$ formed in this way, by running over all forms K of $C(a_2, \ldots, a_{n-1})$ and all possible choices of $\underline{a_1}$ and $\overline{a_n}$.

Remark B.8. In a diagram, this is the composite

$$\begin{tikzcd} \boxed{\Sigma^{n-2} X_n} \arrow[r, "a_n"] & \boxed{\Sigma^{n-2} X_{n-1}} \arrow[d, dash, "a_{n-1}"] & \\ & \vdots \arrow[d, dash, "a_2"] & \\ & \boxed{X_1} \arrow[r, "a_1"] & \boxed{X_0} \end{tikzcd}$$

Lemma B.9. *Let $F\colon \mathcal{C} \to \mathcal{D}$ be a unital exact functor between symmetric monoidal stable ∞-categories. For any morphisms $a_i \in \mathcal{C}$ as above, there is a containment*

$$F(\langle a_1, \dots, a_n \rangle) \subseteq \langle F(a_1), \dots, F(a_n) \rangle.$$

Proof. The assumptions guarantee that F preserves forms of $C(a_1, \dots, a_n)$. ∎

Lemma B.10. *Up to signs, one has*

$$\langle a_1, \dots, a_n \rangle = \langle a_n, \dots, a_1 \rangle.$$

Proof. This follows immediately from the definition using Example B.5. ∎

The set $\langle a_1, \dots, a_n \rangle$ may be empty. Indeed, a form of $C(a_2, \dots, a_{n-1})$ need not exist when $n > 3$, nor the maps $\underline{a_1}$ and $\overline{a_n}$. We therefore give some general existence statements for forms of iterated cones and non-emptiness of Toda brackets. In the following, a *sub-bracket* of $\langle a_1, \dots, a_n \rangle$ refers to any bracket of the form $\langle a_i, a_{i+1}, \dots, a_{i+k} \rangle$.

Proposition B.11.

(1) *A form of $C(a_1, a_2, a_3)$ exists if and only if $a_1 a_2 = a_2 a_3 = 0$ and $0 \in \langle a_1, a_2, a_3 \rangle$. The Toda bracket $\langle a_1, a_2, a_3 \rangle$ is nonempty if and only if $a_1 a_2 = a_2 a_3 = 0$.*

(2) *Let $n > 3$. A form of $C(a_1, \dots, a_n)$ exists if and only if all sub-brackets of $\langle a_1, \dots, a_n \rangle$ contain zero.*

(3) *Let $n > 3$. If $\langle a_1, \dots, a_k \rangle = \{ 0 \}$ for all $k < n$, and $0 \in \langle a_2, \dots, a_n \rangle$, then $\langle a_1, \dots, a_n \rangle$ is nonempty. Conversely, if $\langle a_1, \dots, a_n \rangle$ is nonempty, then all sub-brackets of $\langle a_1, \dots, a_n \rangle$ of length $< n$ contain zero.*

Proof. See [CDvN24, Theorem 3.11]. ∎

Proposition B.12. *The following shuffling formulas hold up to sign.*

(1) *If $\langle a_i, \dots, a_n \rangle = \langle a_1, \dots, a_i \rangle = \{ 0 \}$ for all $3 \leqslant i < n-1$, then*

$$a_1 \langle a_2, \dots, a_n \rangle = \langle a_1, \dots, a_{n-1} \rangle a_n.$$

(2) *Let $a\colon X \to X_n$ be a map. Then $\langle a_1, \dots, a_n \rangle a \subseteq \langle a_1, \dots, a_n a \rangle$.*

(3) *Let $a\colon X \to X_n$ be a map. Then*

$$\langle a_1, \dots, a_n a \rangle \subseteq \langle a_1, \dots, a_{n-1} a, a_n \rangle.$$

(4) *Let $2 < k < n$ and let $a : X_{k-1} \to X$ be a map. If every sub-bracket of the Toda bracket $\langle a_2, \dots, a_{k-1} a, a_k, \dots, a_n \rangle$ contains zero and every sub-brackets of the Toda bracket $\langle a_1, \dots, a_{k-1}, a a_k, \dots, a_{n-1} \rangle$ is equal to $\{ 0 \}$, then*

$$\langle a_1, \dots, a_{k-1} a, a_k, \dots, a_n \rangle \cap \langle a_1, \dots, a_{k-1}, a a_k, \dots, a_n \rangle \neq \varnothing.$$

Proof. See [CDvN24, Proposition 3.13]. ■

B.2 Moss's Theorem

Moss's Theorem gives conditions for Toda brackets in $\pi_* R$ to be detected by Massey products in a spectral sequence converging to $\pi_* R$. The classical reference is [Mos70]; see [BK25] for a modern reinterpretation. We will discuss a version of Moss's Theorem in the setting of synthetic spectra. Throughout this section, we fix a homotopy-associative ring spectrum E of Adams type, and write Syn_E for Syn.

This result is special to the synthetic setting, and is not available in this form in the filtered setting. The reason is that the special fibre of Syn is given by a derived category, where Toda brackets are computed by Massey products; see Proposition B.14 below.

Let $R \in \mathrm{Syn}$ be a synthetic $\mathbf{E}_2$-ring, then we may work with Toda brackets in the monoidal ∞-category of (left) modules over R. Given a Toda bracket

$$\langle a_1, \ldots, a_n \rangle \subseteq \pi_{*,*}\, R,$$

we will describe a general technique for determining the image of $\langle a_1, \ldots, a_n \rangle$ along the map $\mathrm{pr}_r \colon \pi_* R \to \pi_*\, R/\tau^r$, following ideas of Robert Burklund [Bur22].

There is one immediate observation.

