A

Stack-theoretic perspective

on

KU-local stable homotopy theory

by

Jeroen W. van der Meer

Universiteit Leiden

A **master thesis** supervised by

Lennart Meier

co-supervised by

Bas Edixhoven

To be defended on July 05, 2019
Mathematisch Instituut
Universiteit Leiden

Contents

Preface

Goal

Stable homotopy theory is the branch of mathematics that studies phenomena in homotopy theory that arise after repeated application of the suspension operation. Classical examples of such phenomena include the Freudenthal suspension theorem (homotopy groups of spheres stabilize after applying the suspension functor sufficiently many times), the suspension isomorphism (reduced cohomology of a space and of its suspension coincide), and Bott periodicity (K-theory is periodic under the suspension operation).

In the same way that algebraic topologists do much of their work in the category of spaces, homotopy theorists studying stable phenomena often work in the stable homotopy category. Roughly speaking, this category is what we obtain when we take the homotopy category of topological spaces, and force the suspension functor to be an equivalence of categories. We are left with only those phenomena that survive repeated applications of the suspension functor, which should be precisely the stable phenomena that the homotopy theorists are after.

The stable homotopy category is rather complicated, and in order to study it, it is helpful to zoom in a bit further and restrict attention to more digestible pieces of this category. This approach has been quite successful, and its success can be seen, for instance, in the chromatic approach to stable homotopy theory. This approach asserts that structures in the stable homotopy category should loosely correspond to structures in the category of sheaves over a particular space that we call the stack of formal groups. This stack of formal groups admits a filtration, and the chromatic point of view dictates that this ought to translate to a filtration of pieces of the stable homotopy category, each piece more complicated than the previous one, but all of them simpler than the entire thing.

For the most part, this thesis will confine itself to only one digestible piece: we will be interested in the aspects of the stable homotopy category that can be detected by K-homology, or the K-local stable homotopy category for short. Building on earlier work due to Adams, this K-local category was first systematically studied several decades ago by Alridge Bousfield who, in 1979, published a paper, [2], in which he gave an algebraic classification of the objects in the K-local stable homotopy category (or, of the K-local spectra for short, in the same way that objects of the stable homotopy category are called spectra), at least when further localized at an odd prime.

The paper of Bousfield is almost entirely algebraic. The fact that we can reduce questions in homotopy theory to algebra is essentially thanks to a tool called the Adams spectral sequence. In the special case of the K-local category, this tool roughly tells us the following: given two spaces X and Y, if one understands enough of the (purely algebraic!) structure of $\mathrm{Hom}(\mathrm{K}_*(X), \mathrm{K}_*(Y))$, where K_* means K-homology, then one can deduce what the homotopy classes of maps from X to Y should be in the K-local stable homotopy category.

Bousfield's paper predates the chromatic approach that we mentioned earlier. Recall that this approach studies a certain filtration of increasingly complicated pieces of the stable homotopy category. As it turns out, the very first piece of this filtration is strongly related to the K-local category that Bousfield was interested in, and when suitably defined, they are in fact the same thing. This brings us to the goal of this thesis: with the modern, chromatic language at our fingertips, we will shed a new light on the techniques used in the original paper, only to find that, at least in hindsight, the complicated algebraic structures that were initially considered are geometric in origin. In fact, they are but a minor reflection of the deep and mysterious geometric structures governing stable phenomena — structures which, for the most part, remain to be understood to this very day.

Outline and dependencies

In Appendix A we study formal group laws over commutative rings. We begin by introducing the basic definitions, and proceed to study heights of formal group laws, ending with a mostly independent section on endomorphisms of formal group laws. Most of the results are well-known, though we occasionally use non-standard terminology when discussing heights.

Chapter 1 serves as an introduction to the basic language of stable homotopy theory that we will need in the rest of this thesis. None of the results in this chapter are new, and the chapter is mostly included for the sake of bookkeeping, not to mention because of the lack of a suitable reference to which to refer.

In Chapter 2 we introduce Hopf algebroids and algebraic stacks. The purpose of this chapter is two-fold. First, it seems that Hopf algebroids are less familiar than algebraic stacks to most algebraic geometers, and vice versa for topologists, so that it is reasonable to give a combined discussion on both notions. Second, we use this chapter to fix definitions that may otherwise be ambiguous due to inconsistent conventions in the literature.

With the language of spectra and algebraic stacks at hand, we are ready to give a detailed study of the geometric structure of the stack of formal groups in Chapter 3. We explain how heights of formal group laws give rise to a filtration of the stack into substacks, and proceed to look into the geometry of each of the layers of this stack. We also discuss Landweber's exact functor theorem, allowing us to construct spectra in the stable homotopy category using maps into the stack of formal groups. We use this theorem to give detailed constructions of various spectra that carry the name 'Morava E-theory'. It is my hope that the statements and proofs in this chapter fill up some gaps in the mathematical literature.

Finally, in Chapter 4 we use the machinery developed in the previous chapter to study the K-local stable homotopy category. We start out by outlining the key constructions and result in Bousfield's paper. After that, we investigate the category of sheaves over certain substacks of the stack of formal groups; by that point we will know that this category is intimately related to the K-local category that we are ultimately interested in. Finally, we use the conclusions of our investigation to revisit Bousfield's results, and indicate how our results may shed a conceptually simplifying light on the original ideas.

In case the reader is interested only in some specific parts of this thesis, the following dependency diagram may be of help.

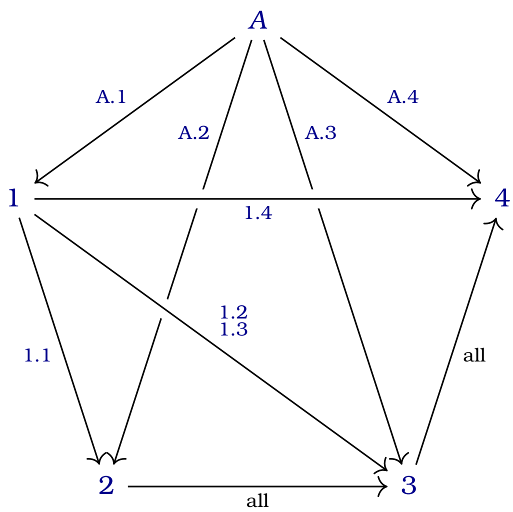

Conventions

Whenever we work with spaces, we assume that we are working inside some convenient category of topological spaces, such as CW complexes or compactly generated weak Hausdorff spaces. At no point will the precise choice of category of spaces be relevant. Likewise, whenever we work with spectra, we will be working in the usual stable homotopy category. Apart from a few constructions in Section 1.2, a particular choice of model for the stable homotopy category will never be relevant. This means that the reader unfamiliar with stable homotopy theory can safely assume the existence of the category as a black box without much loss of continuity.

In order for our story to fit within the framework of classical algebraic geometry, all our rings are assumed to be commutative and with unity. In practice, stable homotopy groups of commutative ring spectra are *anti*-commutative graded rings rather than being truly commutative rings. However, all spectra that we are interested in (most notably, KU, $E(n)$ and MU) are evenly graded, so that our conventions will not cause any problems.

The algebraic stacks that we define in Section 2.3 are rather general: they are fpqc stacks fibred in groupoids over the category of (affine) schemes, without any assumptions on the diagonal, or on the existence of an fpqc atlas. We have chosen this convention so that all algebro-geometric objects that we encounter will fit within our framework. As a special case, we define Adams stacks to be algebraic stacks with affine diagonal admitting an fpqc atlas, and Noetherian stacks to be algebraic stack admitting an fppf atlas from an affine Noetherian scheme, both of which are at a level of generality closer in spirit to the usual theory of algebraic stacks that one encounters in algebraic geometry.

We will often be working with formal group laws, particularly in Appendix A. For us, a formal group law will always mean a commutative, one-dimensional formal group law over a commutative ring with unity. There are also various generalizations of formal group laws that essentially boil down to things that are (Zariski- or fpqc-)locally a formal group law. We will never use these concepts.

Finally, we remark that set-theoretical issues may or may not arise at some points, so just to be certain, we invoke the usual disclaimer mentioning universes and how it ultimately doesn't matter.

Acknowledgements

I would like to thank my supervisor, Lennart Meier, for the discussions we've had. I could not have written this thesis without his valuable aid and advice throughout the last months.

A final word

Should the reader have something interesting to remark — say, a question or comment related to the content, an error in the proofs, or neat results which are relevant to the topic of this thesis, feel free to send me an e-mail at meer@vivaldi.net.