Proposition B.13. *There is an inclusion*

$$\mathrm{pr}_r(\langle a_1, \ldots, a_n \rangle) \subseteq \langle a_1, \ldots, a_n \rangle \subseteq \pi_{*,*}\, R/\tau^r.$$

Proof. The map $R \to R/\tau^r$ is one of $\mathbf{E}_2$-rings, so the claim follows from Lemma B.9. ■

Suppose C_* is a dga in an abelian category $\mathcal{A}$. Recall that, when we regard C_* as an object of the derived ∞-category $\mathcal{D}(\mathcal{A})$, then Toda brackets in the homotopy groups of a dga C_* coincide with Massey products in C_*. This immediately implies the following.

Proposition B.14. *If the $\mathbf{E}_2$-ring R/τ is the underlying object in* Stable *of a dga C_*, then the set $\mathrm{pr}_1(\langle a_1, \ldots, a_n \rangle)$ is contained in the Massey product $\langle a_1, \ldots, a_n \rangle \subseteq \mathrm{H}_* C_*$.*

However, this is not particularly useful if one of the a_i's is divisible by τ^r for $r \geqslant 1$. In that case, the bracket in R/τ^r will contain zero and have large indeterminacy generically.

Instead of proving a version of Moss's Theorem in this context, we describe a general approach to determining the image $\mathrm{pr}_r(\langle a_1, \ldots, a_n \rangle)$ that works well in most cases.

We outline how this works in a simple situation that happens to cover all of our applications in this article; see Section 9.9.

Theorem B.15. *Let R be an $\mathbf{E}_2$-ring in* Syn. *Let $a_1, a_2, a_3 \in \pi_{*,*} R$, and let $r \geqslant s \geqslant 0$ be minimal such that*
$$\tau^r a_1 a_2 = \tau^s a_2 a_3 = 0 \qquad \text{in } \pi_{*,*} R.$$

(1) *Suppose that $r, s > 0$. Then there exist $H_0, H_1 \in \pi_{*,*}(R/\tau^s)$ such that*
$$d_{r+1}(H_0) = a_1 a_2$$
$$d_{s+1}(H_1) = a_2 a_3,$$
and the Toda bracket $\langle a_1, \tau^r a_2, a_3 \rangle \subseteq \pi_{*,*} R$ contains a lift of
$$\tau^{r-s} a_1 H_1 \pm H_0 a_3 \in \pi_{*,*}(R/\tau^s).$$
(*Note that here we do not change the name of an element along either of the maps in the composite*
$$R \longrightarrow R/\tau^n \longrightarrow R/\tau$$
for $n = r, s$.)

In particular, if $r = s$, then the Toda bracket $\langle a_1, \tau^r a_2, a_3 \rangle$ contains a lift of the element $a_1 H_1 \pm H_0 a_3$ in $\pi_{*,*}(R/\tau)$. If $r > s$, then the Toda bracket $\langle a_1, \tau^r a_2, a_3 \rangle$ contains a lift of the element $H_0 a_3$ in $\pi_{*,*}(R/\tau)$. (*Note that in both cases, these linear combinations are contained in the Massey product $\langle a_1, a_2, a_3 \rangle$ formed in the dga* $(\mathrm{E}_{r+1}, d_{r+1})$ *in the spectral sequence underlying σR.*)

(2) *Suppose that $r > 0$ while $s = 0$. Then there exists $H_0 \in \pi_{*,*}(R/\tau^r)$ such that*
$$d_{r+1}(H_0) = a_1 a_2$$
and the Toda bracket $\langle a_1, \tau^r a_2, a_3 \rangle \subseteq \pi_{*,*} R$ contains a lift of
$$az \in \pi_{*,*}(R/\tau^r).$$
In particular, the Toda bracket contains a lift of an element in the Massey product $\langle a_1, a_2, a_3 \rangle$ formed in the dga $(\mathrm{E}_{r+1}, d_{r+1})$ *in the spectral sequence underlying σR.*

Proof. See [CDvN24, Theorem 3.16]. ∎

Remark B.16. Usually, any version of Moss's Theorem contains delicate assumptions about crossing differentials. This does not appear here because we are only working synthetically and we assume that we have elements $a_i \in \pi_{*,*} R$ such that $\tau^r a_1 a_2 = \tau^s a_2 a_3 = 0$. If, on the other hand, we started with elements $b_i \in \pi_* \tau^{-1} R$ such that $b_1 b_2 = b_2 b_3 = 0$, with $d_{r+1}(H_0) = b_1 b_2$ and $d_{s+1}(H_1) = b_2 b_3$ in the spectral sequence σR, then it is not automatic that there exists lifts a_i of b_i along τ^{-1} so that both $\tau^r a_1 a_2$ and $\tau^s a_2 a_3$ are zero.

Remark B.17. The theorem above gives conditions for when certain Toda brackets in $\pi_{*,*} R$ contain lifts of certain elements in $\pi_{*,*}(R/\tau^s)$. We will also need similar statements about lifts along truncation maps of the form

$$\pi_{*,*}(R/\tau^{s+k}) \longrightarrow \pi_{*,*}(R/\tau^s)$$

The theorem does not apply verbatim to this case as $(R/\tau^{s+k})/\tau^s \not\simeq R/\tau^s$. However, there is essentially no difference between these for the purposes of the theorem. Indeed, there is a canonical splitting

$$(R/\tau^{s+k})/\tau^s \cong R/\tau^s \oplus \Sigma^{1,-s-k-1} R/\tau^s.$$

Moreover, the spectral sequence arising from the τ-adic tower that begins with the homotopy groups of $(R/\tau^{s+k})/\tau^s$ and converges to the homotopy groups of R/τ^{s+k} splits in an analogous way into two copies of the truncated Bockstein spectral sequence beginning with $\pi_{*,*}(R/\tau^s)$ and converging to $\pi_{*,*}(R/\tau^{s+k})$. In particular, the differentials in one of these two spectral sequences determine the other, hence the same is true of the Massey products formed in the spectral sequences.

Appendix C

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 $X\colon \mathbf{Z}^{\mathrm{op}} \to \mathrm{Sp}$. 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.

C.1 The reconstruction problem

We have a cofibre sequence

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

leading to a long exact sequence

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

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

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

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

This way of thinking only goes so far: it presupposes that we understand the elements of $\pi_n X^s$, which we usually do not. In 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 $\pi_n \operatorname{Gr}^s X$ for all s, and then try to piece the $\pi_n X^s$ back together from this data. This presents two issues:

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

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

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

C.2 Differentials: obstructions to lifting

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

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

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

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

$$\begin{tikzcd}
\cdots \arrow[r] & X^2 \arrow[r] & X^1 \arrow[r] & X^0 \arrow[d] \\
 & & & \mathrm{Gr}^0 X, \arrow[ul, dashed, "\partial"]
\end{tikzcd}$$

where the dashed arrow indicates that the map is of degree 1: it is the boundary map $\partial\colon \mathrm{Gr}^0 X \to \Sigma X^1$ of a cofibre sequence. By exactness, an element $x \in \pi_n \mathrm{Gr}^0 X$ 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 $X^1 \to \mathrm{Gr}^1 X$ is zero:

$$\begin{tikzcd}
\cdots \arrow[r] & X^2 \arrow[r] & X^1 \arrow[r] \arrow[d] & X^0 \arrow[d] \\
 & & \mathrm{Gr}^1 X & \mathrm{Gr}^0 X. \arrow[ul, dashed, "\partial"]
\end{tikzcd}$$

Write $d_1(x)$ for the image of ∂x in $\pi_{n-1} \mathrm{Gr}^1 X$. 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 $\pi_{n-1} \mathrm{Gr}^2 X$:

$$\begin{tikzcd}
\cdots \arrow[r] & X^2 \arrow[r] \arrow[d] & X^1 \arrow[r] & X^0 \arrow[d] \\
 & \mathrm{Gr}^2 X & & \mathrm{Gr}^0 X. \arrow[ul, dashed, "\partial"]
\end{tikzcd}$$

This choice of lift will not be unique, and neither will the resulting class in $\pi_{n-1} \mathrm{Gr}^2 X$; the class in $\pi_{n-1} \mathrm{Gr}^2 X$ is only well defined up to the image of d_1. We write $d_2(x)$ for this element in $(\pi_{n-1} \mathrm{Gr}^2 X)/d_1$. 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 $\pi_{n-1} \mathrm{Gr}^3 X$ (only well defined up to d_1 and d_2), et cetera.

We obtain inductively defined elements $d_r(x)$ for $r \geqslant 1$. 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 $r \geqslant 1$. This, in turn, means that we can answer the question whether $x \in \pi_n \mathrm{Gr}^0 X$ 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 $\pi_n \operatorname{Gr}^0 X$, the same applies when starting with an element of $\pi_n \operatorname{Gr}^s X$, which lifts to $\pi_n X^s$ if and only if all differentials on it vanish.

C.3 Differentials: kernels of transition maps

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

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

so by exactness, α maps to zero in $\pi_n X^{s-1}$ if and only if it is in the image of the map $\pi_{n+1} \operatorname{Gr}^s X \to \pi_n X^s$. Notice that in terms of x, this is equivalent to the existence of an element $y \in \pi_{n+1} \operatorname{Gr}^{s-1}$ 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 $d_1, \ldots, d_s$.

Remark C.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 $x \in \pi_n \operatorname{Gr}^s X$. 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 $r \leqslant s$ 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 $\pi_n \operatorname{Gr}^s X$. 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 $\pi_n \operatorname{Gr}^s X$ only sees phenomena *up to higher filtration*.

This is a problem that we simply have to live with if all we understand is the associated graded. It can be delicate matter to check that a lift of an element hit by a d_r-differential is the lift that dies r filtrations down. In 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.

C.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

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

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

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

as the boundary map $\mathrm{E}_1^{n,s} \to \pi_{n-1} X^{s+1}$ followed by the projection $\pi_{n-1} X^{s+1} \to \mathrm{E}_1^{n-1,\, s+1}$.

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 $d_1, \ldots, d_{r-1}$. We therefore inductively define, for $r \geqslant 2$,

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

Roughly speaking, doing this process infinitely many times results in page ∞, denoted $\mathrm{E}_\infty^{n,s}$. In good cases (the other part of convergence issues), this is isomorphic to the associated graded of the induced strict filtration on $\pi_n X^0$:

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

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 C.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 $\mathbf{Z} \to \mathrm{Sp}$, 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 $\pi_n X^s \to \pi_n X^{s-1}$. Taking homology for a d_r-differential removes both the fake elements, and kills elements that die under $X^s \to X^{s-r}$. Letting all differentials run brings us to the associated graded of the filtration we were trying to understand.

C.5 Reformulation in terms of τ

It would be useful to introduce some notation to make it easier to describe transitions 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 $X^s \to X^{s-1}$. 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 $x \in \pi_n \operatorname{Gr}^s X$ 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

[AOV08] D. Abramovich, M. Olsson and A. Vistoli. 'Tame Stacks in Positive Characteristic'. *Annales de l'institut Fourier* 58.4 (2008), pp. 1057–1091. DOI: `10.5802/aif.2378`.

[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.

[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.

[BH21] T. Bachmann and M. Hoyois. 'Norms in Motivic Homotopy Theory'. *Astérisque* 425 (1st Sept. 2021). DOI: `10.24033/ast.1147`.

[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. Pstrągowski. '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.