Chapter 1

Spectra

The primary goal of this chapter is to introduce some well-known ideas and methods from stable homotopy theory. We begin by defining the stable homotopy category, and more importantly, by describing some of its properties. We then give some examples of objects in this category, particularly those that will become relevant later in this thesis.

1.1 The stable homotopy category

The stable homotopy category, which we denote by $\mathrm{Ho}(\mathsf{Sp})$, may be thought of as the stabilization of the classical homotopy category of topological spaces under the suspension operation. It provides us with a crucial framework in which much of modern stable homotopy theory is being developed. We start out in this section by stating many basic properties that $\mathrm{Ho}(\mathsf{Sp})$ satisfies. We then proceed to give, without proofs, an explicit construction of this category, via the language of CW spectra. Many other more sophisticated constructions of the category exist, but for us this one construction will suffice — in fact, the fact that the category $\mathrm{Ho}(\mathsf{Sp})$ exists in the first place, without having an explicit construction of it, will often suffice.

To the reader who is not familiar with the stable homotopy category, let me also kindly refer you to [6] and [12], both of which have been of great help when I first learned about this, and on which much of the contents of this section are based. Without further ado, let's start by listing the desired properties.

There exists a functor Σ^∞ from the pointed homotopy category $\mathrm{Ho}(\mathsf{Top}_*)$ of spaces to the stable homotopy category $\mathrm{Ho}(\mathsf{Sp})$, meaning that spaces can be viewed as objects in $\mathrm{Ho}(\mathsf{Sp})$. This functor Σ^∞ is fully faithful, and admits a right adjoint, denoted Ω^∞. Given a space X, the object $\Sigma^\infty X$ is called the **suspension spectrum** of X. More generally, objects in $\mathrm{Ho}(\mathsf{Sp})$ are called **spectra**.

There exists an *equivalence* of categories $\Sigma\colon \mathrm{Ho}(\mathsf{Sp}) \to \mathrm{Ho}(\mathsf{Sp})$, called the **suspension functor**; its inverse equivalence is denoted Ω and called the **loop space functor**. They commute with the classical suspension and loop space functors on spaces in the sense that we have a commutative diagram

$$\begin{array}{ccc} \mathrm{Ho}(\mathsf{Top}_*) & \underset{\Omega}{\overset{\Sigma}{\rightleftarrows}} & \mathrm{Ho}(\mathsf{Top}_*) \\ {\scriptstyle \Omega^\infty}\uparrow\downarrow{\scriptstyle \Sigma^\infty} & & {\scriptstyle \Omega^\infty}\uparrow\downarrow{\scriptstyle \Sigma^\infty} \\ \mathrm{Ho}(\mathsf{Sp}) & \underset{\Omega}{\overset{\Sigma}{\rightleftarrows}} & \mathrm{Ho}(\mathsf{Sp}) \end{array}$$

The suspension and loop space functors should be regarded as the stable analogue of the classical suspension and loop space functors, as both the names and the commutativity of the above diagram already suggest. The fact that these functors are now equivalences may well be regarded as the sole distinguishing feature of the stable homotopy category. One consequence of this feature is that negative dimensions make sense. To illustrate this, if we define the **sphere spectrum** $\mathbb{S}$ to be $\Sigma^\infty S^0$, then for all $n \geq 0$, $\Sigma^n\mathbb{S} \cong \Sigma^\infty S^n$, but for $n < 0$, $\Sigma^n\mathbb{S}$ exist as well, and these essentially behave as negative-dimensional spheres.

The stable homotopy category is additive. Let us expand on what means. First and foremost, the Hom sets in $\mathrm{Ho}(\mathsf{Sp})$ carry abelian group structures. We often write $[E, E']$ for the set of maps from the spectrum E to the spectrum E', endowed with its group structure. Moreover, we often extend the abelian group $[E, E']$ into a *graded* abelian group $[E, E']_*$ containing $[E, E']$ as the 0-th level by defining $[E, E']_n = [\Sigma^n E, E']$. Second, there exists a zero object, often denoted $*$, which is in fact Σ^∞ applied to the one-point topological space. Third, the stable homotopy category admits finite coproducts and finite products, which coincide up to isomorphism; we denote the (co-)product of two spectra X and Y by $X \vee Y$, and refer to this product as the **wedge sum**. Actually, $\mathrm{Ho}(\mathsf{Sp})$ has all infinite products and coproducts, too, but then they need no longer coincide.

The enrichment is not surprising when you think of the equivalence Σ as the stable analogue of the classical suspension functor. Indeed, if we have two spaces X and Y, then the set of homotopy classes of maps from $\Sigma^2 X$ to Y carries a group structure, obtained by 'concatenating' (representatives of classes of) maps; this group structure is abelian by the usual shrink-and-move argument. But in the stable homotopy category, every spectrum X is the two-fold suspension of some other spectrum (namely of $\Omega^2 X$).

The stable homotopy category admits an internal Hom, and a smash product acting as its left adjoint, together turning the category into a closed symmetric monoidal category. We denote the smash product by $\wedge$, and the internal Hom by $F(\,\cdot\,,\,\cdot\,)$. The unit of the smash product is given by the sphere spectrum $\Sigma^\infty S^0$.

The symmetric monoidal structure on $\mathrm{Ho}(\mathsf{Sp})$ is compatible with the classical monoidal structure on $\mathrm{Ho}(\mathsf{Top}_*)$, in the sense that Σ^∞ is strong monoidal and Ω^∞ is (at least) lax monoidal. The monoidal structure on Ho(Sp) is compatible with the suspension and loop space functors too. In fact, we can *define* the suspension and loop space functors explicitly in terms of the symmetric monoidal structure, by setting

$$\Sigma X = (\Sigma^\infty S^1) \wedge X \qquad \text{and} \qquad \Omega X = F(\Sigma^\infty S^1, X)\,.$$

From this perspective, it is easy to see that we have natural isomorphisms

$$\Sigma X \wedge Y \cong \Sigma(X \wedge Y) \cong X \wedge \Sigma Y \qquad \text{and} \qquad \Omega F(X, Y) \cong F(\Sigma X, Y) \cong F(X, \Omega Y)\,.$$

Another feature of the stable homotopy category, which is not shared by the homotopy category of spaces, is that retracts are summands of the original space. More precisely, let X be a spectrum, and let A be a spectrum such that there exist arrows $A \to X \to A$ that compose to the identity. We also say that A is a **retract** of X. Then X contains A as a summand, in other words, there exists a spectrum B such that $X \cong A \vee B$.

There is a notion of homotopy groups in Ho(Sp), or more precisely, of *stable* homotopy groups. Given a spectrum E, we define the n-th **stable homotopy group** of E to be

$$\pi_n^{\mathrm{S}}(E) = [\Sigma^n \mathbb{S}, E]\,.$$

Notice that n may well be negative. We point out that the superscript S is often omitted from notation. If E is the suspension spectrum $\Sigma^\infty X$ of a based space X, then $\pi_n^{\mathrm{S}}(\Sigma^\infty X)$ is naturally isomorphic to the usual stable homotopy group of X. Of course, in that case, the negative homotopy groups will be trivial. More generally, any spectrum whose negative stable homotopy groups vanish is called a **connective spectrum**.

Recall that by Whitehead's Theorem, weak homotopy equivalences between CW complexes induce isomorphisms in the homotopy category. The stable homotopy category admits an analogous result.

Theorem 1.1.1 (Stable Whitehead Theorem). If $f\colon X \to Y$ is a morphism in Ho(Sp) whose induced maps on stable homotopy groups are isomorphisms of abelian groups, then f is an isomorphism in Ho(Sp). □

Yet another key aspect of the stable homotopy category is that the objects of Ho(Sp) define (reduced) homology and cohomology theories on the category of based spaces. (When we say "reduced (co-)homology", we mean functors admitting suspension isomorphisms, and satisfying homotopy invariance, the exactness property, *and* the wedge axiom.) Here's how that works. If E is a spectrum in Ho(Sp), then for any based space X, we define the groups

$$\widetilde{E}_n(X) = \pi_n^{\mathrm{S}}((\Sigma^\infty X) \wedge E) \cong [\mathbb{S}, (\Sigma^\infty X) \wedge E]_n$$

and

$$\widetilde{E}^n(X) = \pi_{-n}^{\mathrm{S}}(F(\Sigma^\infty X, E)) \cong [\Sigma^\infty X, E]_{-n}$$

for all n. These definitions satisfy the axioms of a generalized reduced homology and cohomology theory. A morphism of spectra in Ho(Sp) can be seen to induce a natural transformation between homology and cohomology functors.

A few remarks are in order. First, we may convert these reduced homology and cohomology theories in their unreduced versions in the usual way, in which case we write E_n and E^n rather than $\widetilde{E}_n$ and $\widetilde{E}^n$. Second, we point out that the definition of reduced (co-)homology extends in an obvious way to the case where X is not a space but a spectrum. Finally, if X is the zero object $*$ in Ho(Sp), then we have $E_n(*) = [\mathbb{S}, E]_n = \pi_n^S(E)$, and $E^n(*) = \pi_{-n}^S(E)$. We often write E_* and E^* to refer to these abelian groups.

So we can turn spectra in both homology and cohomology theories. What's more, this process can be reversed. That is, given a (co-)homology theory, there will be a spectrum representing it. The cohomological result is known as the Brown representability theorem, and one of its many versions states the following.

Theorem 1.1.2 (Brown Representability Theorem). Let $\widetilde{E}^*$ is a reduced cohomology theory satisfying the wedge axiom defined on the homotopy category of *connected* CW complexes. Then there exists a unique spectrum E such that $\widetilde{E}^n(X) = [\Sigma^\infty X, E]_{-n}$. Moreover, a natural transformation of two such reduced cohomology theories gives rise to a morphism of spectra in Ho(Sp). □

The homological analogue is due to Adams, and it is the same except that we need to impose an additional assumption on our homology functor, namely the direct limit axiom. While it may rightfully be called the Adams representability theorem, the term "Brown representability theorem" is also often used for the homological version.

Theorem 1.1.3 (Adams Representability Theorem). Let $\widetilde{E}_*$ be a reduced homology theory satisfying the wedge and direct limit axioms. Then there exists a unique spectrum E such that $\widetilde{E}_n(X) = [\mathbb{S}, (\Sigma^\infty X) \wedge E]_n$. Moreover, a natural transformation of two such reduced homology theories gives rise to a morphism of spectra in Ho(Sp). □

It would appear, then, that spectra 'are' (co-)homology theories. But there's one more subtlety that is worth mentioning. Morphisms in the stable homotopy category do not correspond exactly to natural transformations of (co-)homology functors. This is due to the existence of so-called phantom maps in the stable homotopy category. A **phantom map** is a non-trivial morphism $E \to E'$ of spectra such that the induced natural transformation on homology theories is zero. We won't go into this much further, but it is important to keep in mind that they exist.

The stable homotopy category carries the structure of a triangulated category. More precisely, let us call any string of maps of the form $X \to Y \to Z \to \Sigma X$ a **triangle**. Then there is a collection of so-called **distinguished triangles** in Ho(Sp) that satisfy the axioms of a triangulated category. In particular, every morphism $X \to Y$ is part of a distinguished triangle $X \to Y \to Z \to \Sigma X$.

The Z occurring in this distinguished triangle is sometimes also denoted by Y/X, and called the **cofibre** of the map $X \to Y$, and it does indeed carry the intuitive interpretation of a quotient, as is exemplified by the long cofibre and fibre sequences that arise from the distinguished triangles.

The smash product and the triangulation are compatible in the sense that smash products preserve cofibres. More precisely, if $X \to Y$ is a map with cofibre Y/X, and we smash the map with a spectrum E so as to obtain a map $E \wedge X \to E \wedge Y$, then the cofibre $(E \wedge Y)/(E \wedge X)$ is given by $E \wedge (Y/X)$. Another nice compatibility is the following. If A is a retract of X, then the natural 'inclusion' $A \to X$ is part of a distinguished triangle $A \to X \to X/A \to \Sigma A$. Recall that retracts are summands, so that $X \cong A \vee B$ for some B. It is reasonable, then, to expect that $B \cong X/A$, and indeed this is the case.

The last property that we mention is an imprecise but nonetheless fascinating one. Some of the properties stated above are very similar to those of the category of abelian groups. Both are additive, both satisfy the property that retracts are summands, and both have a closed symmetric monoidal structure. Based on this analogy, we make the following definition. In the same way that a (not necessarily commutative but still unital) ring is nothing but a monoid object in the monoidal category of abelian groups, we define a **ring spectrum** to be a monoid object in the stable homotopy category. More precisely, a ring spectrum R is a spectrum together with two morphisms, $m\colon R \wedge R \to R$ and $e\colon \mathbb{S} \to R$, such that the diagrams

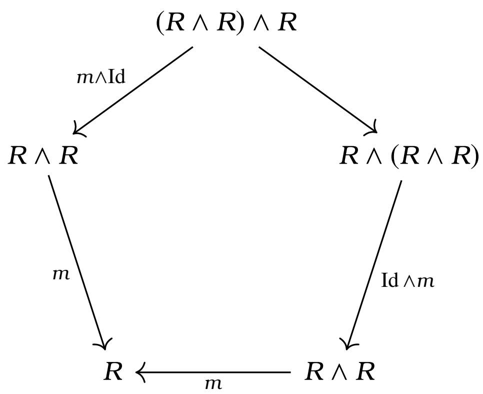

and

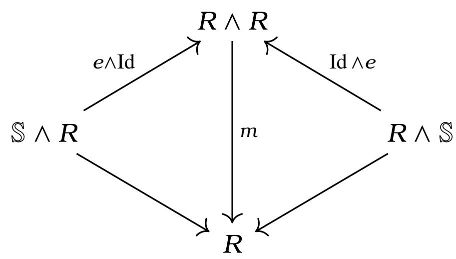

are commutative. A **morphism of ring spectra** is just a morphism of spectra such that all relevant structures are compatible. It is obvious what it means for a ring spectrum to be **commutative** as well. All ring spectra that we are interested in will be commutative.

Given a commutative ring spectrum E and a space X, the graded groups $\widetilde{E}^*(X) = [\Sigma^\infty X, E]_*$ inherit a multiplicative structure in the following way. Start with two elements of the cohomology groups. Smash them together, compose them with the diagonal $\Sigma^\infty X \to \Sigma^\infty X \wedge \Sigma^\infty X$ and with the multiplication $E \wedge E \to E$, and we end up with a new element. This structure turns $\widetilde{E}^*(X)$ into what unfunny people call a rng — a structure that is almost a (commutative) ring but lacks a multiplicative identity element. The unreduced version $E^*(X) = \widetilde{E}^*(X^+)$ does have a unit, namely the map $X^+ \to S^0$ sending X to the point that is not the base-point of S^0. With some more work, one can show that E^* defines a multiplicative cohomology theory. Conversely, if we have a multiplicative cohomology theory E^*, Brown representability (1.1.2) ensures that the structures lift to a structure on the spectrum E which turns it into a multiplicative cohomology theory.

For completeness, we remark that there is another definition of ring spectra which is not equivalent to the one presented above. The stable homotopy category is really the underlying homotopy category of an ∞-category Sp. Rather than asking for a spectrum to be a commutative monoid object in $\mathrm{Ho}(\mathsf{Sp})$, we can also ask for it to be a commutative monoid in Sp. Such spectra are sometimes called E_∞**-ring spectra**.

In the same way that modules over rings exist, we can define a **left module spectrum** over a ring spectrum R to be a module object over R. More precisely, it is a spectrum M together with a morphism $\rho\colon R \wedge M \to M$ such that the diagrams

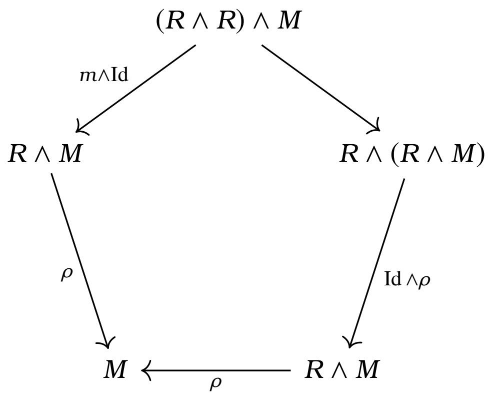

and

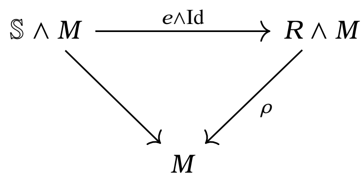

commute. There is obviously a notion of a **right module spectrum** as well, and if the ring spectrum is commutative, these two notions coincide.

We now turn to the task of constructing the stable homotopy category. There are various known approaches to this, all of them with their own advantages and disadvantages. Following [13, Ch. 8], we'll sketch one construction below, which is historically one of the first constructions and has the advantage of being reasonably simple.