[BGH20] C. Barwick, S. Glasman and P. Haine. *Exodromy*. 23rd Aug. 2020. arXiv: 1807.03281v7 [math]. Preprint.

[Bau08] T. Bauer. 'Computation of the homotopy of the spectrum tmf'. In: *Groups, homotopy and configuration spaces (Tokyo 2005)*. Groups, homotopy and configuration spaces, in honour of Fred Cohen's 60th birthday. Mathematical Sciences Publishers, 22nd Feb. 2008, pp. 11–40. DOI: 10.2140/gtm.2008.13.11.

[BL10] M. Behrens and T. Lawson. *Topological Automorphic Forms*. Online-Ausg. Memoirs of the American Mathematical Society volume 204, number 958. American Mathematical Society, 2010.

[BK25] E. Belmont and H. J. Kong. 'A Toda Bracket Convergence Theorem for Multiplicative Spectral Sequences'. *Peking Mathematical Journal* (1st June 2025). DOI: 10.1007/s42543-025-00099-x.

[BMS19] B. Bhatt, M. Morrow and P. Scholze. 'Topological Hochschild homology and integral p-adic Hodge theory'. *Publications mathématiques de l'IHÉS* 129.1 (June 2019), pp. 199–310. DOI: 10.1007/s10240-019-00106-9.

[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.

[BR] R. R. Bruner and J. Rognes. *The Adams Spectral Sequence for Topological Modular Forms*. Mathematical Surveys and Monographs volume 253. American Mathematical Society, 2021. 690 pp. URL: `https://www.mn.uio.no/math/personer/vit/rognes/papers/surv-253.pdf`.

[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. Pstrągowski. 'Quivers and the Adams Spectral Sequence'. *Advances in Mathematics* 471 (June 2025), p. 110270. DOI: `10.1016/j.aim.2025.110270`.

[Car22] C. Carrick. 'Stacks and Real-Oriented Homotopy Theory'. PhD thesis. University of California, 2022. URL: `https://escholarship.org/uc/item/4bz5k1pb`.

[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`.

[CE16] H. Cartan and S. Eilenberg. *Homological Algebra (PMS-19)*. Princeton Mathematical Series 19. Princeton University Press, 2016. 1 p. 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'. Master's thesis. Utrecht University, 22nd Jan. 2025. URL: https://studenttheses.uu.nl/handle/20.500.12932/48571.

[Dav23] J. M. Davies. 'Elliptic Cohomology Is Unique up to Homotopy'. *Journal of the Australian Mathematical Society* 115.1 (Aug. 2023), pp. 99–118. DOI: 10.1017/S1446788722000209.

[Dav24a] J. M. Davies. 'Hecke Operators on Topological Modular Forms'. *Advances in Mathematics* 452 (Aug. 2024), p. 109828. DOI: 10.1016/j.aim.2024.109828.

[Dav24b] J. M. Davies. 'On Lurie's Theorem and Applications'. *Mathematische Zeitschrift* 309.2 (23rd Dec. 2024), p. 24. DOI: 10.1007/s00209-024-03626-1.

[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/.

[Del75] P. Deligne. 'Courbes Elliptiques: Formulaire d'après J. Tate'. In: *Modular Functions of One Variable IV*. Ed. by B. J. Birch and W. Kuyk. Red. by A. Dold and B. Eckmann. Vol. 476. Springer Berlin Heidelberg, 1975, pp. 53–73. DOI: 10.1007/BFb0097583.

[Dev18] S. Devalapurkar. 'Wood's Theorem (Notes)'. 2018. URL: https://sanathdevalapurkar.github.io/files/wood.pdf. Draft.

[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.

[DFHH] C. Douglas, J. Francis, A. Henriques and M. Hill, eds. *Topological Modular Forms*. Vol. 201. Mathematical Surveys and Monographs. American Mathematical Society, 3rd Dec. 2014. DOI: 10.1090/surv/201.

[Dua+24] Z. Duan, H. J. Kong, G. Li, Y. Lu and G. Wang. '$RO(G)$-Graded Homotopy Fixed Point Spectral Sequence for Height 2 Morava E-Theory'.

Peking Mathematical Journal (27th May 2024). DOI: 10.1007/s42543-024-00087-7.

[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.

[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. '$\mathbb{C}$-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.

[Goe10] P. G. Goerss. 'Topological modular forms [after Hopkins, Miller, and Lurie]'. *Société mathématique de France*. Astérisque Séminaire Bourbaki : volume 2008/2009 exposés 997-1011 (2010), pp. 221–255. URL: http://www.numdam.org/item/AST_2010__332__221_0/.

[Goe+05] P. Goerss, H.-W. Henn, M. Mahowald and C. Rezk. 'A resolution of the $K(2)$-local sphere at the prime 3'. *Annals of Mathematics* 162.2 (1st Sept. 2005), pp. 777–822. DOI: 10.4007/annals.2005.162.777.

[Goe09] P. G. Goerss. 'Realizing Families of Landweber Exact Homology Theories'. In: *New Topological Contexts for Galois Theory and Algebraic Geometry (BIRS 2008)*. New Topological Contexts for Galois Theory and Algebraic Geometry. Mathematical Sciences Publishers, 16th June 2009, pp. 49–78. DOI: 10.2140/gtm.2009.16.49.

[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.

[Gre21] R. Gregoric. *Moduli Stack of Oriented Formal Groups and Periodic Complex Bordism*. 30th Nov. 2021. arXiv: 2107.08657v2 [math]. Preprint.

[HRW24] J. Hahn, A. Raksit and D. Wilson. *A Motivic Filtration on the Topological Cyclic Homology of Commutative Ring Spectra*. 3rd Oct. 2024. arXiv: 2206.11208v2 [math]. Preprint.

[Hed20] A. P. Hedenlund. 'Multiplicative Tate Spectral Sequences'. PhD thesis. University of Oslo, 2020. URL: https://www.mn.uio.no/math/personer/vit/rognes/theses/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.

[HHR16] M. Hill, M. Hopkins and D. C. Ravenel. 'On the Nonexistence of Elements of Kervaire Invariant One'. *Annals of Mathematics* 184.1 (1st July 2016), pp. 1–262. DOI: 10.4007/annals.2016.184.1.1.

[HL16] M. Hill and T. Lawson. 'Topological Modular Forms with Level Structure'. *Inventiones mathematicae* 203.2 (Feb. 2016), pp. 359–416. DOI: 10.1007/s00222-015-0589-5.

[Hop95] M. J. Hopkins. 'Topological Modular Forms, the Witten Genus, and the Theorem of the Cube'. In: *Proceedings of the International Congress of Mathematicians*. Ed. by S. D. Chatterji. Birkhäuser Basel, 1995, pp. 554–565. DOI: 10.1007/978-3-0348-9078-6_49.

[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.

[Isa+24] D. C. Isaksen, H. J. Kong, G. Li, Y. Ruan and H. Zhu. '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.

[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.

[Koc] S. O. Kochman. *Stable Homotopy Groups of Spheres*. Vol. 1423. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 1990. DOI: 10.1007/BFb0083795.

[Kon12] J. Konter. *The Homotopy Groups of the Spectrum Tmf*. 14th Dec. 2012. arXiv: 1212.3656 [math]. Preprint.

[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.

[EC2] J. Lurie. 'Elliptic Cohomology II: Orientations'. 26th Apr. 2018. URL: https://www.math.ias.edu/~lurie/papers/Elliptic-II.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.

[MR09] M. Mahowald and C. Rezk. 'Topological Modular Forms of Level 3'. *Pure and Applied Mathematics Quarterly* 5.2 (2009), pp. 853–872. DOI: 10.4310/PAMQ.2009.v5.n2.a9.

[Man21] L. Mantovani. *Localizations and completions of stable ∞-categories*. 5th May 2021. arXiv: 2105.02285v1 [math]. Preprint.

[Mar24] P. Marek. $\mathrm{H}\mathbb{F}_2$*-synthetic homotopy groups of topological modular forms*. 22nd Feb. 2024. arXiv: 2202.11305v2 [math]. Preprint.

[Mat16] A. Mathew. 'The homology of tmf'. *Homology, Homotopy and Applications* 18.2 (2016), pp. 1–29. DOI: 10.4310/HHA.2016.v18.n2.a1.

[MM15] A. Mathew and L. Meier. 'Affineness and Chromatic Homotopy Theory'. *Journal of Topology* 8.2 (June 2015), pp. 476–528. DOI: 10.1112/jtopol/jtv005.

[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/CBO9780511626289.

[Mei16] L. Meier. *Review of "The Homology of Tmf"*. 2016. URL: https://mathscinet.ams.org/mathscinet/article?mr=3515195.

[Mei17] L. Meier. 'Relatively Free TMF-modules'. 2017. URL: https://webspace.science.uu.nl/~meier007/RelativelyFree4.pdf. Preprint.

[Mei21] L. Meier. 'Topological Modular Forms with Level Structure: Decompositions and Duality' (24th July 2021). DOI: 10.1090/tran/8514.

[Mei22] L. Meier. 'Additive Decompositions for Rings of Modular Forms'. *Documenta Mathematica* 27 (2022), pp. 427–488. DOI: 10.4171/dm/874.

[Mei23] L. Meier. 'From Elliptic Genera to Topological Modular Forms'. Lecture Notes. 18th July 2023. URL: https://webspace.science.uu.nl/~meier007/TMF-Lecture.pdf.

[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.

[Mos70] R. M. F. Moss. 'Secondary Compositions and the Adams Spectral Sequence'. *Mathematische Zeitschrift* 115.4 (1970), pp. 283–310. DOI: 10.1007/BF01129978.

[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.

[Nau07] N. Naumann. 'The Stack of Formal Groups in Stable Homotopy Theory'. *Advances in Mathematics* 215.2 (Nov. 2007), pp. 569–600. DOI: 10.1016/j.aim.2007.04.007.

[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. Pstrągowski. *Adams Spectral Sequences and Franke's Algebraicity Conjecture*. 9th July 2023. arXiv: 2110.03669v2 [math]. Preprint.

[Pst22] P. Pstrągowski. 'Synthetic Spectra and the Cellular Motivic Category'. *Inventiones mathematicae* (20th Dec. 2022). DOI: 10.1007/s00222-022-01173-2.

[Pst25] P. Pstrągowski. 'Perfect Even Modules and the Even Filtration'. *Journal of the European Mathematical Society* (11th July 2025). DOI: 10.4171/jems/1669.

[PV22] P. Pstrągowski 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`.

[Rez07] C. Rezk. 'Supplementary Notes for Math 512 (Version 0.18)'. Lecture Notes. 19th July 2007. URL: `https://rezk.web.illinois.edu/512-spr2001-notes.pdf`.

[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.

[Sil] J. H. Silverman. *The Arithmetic of Elliptic Curves*. 2nd ed. Graduate Texts in Mathematics 106. Springer, 2010. 513 pp.

[ext-rs] Spectral Sequences Project. *ext-rs*. GitHub repository. 2024. URL: `https://github.com/SpectralSequences/sseq/tree/master/ext`.

[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`.

[WX17] G. Wang and Z. Xu. 'The Triviality of the 61-Stem in the Stable Homotopy Groups of Spheres'. *Annals of Mathematics* 186.2 (1st Sept. 2017). DOI: 10.4007/annals.2017.186.2.3.

[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.