A **CW spectrum** consists of based CW complexes $X_0, X_1, \ldots$, together with cellular inclusion maps $\Sigma X_n \hookrightarrow X_{n+1}$ turning ΣX_n into a subcomplex of X_{n+1}. With this definition, every $(k+n)$-cell of E_n may be identified with a certain $(k+n+1)$-cell of E_{n+1}, so that it makes sense to define a **stable** k**-cell** to be a $(k+n)$-cell of E_n, modulo said identification. Notice in particular that k can be negative, and recall that the stable homotopy category aims to make sense of negative dimensions. In the literature, CW spectra are also called CW prespectra, sequential spectra, or even just spectra. The terms have also been used for several very similar notions. For instance one may drop the requirement that the structure maps are cellular inclusions, one might allow the spaces to be arbitrary topological spaces, or one might also want to have negatively graded spaces. Such changes usually do not make a difference, though this is not always easy to prove.

The morphisms that we are interested in are *not* what you might expect. Our definition of a **morphism of CW spectra** will be a bit sketchy to keep it short, but hopefully the idea is conveyed. A morphism $f\colon X \to Y$ is defined as follows. For each stable m-cell on X, there should be an N such that for all $n \geq N$, we have an explicit continuous map from the $(m+n)$-cell on X_n realizing the stable m-cell, to Y_n. We want the continuous maps to be compatible with the structure maps $\Sigma X_n \to X_{n+1}$, in the sense that if f defines a map on an $(m+n)$-cell of X_n, then the map on the corresponding $(m+n+1)$-cell of X_{n+1} should be the suspension of the map of the $(m+n)$-cell. So in short, we could say that maps of CW spectra are 'eventually-defined maps'. We end up with a category CW–Sp of CW spectra, where one should think for a moment to convince oneself that compositions are well-defined.

There is a notion of homotopy of maps, defined as follows. Given two maps $f, g\colon X \to Y$ of CW spectra, a homotopy should be a map of CW spectra from the CW spectrum $(X \wedge I_+)_n = X_n \wedge I_+$ to Y, in such a way that it restricts to the maps f and g at the end-points. That is, upon composing with the two inclusions $X \to X \wedge I_+$, we get f and g back. One can verify that this is an equivalence relation.

We now define a category Ho(CW–Sp) whose objects are CW spectra, and whose morphisms are *homotopy classes* of maps. Here too one needs to spend some quality time with the definition of compositions. This category is equivalent to the stable homotopy category, or if you wish, it defines it. In particular, it satisfies all the desired properties we discussed before.

The reader should feel free to try and verify some of the properties. This immediately brings us to the downside of this particular construction. While it is arguably the easiest conception of the stable homotopy category, it is also very hard to work with. For instance, if the reader has indeed tried to prove that compositions are well-defined, he will have found that properties such as associativity are tedious to verify. The same is true for many of the other properties.

A slightly more sophisticated but equally important issue with this construction is the lack of a decent smash product. The category Ho(CW–Sp) is closed symmetric monoidal (as was one of our desired properties), but this structure does not come from a symmetric monoidal structure on CW–Sp. Ultimately, when trying to prove things, it is important to be able to refer back to a concrete point-set model such as CW–Sp, and it would then be desirable to have the smash product directly available.

There are other constructions which do not suffer from the aforementioned disadvantages that our construction suffers, but which are technically more difficult to construct. For us, the disadvantages stated above will not pose an issue. We will simply take for granted that the category of CW spectra satisfies all properties we wanted, and we will never need a point-set description of a smash product of two spectra. We will therefore not consider any of the other constructions.

1.2 First examples of spectra

We consider a few examples that we will need later on. If he wishes, the reader can find many other examples in [12].

Example 1.2.1. The most important family of example was already considered in the previous section. To every based space X we can associate a suspension spectrum, which we denote by $\Sigma^\infty X$. The homotopy groups $\pi_*^S(\Sigma^\infty X)$ are precisely the stable homotopy groups of the space X.

Example 1.2.2. Recall that objects in the stable homotopy category may be interpreted as (generalized) cohomology theories. The standard example of cohomology is ordinary cohomology, with coefficients in some abelian group G. It is represented by a spectrum, called the **Eilenberg–MacLane spectrum**, denoted HG. If we take our coefficients not in an abelian group G but a commutative ring R, the cohomology inherits a cup product, and this implies that the spectrum HR is actually a commutative ring spectrum.

Recall that Eilenberg–MacLane spaces, from which these spectra are constructed, have only one non-trivial homotopy group. Using this, one can deduce that $\pi_*^S(HG)$ is concentrated in degree 0, and $\pi_0^S(HG) = G$.

Example 1.2.3. The next spectrum will be used as an example in Section 1.4. Start out with the identity map $\mathbb{S} \to \mathbb{S}$ in the stable homotopy category. As this category is additive, we can add up this map p times so as to yield what we call the multiplication-by-p map $\mathbb{S} \xrightarrow{\times p} \mathbb{S}$. The cofibre of this map (or equivalently, Σ^{-1} of the fibre) is what we call the **mod-p Moore spectrum**, denoted $\mathbb{S}/p$, but also $S\mathbb{Z}/p$ or $M\mathbb{Z}/p$. The cohomological functor $\mathrm{Hom}(\mathbb{S}, \cdot)$ gives rise to a long fibre sequence on homotopy groups, from which we can deduce that $\pi_0^S(\mathbb{S}/p) = \mathbb{Z}/p$, and that, for negative n, $\pi_n^S(\mathbb{S}/p)$ is trivial. On the other hand, let's take our distinguished triangle $\mathbb{S} \xrightarrow{\times p} \mathbb{S} \to \mathbb{S}/p \to \Sigma\mathbb{S}$, and this time smash it with the Eilenberg–MacLane spectrum $H\mathbb{Z}$. As mentioned earlier, the triangulation is compatible with smash products, so we end up with another distinguished triangle. The same procedure yields another long fibre sequence, this time on homology groups because $[\mathbb{S}, X \wedge H\mathbb{Z}]_* = H_*(X; \mathbb{Z})$ by definition. By investigating this long exact sequence, one easily finds that, for positive n, $H_n(\mathbb{S}/p; \mathbb{Z}) = 0$.

Example 1.2.4. More generally, if G is an abelian group, then the **Moore spectrum** SG or MG is a particular spectrum characterized by the properties $\pi_0(SG) = G$, $\pi_{<0}(SG) = 0$, and $H_{>0}(SG; \mathbb{Z}) = 0$. The above discussion gave a construction in the special case where $G = \mathbb{Z}/p\mathbb{Z}$. For more general Moore spectra, constructions become more difficult, but it suffices for us to know that they exist.

Example 1.2.5. The next example we'll discuss is complex K-theory. Let X be a compact Hausdorff and connected topological space. Define the **complex K-theory** of X, denoted, $\mathrm{KU}^0(X)$, to be the group completion of the monoid of isomorphism classes of finite-dimensional complex vector bundles on X.

Lemma 1.2.6. With the notation as above, we have an equality $\mathrm{KU}^0(X) \cong [X, \mathrm{BU} \times \mathbb{Z}]$. □

Proof: As X is compact Hausdorff, any map $f\colon X \to \mathrm{BU}\times\mathbb{Z}$ factors through some $\mathrm{BU}(n)\times\{m\}$. (Indeed, suppose this weren't true. Then we could pick a point $x_k \in f(X)\cap \mathrm{BU}(k)\times\{m\}$ for infinitely many k. The space BU being normal, this would yield an infinite discrete subset of $f(X)$, which is impossible by compactness.) So we can associate an n-dimensional complex vector bundle V, and a number m, to the homotopy class $[f]$ of f. In turn, we send the pair (V, m) to the pair $V - \varepsilon^{n-m}$ in $\mathrm{KU}^0(X)$, where ε^N denotes the trivial vector bundle of rank N. This defines a map $[X, \mathrm{BU} \times \mathbb{Z}] \to \mathrm{KU}^0(X)$. This map is easily seen to be injective.

To prove surjectivity, we first point out that any complex vector bundle over a compact Hausdorff space is a summand of some trivial bundle. That is, if W is a complex bundle on X, then there is some W' so that $W \oplus W'$ is isomorphic to the trivial bundle ε^N of some rank N. We use this as follows. Take any element $V - W$ of $\mathrm{KU}^0(X)$. Take a bundle W' so that $W \oplus W' \cong \varepsilon^N$. Now $V - W = (V \oplus W') - \varepsilon^N$. Write $m = \mathrm{rank}(V \oplus W') - N$. The element $V - W$ is now seen to be the image of a map $f\colon X \to \mathrm{BU} \times \mathbb{Z}$ factoring through $\mathrm{BU}(\mathrm{rank}(V \oplus W')) \times \{m\}$. This proves the result. ■

We have defined complex K-theory for compact Hausdorff spaces, but we can extend the definition to more general spaces by representability. Moreover, we can extend this to a cohomology theory represented by a spectrum KU, called the **complex K-theory spectrum**. In fact, KU can be constructed as a CW spectrum, made out of $\mathrm{BU} \times \mathbb{Z}$ in even degrees, and of U in odd degrees. The following classical result justifies this.

Theorem 1.2.7 (Bott Periodicity Theorem). We have homotopy equivalences $\Omega(\mathrm{BU}\times\mathbb{Z}) \cong \mathrm{U}$ and $\Omega\mathrm{U} \cong \mathrm{BU} \times \mathbb{Z}$. In both cases we have picked base-points for our spaces. □

Recall from Section 1.1 that any spectrum E yields an unreduced cohomology theory by writing $E^n(X) = [\Sigma^\infty X^+, E]_{-n}$. (More generally, there is a definition for pairs of spaces, which we ignore.) In the case $E = \mathrm{KU}$ and $n = 0$, we find that $\mathrm{KU}^0(X) = [\Sigma^\infty X^+, \mathrm{KU}]$. Let us sketch a proof that this KU^0 coincides with the one we defined earlier, at least for compact connected spaces.

Work in the framework of CW spectra. Thanks to compactness of X, a morphism $\Sigma^\infty X^+ \to \mathrm{KU}$ is precisely some based map $\Sigma^k X^+ \to \mathrm{KU}_k$, where KU_k denotes the k-th space in the CW complex that defines KU (i.e., either U or $\mathrm{BU}\times\mathbb{Z}$). By the suspension–loop space adjunction, we may as well write $X^+ \to \Omega^k \mathrm{KU}_k$, but by construction, $\Omega^k\mathrm{KU}_k \cong \mathrm{KU}_0 = \mathrm{BU} \times \mathbb{Z}$. Thus every map $\Sigma^\infty X^+ \to \mathrm{KU}$ corresponds to a map $X^+ \to \mathrm{BU} \times \mathbb{Z}$. Similar arguments show that homotopies of maps of spectra $\Sigma^\infty X^+ \to \mathrm{KU}$ correspond to homotopies of maps $X^+ \to \mathrm{BU} \times \mathbb{Z}$.

Example 1.2.8. The last family of spectra we will define are the Thom spectra. We will only briefly consider their construction, but we cannot do much justice to its profound historical significance.

Before we introduce them, we need the following notion, which will also be used in the next section. Given an arbitrary vector bundle $V \to B$ over a paracompact base space, we define the **Thom space** $T(V)$ as follows. First, using partitions of unity, construct a metric on the vector bundle V, so that we may asssociate to it a disk bundle $D(V)$ and a sphere bundle $S(V)$. Define $T(V)$ to be the space obtained by collapsing the subspace $S(V)$ of $D(V)$ to a point. If you don't care for metrics, the following definition might be more appealing. Start out with the set $V \sqcup \{\infty\}$. Topologize this set by declaring a neighbourhood of ∞ to be open if and only if its complement is closed in V *and* the intersection with every fibre of V is compact.

We begin our construction with a space X, and a map $X \to \mathrm{BO}$, where X is usually something fair like BU or BSO. Write $X(k)$ for the pullback $\mathrm{BO}(k) \times_{\mathrm{BO}} X$. The tautological bundle ξ_k on $\mathrm{BO}(k)$ pulls back to a bundle ξ_k^X on $X(k)$. Notice that the universal property of $\mathrm{BO}(k+1)$ tells us that the bundle $\xi_k^X \oplus \varepsilon^1$ gives rise to a map $X(k) \to X(k+1)$. We know write $\mathrm{M}X(k)$ for the Thom space $T(\xi_k^X)$ of the bundle ξ_k^X over $X(k)$. The map $X(k) \to X(k+1)$ induces a map $T(\xi_X^k \oplus \varepsilon^1) \to \mathrm{M}X(k+1)$ on Thom spaces. As $T(\xi_X^k \oplus \varepsilon^1)$ is equivalent to $\Sigma T(\xi_X^k)$, we can write it as a map $\Sigma \mathrm{M}X(k) \to \mathrm{M}X(k+1)$.

The spaces $\mathrm{M}X(k)$, along with these maps $\Sigma \mathrm{M}X(k) \to \mathrm{M}X(k+1)$, yield a CW spectrum, called the **Thom spectrum** and denoted $\mathrm{M}X$. In the special case where X is BO we write MO instead of MBO, and likewise for the cases where X is, say, BU or BSO. The spectrum MU that arises in this way is particularly important, and will end up being central in the next section. Without proof, we mention one relevant property of Thom spectra. The proof can be found in [7, Prop. 2.17].

Lemma 1.2.9. Let $X \to \mathrm{BO}$ be a map of H-spaces, where BO is equipped with the H-space structure classifying direct sums. Then the Thom spectrum $\mathrm{M}X$ admits the structure of a ring spectrum. If the H-space structure on X is commutative, then $\mathrm{M}X$ is commutative as well. □

Before we close off, I want to briefly talk about the aforementioned historical equivalence. The reader can safely skip this paragraph if he wishes. Thanks to the work by Thom, we know that there is a deep relationship between the Thom spectra, and cobordism of manifolds. Most prominently, the **Pontryagin–Thom construction** establishes an isomorphism between the bordism group Ω_n^X of n-manifolds with an X-structure on their stable normal bundle, and the n-th stable homotopy group $\pi_n^S(\mathrm{M}X)$. The construction allows us to view the bordism groups from a homotopy-theoretic viewpoint, and this has led to significant results in the calculation of bordism groups.

1.3 Complex orientations

We start out this section with the following well-known and classical theorem, known as the Thom isomorphism theorem.

Theorem 1.3.1 (Thom Isomorphism Theorem). Let $p\colon V \to B$ be a rank-n oriented real vector bundle over a paracompact base space. There exists an element $u \in H^n(V, V \setminus B)$, called the **Thom class** of V, such that, for every fibre V_b the restriction of u to $(V_b, V_b \setminus \{0\})$ is the class induced by the orientation of V. Moreover, the map $H^k(V) \to H^{k+n}(V, V \setminus B)$ obtained by cupping with u is an isomorphism. □

Proof sketch: The Thom class can be constructed directly by invoking the orientation on our bundle. Given the Thom class, the isomorphism $H^k(V) \cong H^{k+n}(V, V \setminus B)$ is given by a direct application of the Leray–Hirsch theorem. ■

Rather than working with relative cohomology, we can invoke Thom spaces instead, leading to the following lemma. It allows us to conclude a direct relation between the cohomology of the total space of a vector bundle, and the cohomology of the associated Thom space.

Lemma 1.3.2. If $V \to B$ is a vector bundle, then $H^p(V, V \setminus B) \cong \widetilde{H}^p(T(V))$. □

Proof: The right-hand side is isomorphic to $H^p(D(V), S(V))$ as $S(V)$ is a nice inclusion into $D(V)$. By homotopy invariance, we can write this as $H^p(D(V), D(V) \setminus B)$. Viewing V as the complex complement of the boundary of $D(V)$, we may apply the excision theorem to conclude the desired result. ■

In this section, we first expand on the above idea, and then use it to define complex orientations. Roughly speaking, a complex orientation on a multiplicative cohomology theory is the assignment of a Thom class to every complex vector bundle. The main question we will then be concerned with is how Thom classes of line bundles behave under taking tensor products. We will see that this behaviour can be described in terms of formal group laws. This will end up being the starting point of a deep connection between formal groups and algebraic topology. Our discussion will closely follow [7, Sections 2.6, 2.7 and 2.9].

If E is a ring spectrum representing a multiplicative cohomology theory, we define a **Thom class** to be any element of $\widetilde{E}^n(T(V))$ whose restriction to every compactified fibre $\overline{V_b}$ yields a generator of $\widetilde{E}^n(\overline{V_b}) \cong \widetilde{E}^n(S^n) \cong \widetilde{E}^0(S^0)$ as a module over itself. A vector bundle admitting such a Thom class is called E**-orientable**, and a choice of a Thom class is also called an E**-orientation**. We point out that, unlike the case of ordinary cohomology, a choice of Thom class need not be unique, even up to sign.