Nederlandse samenvatting

Het doel van dit proefschrift is om een belangrijke stelling (bekend als de *Gap Theorem*) te bewijzen over wiskundige objecten die bekend staan als *topologische modulaire vormen*. Gezien het karakter en de beperkte lengte van deze samenvatting, is een grondige uitleg van wat deze objecten zijn niet mogelijk. In plaats daarvan zullen we kijken naar de context waarin deze objecten tevoorschijn komen, en het wiskundige gereedschap dat in dit proefschrift gebruikt wordt om deze stelling te bewijzen.

Topologie

Topologische modulaire vormen komen voor in de specialisatie *homotopietheorie*, soms ook *topologie* genoemd.[1] Homotopietheorie bestudeert meetkundige objecten door te doen alsof ze allemaal uit te rekken en te kneden zijn. Een eerste voorbeeld: in homotopietheorie is een koffiekop niet te onderscheiden van een donut, omdat we het kopje kunnen kneden tot een steel met een oor, wat vervolgens tot een donut kan worden omgekneed. Een donut is echter wel een andere vorm dan een bol: we zouden een donut enkel kunnen omvormen tot een bol door het gat dicht te maken. Een gat is daarmee een essentiële eigenschap wat topologie betreft.

Misschien vraagt de lezer zich af waarom dit een interessante sport is om te bedrijven. De reden is dat in veel verschillende gebieden van de wiskunde zich problemen voordoen die makkelijker te begrijpen zijn vanuit dit perspectief. Zulke voorbeelden kom je ook tegen in de echte wereld. Als je je fiets aan de fietsenstalling wilt vastmaken, moet je een ketting door het wiel en frame van je fiets doen. De precieze locatie van de ketting maakt dan niet uit, zolang deze maar door het wiel en het frame heengaat. De vraag of je fiets veilig is tegen dieven, is in deze zin dus een topologisch vraagstuk.

[1]Strikt genomen zijn dit verschillende disciplines, maar in deze samenvatting zullen we dit onderscheid negeren.

Invarianten

De donut en de bol zijn slechts de meest eenvoudige voorbeelden van meetkundige vormen. Meestal bestuderen wiskundigen objecten in hogere dimensies, en soms zelfs oneindig veel dimensies. Dit gaat het verbeeldingsvermogen te boven, en vereist daarom andere methoden om deze objecten toch te kunnen 'zien'. Hét middel naar keuze is om *algebra* hiervoor te gebruiken. Om een voorbeeld te geven: als een topoloog wil bewijzen dat een bol en een donut verschillend zijn, rekent hij van beide uit hoeveel 'gaten' deze hebben; deze hoeveelheid wordt het *geslacht* genoemd. Dit getal verandert niet als we de vorm uitrekken of krimpen; we noemen zo'n getal daarom een *invariant*. Met een beetje moeite kun je uitrekenen dat een bol geslacht nul heeft, en een donut geslacht één. Dit komt natuurlijk overeen met onze intuïtie, maar vertelt ons in dit geval ook een wiskundig feit: de bol is niet om te vormen in een donut (en vice versa), simpelweg omdat $0 \neq 1$. Het vergelijken van twee getallen is een stuk eenvoudiger dan twee meetkundige objecten direct vergelijken. Dit maakt deze techniek tot een belangrijk instrument voor een topoloog.

Hierop voortbouwend hebben topologen geprobeerd om andere invarianten te verzinnen. Eén manier om een krachtigere invariant te maken, is om niet te werken met getallen, maar met ingewikkeldere algebraïsche objecten. Zo kunnen moderne invarianten een heel *getallensysteem* als output hebben. Een getallensysteem is een wereld waarin getallen leven. Getallen zoals $0, 1, 2$, etc., en ook $-1, -2$, etc., leven in het systeem van de *gehele getallen*, aangeduid met $\mathbf{Z}$ (naar het Duitse woord *Zahlen*, oftewel getallen). De breuken, zoals $\frac{1}{2}$ of $-\frac{3}{7}$, leven niet in $\mathbf{Z}$, maar in het grotere systeem van de *rationale getallen*, aangeduid met $\mathbf{Q}$ (naar het Duitse *Quotient*, oftewel breuk). Er blijken nog veel meer getallensystemen te bestaan naast deze (waarschijnlijk bekende) systemen. Deze komen overal voor in de wiskunde, en voor topologie in het bijzonder worden ze dus veelal gebruikt als de output van ingewikkeldere invarianten.

Topologische modulaire vormen

Modulaire vormen zijn objecten die al decennia lang worden bestudeerd en gebruikt in met name de getaltheorie. Het zijn transformaties van het tweedimensionale vlak met speciale symmetrieën, die mooie plaatjes opleveren in een poging om ze weer te geven. In de jaren 90 heeft de beroemde wiskundige Mike Hopkins, samen met zijn coauteurs, een variant ontdekt genaamd *topologische modulaire vormen*. Deze vormen een soort combinatie tussen getaltheorie en topologie, maar zijn tegelijkertijd mysterieuze objecten: we snappen nog steeds niet goed wat een individuele topologische modulaire vorm precies moet voorstellen. Wat we wel enigszins begrijpen, is de collectie van alle topologische modulaire vormen tezamen. Samen maken deze deel uit van een oneindigdimensionaal object; de wiskundige naam voor dit object is Tmf, naar de Engelse benaming *topological modular forms*. Aangezien hier geen plaatjes meer van te tekenen zijn, stelt een topoloog zich als eerste de vraag: welke invarianten van Tmf kunnen we uitrekenen?