Theorem 1.3.3. Let $p\colon V \to B$ be an E-orientable rank-n vector bundle, with Thom class $u \in \widetilde{E}^n(T(V)) \cong E^n(V, V \setminus B)$. Then cupping with u yields an isomorphism $E^k(V) \xrightarrow{\sim} E^{k+n}(V, V \setminus B)$, hence $E^k(V) \cong \widetilde{E}^{k+n}(T(V))$. □

Proof: The proofs of both Theorem 1.3.1 and Lemma 1.3.2 carry over to generalized cohomology theories. In particular, invoking the Leray–Hirsch theorem is still valid — see [13, Thm. 15.47]. ■

Thom classes can be pulled back. Consider a map of spaces $f\colon X \to Y$, and let $V \to Y$ be an E-orientable rank-n vector bundle. We can pull this bundle back to a bundle f^*V over X. The morphism $f^*V \to V$ induces a morphism of corresponding Thom classes, and hence a morphism of E-cohomology groups in the opposite direction. Under this morphism, a Thom class $u \in \widetilde{E}^n(T(V))$ gets sent to a class in $\widetilde{E}^n(T(f^*V))$, which is a Thom class for f^*V.

Let E be a multiplicative cohomology theory. Write BU(1) for the classifying space of complex line bundles. It comes equipped with a line bundle γ, often called the tautological bundle. We say E is **complex-orientable** if the line bundle γ is E-orientable. A choice of E-orientation for γ is what we call a **complex orientation** of E. For any decent space X, then, complex line bundles over X correspond to homotopy classes of maps $X \to \mathrm{BU}(1)$, and the bundle over X can be obtained by pulling back the Thom class of BU(1) to a Thom class of the line bundle of X. Thus, complex orientations give a natural choice of Thom classes for all complex line bundles, over every space.

To demontrate the computational advantages of a complex orientation, we compute the cohomology rings of complex projective space.

Lemma 1.3.4. Let E be a complex-oriented cohomology theory. Denote by u_n the Thom class in $E^2(T(V_n))$ of the tautological bundle $V_n \to \mathbb{CP}^n$ of $\mathbb{CP}^n$, and write i_n for the inclusion of the zero section $\mathbb{CP}^n \hookrightarrow T(V_n)$, inducing a map on cohomology $i_n^*\colon \widetilde{E}^*(T(V_n)) \xrightarrow{\sim} E^*(V_n, V_n \setminus \mathbb{CP}^n) \to E^*(\mathbb{CP}^n)$. For all finite n, we have a ring isomorphism $E^*(\mathbb{CP}^n) \cong E^*[i_n^* u_n]/((i_n^* u_n)^{n+1})$, and in the infinite case, we have $E^*(\mathbb{CP}^\infty) \cong E^*[[u_\infty]]$. □

A few comments might be useful. First, we emphasize that the Thom class has degree 2 rather than degree 1, because the bundles over $\mathbb{CP}^n$ are complex. Second, in the ring isomorphism, we have used E^* as a shorthand for the coefficient ring of E. Finally, we remark that the double brackets in the cohomology of $\mathbb{CP}^\infty$ may be surprising, as it would imply that, for ordinary cohomology, $H^*(\mathbb{CP}^\infty)$ is $\mathbb{Z}[[u]]$ rather than $\mathbb{Z}[u]$. This is because we adopt the convention that the cohomology ring $E^*(\cdot)$ is a product over the $E^k(\cdot)$, rather than a direct sum.

Proof sketch of Lemma 1.3.4: The key observation is that $\mathbb{CP}^{n+1}$ is homotopy equivalent to the Thom space $T(V_n)$, hence Theorem 1.3.3 tells us that $E^k(\mathbb{CP}^n) \cong \widetilde{E}^{k+2}(\mathbb{CP}^{n+1})$. This allows us to relate the cohomology of $\mathbb{CP}^n$ to that of $\mathbb{CP}^{n+1}$, so that we can proceed by induction. The base case is given by $n = 0$. The space $\mathbb{CP}^0$ is a single point, so $E^*(\mathbb{CP}^0) = E^*(*) = \widetilde{E}^*(S^0)$, as desired.

For the case $n = \infty$, there are two ways to proceed. The Milnor sequence states that if we have a diagram of spaces $X_0 \to X_1 \to X_2 \to \cdots$ whose homotopy colimit is X, then there's a short exact sequence

$$0 \longrightarrow \lim^1 E^{*-1}(X_i) \longrightarrow E^*(X) \longrightarrow \varprojlim_i E^*(X_i) \longrightarrow 0$$

Our tower of maps is surjective, so that the Mittag-Leffler condition is satisfied and $\lim^1$ vanishes. Alternatively, it is known that being complex-orientable is equivalent to the degeneration of the Atiyah–Hirzebruch spectral sequence at the second page (see [5, Lecture 1]) which allows us to effectively compute the cohomology of $\mathbb{CP}^\infty$ as well. ∎

More elaborate versions of the argument above allows us to compute the cohomology groups for some related spaces as well. For example, the cohomology $E^*((\mathbb{CP}^\infty)^{\times n})$ of the n-fold product of the infinite-dimensional projective space is $E^*[\![x_1, \ldots, x_n]\!]$, where the generators are the pullbacks of the generator of $E^*(\mathbb{CP}^\infty)$ along the various projection maps.

We know that complex orientations give rise to a choice of Thom class for line bundles. But more is true than that: they yield a choice of Thom class for all complex vector bundles, no matter what rank. This can be done thanks to the splitting principle: for any rank-n complex vector bundle $V \to B$, one can find a fibration $p\colon \mathbb{P}V \to B$ with fibre $\mathbb{CP}^{n-1}$ such that the pullback bundle $p^*V \to \mathbb{P}V$ splits into a direct sum $L_1 \oplus \cdots \oplus L_n$ of line bundles. For each of these line bundles, the complex orientation allows one to give a class in $E^2(T(L_i))$, which may be cupped to a class in $E^{2n}(T(L_1 \oplus \cdots \oplus L_n))$. One can then argue that this class must in fact be in the image of the map $E^{2n}(T(V)) \to E^{2n}(T(L_1 \oplus \cdots \oplus L_n))$. To do this, one needs to relate the E-cohomology of $\mathbb{P}V$ with that of B, either by invoking the Atiyah–Hirzebruch spectral sequence or using the Leray–Hirsch theorem for generalized multiplicative cohomology (see [13, Thm. 15.47]).

We now explain how formal group laws pop up from complex orientations. Up to homotopy, there is a natural map $m\colon \mathbb{CP}^\infty \times \mathbb{CP}^\infty \to \mathbb{CP}^\infty$ classifying the tensor product bundle $\mathrm{pr}_1^* \eta_\infty \otimes \mathrm{pr}_2^* \eta_\infty$. It carries a certain ubiquity because in a sense, it encodes tensor products. To make this precise, let L and L' be two line bundles on a space X, classified by maps $f_1\colon X \to \mathbb{CP}^\infty$ and $f_2\colon X \to \mathbb{CP}^\infty$. Then the tensor product $L \otimes L'$ is classified by the composition of the product map $f_1 \times f_2$ with the multiplication map m.

The map m induces a map $E^*[\![x]\!] \to E^*[\![x_1, x_2]\!]$ on cohomology. As we will see in all that follows, lots of deep information about E is contained within this particular map on cohomology, so we make sure to spend a good amount of time on it. We may look at the image of the monomial x, which is a formal power series $f(x_1, x_2)$. In decent cases, this formal power series entirely determines the map on cohomology, and in any case, it determines a good chunk of it.

The properties of the tensor product operation manifest themselves as properties of the power series $f(x_1, x_2)$. For instance, consider the map $\mathbb{CP}^\infty \to \mathbb{CP}^\infty \times \mathbb{CP}^\infty$ obtained by inclusion into the first factor, and compose it with m. Essentially by definition this is homotopic to the identity map. One can check that this translates into the requirement that $f(x_1, 0) = x_1$; by symmetry, we also have $f(0, x_2) = x_2$. In a similar vein one can show that the associativity of m (up to homotopy, that is) translates into the requirement that $f(x_1, f(x_2, x_3)) = f(f(x_1, x_2), x_3)$, and that the commutativity yields $f(x_1, x_2) = f(x_2, x_1)$. Taken together, these restrictions are precisely what defines a one-dimensional formal group law over E_*, which we investigated in Appendix A.