Dit brengt ons bij het belang van het bestaan van Tmf: het valt precies in het midden in termen van hoe complex het is, en hoeveel we erover kunnen uitrekenen. Een algemeen principe is dat, hoe ingewikkelder een object is, des te lastiger het is om invarianten ervan uit te rekenen, maar ook des te interessanter de informatie is die in deze invarianten bevat zitten. Het wonder van Tmf is dat veel invarianten volledig uit te rekenen zijn, en dat deze toch ongelofelijk veel informatie bevatten. Deze informatie kan vervolgens gebruikt worden om andere objecten te bestuderen en vragen van tientallen jaren oud te beantwoorden. Sinds de ontdekking van Tmf hebben deze ideeën geleid tot vele doorbraken in topologie.

Toch zit er een addertje onder het gras, welke het onderwerp is van dit proefschrift. De wiskundige literatuur over Tmf is namelijk incompleet: het bewijs van deze berekeningen is niet aan elkaar te plakken zonder een cirkelredenering te introduceren. Bij de ontdekking van Tmf in de jaren 90 zijn veel bewijzen niet gepubliceerd in onderzoeksartikelen; pas in de twee daarop volgende decennia zijn delen van de theorie opgeschreven in artikelen en boeken. Echter zijn niet alle nodige puzzelstukjes hier opgeschreven, en verhouden al deze artikelen zich op verschillende wijzen tot elkaar. Uit een nauwkeurige analyse blijkt dat uit de geschreven literatuur geen lineair verhaal op te bouwen is. Zeker gezien het belang van Tmf mag dit niet zo blijven.

In dit proefschrift wordt deze berekening op vaste grond geplaatst, door een nieuw bewijs te geven dat deze cirkelredenering omzeilt. De hoofdstelling is Stelling A (zie de Inleiding van dit proefschrift), bekend als de *Gap Theorem*; grof gezegd houdt deze stelling in dat het object Tmf geen gaten van een specifiek aantal dimensies heeft. Het nieuwe inzicht dat gebruikt is om deze stelling te bewijzen, is het stukje wiskundig gereedschap dat in de laatste paar jaar is ontwikkeld: de zogeheten *synthetische objecten*.

Synthetische objecten

Synthetische objecten[2] zijn, in zekere zin, fictieve objecten, ontworpen enkel en alleen om een hulpmiddel te zijn in topologische berekeningen. Het is een zeer recente theorie: de eerste aanwijzingen van het bestaan ervan verschenen rond 2014, en het werd daadwerkelijk uitgewerkt door Piotr Pstrągowski in 2018, met verdere ontwikkelingen sindsdien.

[2]Meer precies, synthetische *spectra*, maar een uitleg van wat een *spectrum* is in deze context zou wederom niet passen in deze samenvatting.

Nieuwe objecten verzinnen is niets nieuws in de wiskunde. De negatieve getallen zijn een bekend voorbeeld hiervan. Als je iemand zou vertellen dat je -1 appels in de fruitmand hebt liggen, word je waarschijnlijk raar aangekeken. Het getal -1 is, in zekere zin, een wiskundige fictie, in het leven geroepen met als enige doel om nul te worden als je er één bij optelt. Desalniettemin zijn negatieve getallen enorm nuttig. Synthetische objecten spelen een soortgelijke rol: het zijn fictieve objecten die bedoeld zijn om 'echte' meetkundige vormen beter te begrijpen.

Een synthetisch object kun je je een beetje voorstellen als een bouwwerk bestaande uit verschillende lagen bovenop elkaar gestapeld. Houd hierbij wel in het achterhoofd dat dit slechts een vergelijking is; deze 'stapel' is niet wezenlijk meetkundig, maar enkel een intuïtieve beschrijving van hoe deze wiskunde werkt. De eerste laag is het meest eenvoudig (en is wiskundig gezien volledig algebraïsch te begrijpen), en elke volgende laag maakt het object ingewikkelder. Als je vervolgens 'van bovenaf' kijkt, zie je de verschillende lagen niet meer, en blijft enkel een ingewikkeld patroon over. Dit patroon is het meetkundige object dat je hoopt te bestuderen, en het synthetische object heeft deze dus 'opgeknipt' in stukken. Wat dit nuttig maakt, is dat heel precies te analyseren is hoe elke volgende laag het resultaat verandert.

Met deze nieuwe theorie blijkt het mogelijk te zijn om een direct bewijs te geven van de *Gap Theorem* over Tmf — herinner dat deze stelling claimt dat Tmf geen gaten heeft van een aantal specifieke dimensies. Dit gaat als volgt. Allereerst moet er een synthetisch object gebouwd worden dat Tmf in 'laagjes' opdeelt. Dit is een van de belangrijke constructies in dit proefschrift; het resultaat is een object dat we *synthetische modulaire vormen* noemen, en aanduiden met Smf. Om in de vorige vergelijking te blijven, is Smf een object dat bestaat uit (oneindig veel) 'laagjes', maar wanneer je het object als geheel beschouwt, zie je het object Tmf weer verschijnen. De eerste laag is, met een beetje werk, volledig te begrijpen. Indien die eerste laag geen 'gaten' zou hebben, zouden we meteen de *Gap Theorem* bewezen hebben, maar dit is niet het geval. Het echte werk is om vervolgens te kijken hoe de verdere lagen die je eraan plakt, de gaten doen verdwijnen. Het blijkt dat na het plakken van 23 lagen, alle gaten die de *Gap Theorem* claimt, ook daadwerkelijk verdwenen zijn. Synthetische objecten werken zo dat verdere lagen geen nieuwe gaten kunnen introduceren, en daarmee is de stelling bewezen.

Het proefschrift is in twee delen opgedeeld. Het eerste betreft een uitleg van wat synthetische objecten zijn, gericht op topologen. Niet alleen biedt dit eerste gedeelte de nodige voorkennis om het bewijs te begrijpen, maar ook draagt het bij aan de wiskundige literatuur over synthetische objecten. Gezien het een recent onderwerp is, zijn er vooralsnog weinig leermaterialen voor; de hoop is dat dit proefschrift een bijdrage levert hieraan. Het tweede gedeelte van het proefschrift bevat zowel de constructie van Smf, als het bewijs van de *Gap Theorem*.