The reader might have noticed that we wrote E_* instead of E^*. As rings, E_* and E^* are the same, so that a formal group law over one is equivalent to a formal group law over the other, hence we have equivalent ring maps $L \to E_*$ and $L \to E^*$. The *gradings* of E_* and E^* however are reversed, and only the ring map $L \to E_*$ is a map of graded rings, with respect to the grading on L introduced below Theorem A.2.2.

The generators x, x_1, and x_2 are not canonical, as they depend on the choice of a Thom class. The formal group law that arises from the map m may therefore not be fixed either. Nonetheless, we have the following result.

Lemma 1.3.5. The formal group law associated to a complex-oriented cohomology theory is independent of the choice of complex orientation, up to isomorphism. □

Proof: Let E be a cohomology theory, with two complex orientations, corresponding to two choices of Thom classes u_1, u_2 of the tautological bundle of $\mathbb{CP}^\infty$. By Lemma 1.3.4 we have $E^*(\mathbb{CP}^\infty) \cong E^*[\![u_1]\!] \cong E^*[\![u_2]\!]$, and the isomorphism $E^*[\![u_1]\!] \xrightarrow{\sim} E^*[\![u_2]\!]$ is uniquely determined by the image of u_1, which can be expressed by some formal power series $h(t)$ with coefficients in E^*. Let $f_1(x, y)$ and $f_2(x, y)$ be the formal group laws associated to the two complex orientations. By naturality of the diagram

$$\begin{array}{ccc} E^*[\![u_1]\!] & \longrightarrow & E^*[\![(u_1)_1, (u_1)_2]\!] \\ \downarrow & & \downarrow \\ E^*[\![u_2]\!] & \longrightarrow & E^*[\![(u_2)_1, (u_2)_2]\!] \end{array}$$

it immediately follows that $h(f_1(x, y)) = f_2(h(x), h(y))$, which, by symmetry, is an isomorphism of formal group laws. ■

Example 1.3.6. The two simplest examples are ordinary cohomology and complex K-theory, both of which are orientable. With purely algebraic arguments one can show that the formal group law associated to $H\mathbb{Z}$ is given by the additive formal group law; the same holds for ordinary cohomology with other coefficients, although the underlying coefficient ring of the formal group law will be different. To determine the formal group law of complex K-theory, one uses Bott periodicity to find it to be the multiplicative formal group law. More details on this can be found in [7, Section 2.7].

Consider a morphism $E \to E'$ of ring spectra. If E is complex-orientable then so is E'; indeed, any choice of Thom class $u \in \widetilde{E}^2(T(\gamma))$ for the tautological bundle γ over BU(1), can be sent to a class in $\widetilde{E'}^2(T(\gamma))$, which one can show is an E'-orientation for γ. We have a commuting diagram

$$\begin{array}{ccc} E^*[[x]] & \longrightarrow & E^*[[x_1, x_2]] \\ \downarrow & & \downarrow \\ E'^*[[x]] & \longrightarrow & E'^*[[x_1, x_2]] \end{array}$$

induced by the tensor product of line bundles, thanks to naturality. This tells us that the formal group law over E'_* is the pullback of the formal group law over E_* along the ring map $E_* \to E'_*$.

A short aside. Just for fun, let's see an easy application of complex orientable to a concrete question about spectra. If pressed for time the reader can skip this result.

Lemma 1.3.7. $H\mathbb{Z}/p \wedge \mathrm{KU} = 0$. □

Proof: The maps $H\mathbb{Z}/p \to H\mathbb{Z}/p \wedge \mathrm{KU}$ (defined by smashing $H\mathbb{Z}/p$ with the sphere spectrum on the right, and considering the map $\mathrm{Id} \wedge 0$) and $\mathrm{KU} \to H\mathbb{Z}/p \wedge \mathrm{KU}$ (defined similarly) induce two complex orientations on the smash product. By Example 1.3.6 and Example 1.3.6, these two formal group laws are the additive and multiplicative formal group law, respectively. By virtue of Lemma 1.3.5, these two formal group laws must be isomorphic.

The map $H\mathbb{Z}/p \to H\mathbb{Z}/p \wedge \mathrm{KU}$ gives rise to a morphism on homotopy groups $\pi_*^S(H\mathbb{Z}/p) \to \pi_*^S(\mathrm{KU})$. In Example 1.2.2, we saw that the stable homotopy groups of $H\mathbb{Z}/p$ form a ring of characteristic p, hence so is $\pi_*^S(H\mathbb{Z}/p \wedge \mathrm{KU})$. In Example A.3.5, however, we determined that in such a ring, the additive and multiplicative formal group law cannot be isomorphic. The only way out is for $\pi_*^S(H\mathbb{Z}/p \wedge \mathrm{KU})$ to be 0, which by Theorem 1.1.1 implies that the spectrum must be trivial. ■

In the previous section, Example 1.2.8, we introduced what we called Thom spectra. It turns out that these Thom spectra are related to complex orientations, in the following way. We refer the reader to [7, Prop. 2.25] for a proof.

Theorem 1.3.8. For any ring spectrum E, complex orientations on E are in one-to-one correspondence with ring spectrum maps $\mathrm{MU} \to E$ in $\mathrm{Ho}(\mathsf{Sp})$. Here we endow MU with the ring structure obtained in Lemma 1.2.9. □

There is a canonical complex orientation on MU, corresponding to the identity map $\mathrm{MU} \to \mathrm{MU}$, and it gives rise to a formal group law over MU_*. Now take any complex orientation on a ring spectrum E. By Theorem 1.3.8, the formal group law over E_* determined by this orientation is the pullback of the formal group law over MU_* along the map $\mathrm{MU}_* \to E_*$. It makes sense, then, to ask what this 'universal' formal group law over MU_* should be. By Theorem A.2.2, our formal group law over MU_* corresponds to a ring map $L \to \mathrm{MU}_*$. We have the following result.

Theorem 1.3.9 (Quillen's Theorem on MU**).** The ring map $L \to \mathrm{MU}_*$ induced by the canonical orientation on MU is an isomorphism of rings. □

A proof of the above theorem can be found in [5, Lectures 8–10]. We use the theorem to endow L with a natural grading, namely, the one corresponding to the grading of MU_*. For any complex-orientable ring spectrum E, the ring map $L \to \pi_* E$ that corresponds to the formal group law over $\pi_* E$ will then automatically be a morphism of *graded* rings.

1.4 Bousfield localizations

In this section we introduce the general notion called Bousfield localization of spectra, and investigate some of its properties. We closely follow Bousfield's original paper [1], with additional remarks and proofs wherever we deem it useful to the reader.

Throughout this section we work within the stable homotopy category. Let E be a spectrum. A morphism $f: A \to B$ in $\mathrm{Ho}(\mathsf{Sp})$ is called an E**-equivalence** if the induced map $f_*: \widetilde{E}_*(A) \to \widetilde{E}_*(B)$ is an isomorphism. Equivalently by Whitehead's theorem (1.1.1) the induced map $E \wedge A \to E \wedge B$ is an isomorphism. A spectrum A is called E**-acyclic** if $\widetilde{E}_*(A) \cong 0$. Again by Theorem 1.1.1, A is E-acyclic if and only if $E \wedge A \cong 0$.

Lemma 1.4.1. A morphism $f: A \to B$ is an E-equivalence if and only if B/A is E-acyclic. □

Proof: In our discussion on the stable homotopy category, we learned that smash products preserve cofibres. That is, the cofibre of the map $E \wedge A \to E \wedge B$ is given by $E \wedge (B/A)$. But the map $E \wedge A \to E \wedge B$ is an isomorphism, so when we write down the long cofibre and fibre sequence associated to the map $E \wedge A \to E \wedge B$ we find that, for any spectrum X, $[E \wedge (B/A), X]_* = 0$ and $[X, E \wedge (B/A)]_* = 0$; both are sufficient to imply that $E \wedge (B/A) \cong 0$, hence B/A is E-acyclic, as desired. Every step in the above proof can be inverted to yield a proof for the opposite direction. ■

A spectrum X is called E**-local** if, for any E-acyclic spectrum A, we have $[A, X]_* \cong 0$.