Acknowledgements

It is impossible to do justice to all the help and encouragement I have received over the past four years, but I will try nonetheless.

First and foremost: Lennart, thank you for all of your time, patience, advice, support, and dedication. From the extensive feedback you gave on everything I sent you, to the way you even once answered an email from a hospital bed, your way of supervising has left a lasting impression on me. It was an honour to be able to ask you so many questions, to learn the ropes of doing research directly from you, and to have you looking over my shoulder throughout the whole process. In more ways than you might realise, the final thesis feels to me like a collection of insights I learned from you. I hope to continue learning from you in the years to come.

Ieke, thank you for being my promotor and for your interest in my research. I learned a lot from our weekly meetings near the end of the project, and I always found the yearly meetings to be very reassuring and comforting.

Thank you to the committee members for reading my thesis, particularly in light of my painfully-apparent inability to use few words. I hope you can forgive me.

Christian and Jack, thank you for everything that you taught me, for the way you mentored me, and for the example you both set by being so free with your time. I have many fond memories of our collaborations, including a few delightful *eureka* moments, and the week of crunch-writing during a conference in Bonn. Both of you welcomed me with open arms in Bonn, and I'm excited for the years ahead.

The eighth floor of the Hans Freudenthal building truly was the place to be, and one that I will miss dearly. To all of my office mates: thank you for fostering such a pleasant working environment. Anna and Miguel, thank you for the many office discussions; whether during morning coffee or at unpredictable moments during the day, whether serious or light-hearted, each one played an essential part in keeping me sane. Anna, thank you as well for your help in administrative matters and thesis writing advice. Max, thank you for sharing and teaching me so much mathematics; it's hard to believe it's been nine years at this point. Miguel, thank you for all the inimitable humour you brought to the office, and to Utrecht as a whole (though my tendency to be loud in my enjoyment of it annoyed Niall, and probably others, every now and then). Ryan, thank you for your encouragements.

So much of my enjoyment of the past years is due to the amazing topology group at Utrecht, and in the Netherlands as a whole. Gijs, thank you for bringing me to Utrecht, for your advice, and for the inspiring mathematics you taught me over the past six years. To all the PhD's and postdocs: thank you for the lunch conversations, the company on conferences all over the continent, your encouragement, the mathematics I learned from you, and so much more. It's a long list, but that is half the fun. From Utrecht: thank you Christian, Guy, Jack, Jaco, Josefien, Léonard, Maite, Marco, Max, Miguel, Niall, Remy, Tobias, Vignesh, and Yuqing. From Nijmegen and Amsterdam: thank you Ahina, Alba, Grégoire, Grigori, Julius, Maxime, Miguel, and Tommy. Particular thanks to Guy and Niall for making sure I was okay when I was feeling particularly discouraged, and to Christian, Guy and Niall for their help when I was applying for postdocs.

I owe a lot to the homotopy theory community at large. So much about synthetic spectra I learned from Robert Burklund, Irakli Patchkoria, and Piotr Pstrągowski; thank you for answering my questions, for the inspiring papers you wrote, and for the interest you showed in my work, which was incredibly encouraging. Large portions of this thesis grew out of my attempts to better understand the Omnibus Theorem of Burklund–Hahn–Senger; their work has been very influential. I am very thankful for my conversations and (ongoing) collaborations with Torgeir Aambø, Shaul Barkan, Marius Nielsen, and Maxime Ramzi; thank you for being so open in sharing ideas. Thank you to Miguel Barata, Bastiaan Cnossen, and Marius Nielsen for reading and giving feedback on Part I of the thesis. I would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for their support and hospitality during the programme *Equivariant homotopy theory in context* (supported by EPSRC grant EP/Z000580/1), where parts of this thesis were written.

Dan zijn er nog een paar mensen die, als ik het goed inschat, liever in het Nederlands worden aangesproken. Vincent, dank je wel voor het ontwerpen van de omslag. Ik heb erg genoten van onze gesprekken erover, en was onder de indruk van hoe goed je door de materie heen prikte.

Mam, pap, Michelle. Jullie zijn er altijd voor me geweest, door alle pieken en dalen van de afgelopen jaren heen. Zoals jullie weten vond ik wiskunde van jongs af aan al interessant, maar dat ik er mijn werk van heb kunnen maken, heb ik aan jullie te danken. Naast het aanmoedigen van mijn interesses, hebben jullie me vanaf het begin — en zonder enige aarzeling — volledig gesteund in mijn studie. Al kon ik moeilijk uitleggen waar ik van dag tot dag nou precies mee bezig was, kon ik bij jullie altijd vinden wat echt belangrijk was. Al die gezelligheid, rust, humor (vooral die van jou, zus) en geborgenheid betekenen alles voor me. Dank jullie wel.

Curriculum vitae

Sven van Nigtevecht was born on the 10th of October 1998 in Veldhoven. After finishing high school in 2016, he enrolled in the double Bachelor's program in mathematics and physics at the University of Amsterdam. His Bachelor's thesis on topological phases and K-theory was supervised by Jan de Boer and Hessel Posthuma. He obtained both degrees *cum laude* in 2019. Afterwards, he started a Master's degree in mathematics, also at the University of Amsterdam, which he obtained *cum laude* in 2021. His Master's thesis on unstable homotopy theory was supervised by Gijs Heuts from Utrecht University. The next four years he spent as a PhD student at Utrecht University supervised by Lennart Meier, resulting in the present thesis. Starting in October 2025, he will start a postdoc position at the University of Bonn under Markus Hausmann.