Lemma 1.4.2. Let X and E be two spectra. Then X is E-local if and only if, for any E-equivalence $f: A \to B$, the induced map $f^*: [B, X]_* \to [A, X]_*$ is an isomorphism. □

Proof: Let C be an E-local spectrum, and let A be E-acyclic. Then the zero morphism $0 \to A$ is an E-equivalence; the induced isomorphism between $[0, C]_*$ and $[A, C]_*$ proves the desired result. Conversely, suppose that $[A, C]_* = 0$ for every E-acyclic spectrum A. Consider an arbitrary E-equivalence $f: A \to B$. Associated to the distinguished triangle $A \to B \to B/A$ is the long cofibre sequence $\cdots \to [\Sigma A, C]_* \to [B/A, C]_* \to [B, C]_* \to [A, C]_* \to [\Sigma^{-1} A, C]_* \to \cdots$. By Lemma 1.4.1, the spectrum B/A is E-acyclic (and therefore so are its shifts $\Sigma^n(B/A)$) hence exactness of the long cofibre sequence implies $[B, C]_* \cong [A, C]_*$, as desired. ■

We remark that the above lemma fails in the non-stable world. This is the first indication that localizations are very well-behaved in the stable world.

Lemma 1.4.3 (E-Whitehead Theorem). Any E-equivalence between E-local spectra is an isomorphism in $\mathrm{Ho}(\mathsf{Sp})$. □

Proof: Denote by $f\colon A \to B$ an E-equivalence between E-local spectra. The induced maps $f^*\colon [B, B]_* \to [A, B]_*$ and $f^*\colon [B, A]^* \to [A, A]^*$ are, by definition of E-locality, isomorphisms. The isomorphisms guarantee the desired existence of a left and right inverse of f. ■

Lemma 1.4.4. The following claims hold.

- Module spectra over a ring spectrum E are E-local.
- If $A \to B \to C \to \Sigma A$ is a distinguished triangle in $\mathrm{Ho}(\mathsf{Sp})$ and two of the three spectra A, B and C are E-local, then so is the third.
- E-local spectra are closed under products, and more generally, under arbitrary small homotopy limits.
- E-local spectra are closed under retracts. □

Proof: Let M be a module spectrum over a ring spectrum E, and let A be an E-acyclic spectrum. Consider any map $f\colon A \to M$, which we aim to show is the zero map. Decompose f as in the diagram below:

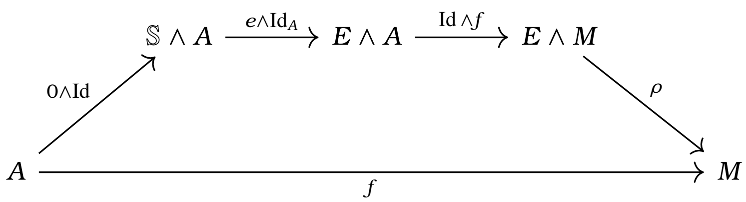

Here the first map exists because $\mathbb{S}$ is the unit of the smash product. Now, the object $E \wedge A$ is trivial, since A is E-acyclic. Hence f must be the zero map. The second part follows by considering the long fibre sequence of the distinguished triangle, obtained by applying $\mathrm{Hom}(X, \cdot)$ for any E-acyclic spectrum X. The claim regarding products is obvious, and we skip the generalization to arbitrary small homotopy limits as we will never need this. Finally, let A be a retract of X. There exist maps $A \to X \to A$ that compose to the identity. For any E-acyclic spectrum B, we have a composition $[B, A]_* \to [B, X]_* \to [B, A]_*$, factoring through the zero object, which composes to the identity. It must be the case that $[B, A]_* \cong 0$. ■

An E-**localization** of a spectrum X is an E-equivalence $X \to L_E X$ such that $L_E X$ is E-local. By Lemma 1.4.3 any E-localization must be unique up to isomorphism. Bousfield's goal was to show that we can make a functorial choice of a E-localized spectrum for every spectrum X.

In other words, there exists a functor $L_E\colon \mathrm{Ho}(\mathsf{Sp}) \to \mathrm{Ho}(\mathsf{Sp})$, called the ***E*-localization functor** together with a natural transformation $\mathrm{Id} \to L_E$, such that $X \to L_E X$ is an E-localization for all X.

Lemma 1.4.5. Every E-localization $f\colon X \to L_E X$ satisfies the following universal conditions. It is initial among those morphisms $X \to Y$ with Y an E-local spectrum; on the other hand, it is terminal among the E-equivalences $X \to Y$. □

Proof: As for the first claim, suppose we have a morphism $X \to Y$ with Y an E-local spectrum. By Lemma 1.4.2, Y being E-local tells us that the E-equivalence $X \to L_E X$ gives rise to an isomorphism $[L_E X, Y]_* \to [X, Y]_*$, from which a unique map $L_E X \to Y$ follows. The other claim is proved similarly. ■

This lemma suggests a way of obtaining E-localizations. For any spectrum X, consider the collection of all E-equivalences $X \to Y$, and take the colimit, or perhaps take all maps $X \to Y$ with E-local and consider the limit inside the category of spaces under X. Unfortunately, the collection of E-equivalences and of E-local spectra may well be proper classes, so taking the colimit is a set-theoretically unsound procedure. As it turns out, this is more or less the *only* obstruction to our naive approach, and a large part of [1] is dedicated to showing that there is a way to get around this issue. In fact, Bousfield shows an even stronger result: every spectrum can be decomposed into an E-local spectrum and an E-acyclic spectrum. We state the result below; it can also be found as Thm. 1.1 of Bousfield's paper.

Theorem 1.4.6. For any spectrum A, there is a natural triangle ${}_EA \to A \to L_E A \to \Sigma({}_EA)$ in $\mathrm{Ho}(\mathsf{Sp})$ such that ${}_EA$ is E-acyclic, and $L_E A$ is E-local. In particular, the localization functor $L_E\colon \mathrm{Ho}(\mathsf{Sp}) \to \mathrm{Ho}(\mathsf{Sp})$ exists. □

Corollary 1.4.7. The full subcategory of the stable homotopy category consisting of only the E-local spectra is equivalent to the category obtained by formally inverting the E-equivalences in $\mathrm{Ho}(\mathsf{Sp})$. □

It is sometimes said that the subcategory of E-local spectra is what is left of the stable homotopy category after considering only the information that E grants us. The above corollary makes this precise: we formally invert anything that the E-homology theory cannot distinguish between. From this point on, we refer to either of the two equivalent categories in the above statement as the ***E*-local stable homotopy category**.

Proof: Just temporarily we'll write $\mathrm{Ho}(\mathsf{Sp})/E$ for the localization at the E-equivalences, and we'll write $\mathrm{Ho}(\mathsf{Sp})_E$ for the subcategory of E-local spectra. By Theorem 1.4.6, there is a functor $L_E\colon \mathrm{Ho}(\mathsf{Sp}) \to \mathrm{Ho}(\mathsf{Sp})_E$, and it admits the inclusion functor $i\colon \mathrm{Ho}(\mathsf{Sp})_E \to \mathrm{Ho}(\mathsf{Sp})$ as its right inverse. By Lemma 1.4.3, we know that L_E passes through $\mathrm{Ho}(\mathsf{Sp})/E$. The functors L_E and i between $\mathrm{Ho}(\mathsf{Sp})/E$ and $\mathrm{Ho}(\mathsf{Sp})_E$ define mutually inverse functors. As i is the right inverse to L_E,

$i \circ L_E \cong \mathrm{Id}$ automatically; on the other hand, the isomorphism $L_E \circ i \cong \mathrm{Id}$ follows from the natural E-equivalences $X \to L_E X$ provided by Theorem 1.4.6. ■

Two spectra E and F are said to be **Bousfield equivalent** if $\widetilde{E}_*(X) = 0$ if and only if $\widetilde{F}_*(X) = 0$ for all spectra X.

Lemma 1.4.8. If L_E and L_F are isomorphic, then E and F are Bousfield equivalent, and conversely, Bousfield equivalent spectra induce isomorphic localization functors. □

Proof: If E and F are Bousfield equivalent spectra, then a spectrum is E-local if and only if it is F-local. Using the first universal property of the localization in Lemma 1.4.5, it must be the case that $L_E X$ and $L_F X$ are naturally isomorphic, for all X. Conversely, suppose there's a natural isomorphism between L_E and L_F. The triangle in Theorem 1.4.6 tells us that a spectrum A is E-acyclic if and only if $L_E A = 0$, while it is F-acyclic if and only if $L_F A = 0$. The natural isomorphism ensures that these two hypotheses are equivalent. ■

Generally speaking, Bousfield localizations of spectra are rather illusive, in the sense that it can be hard to describe them explicitly — and indeed the general existence of proof in [1] is rather non-constructive. Under some additional conditions, however, there is more to be said. For instance, in the same paper, Bousfield shows that for connective spectra, Bousfield localizations admit a reasonably accessible description, which we won't need and therefore won't state. Another situation in which localization becomes nice is when the localization is smashing. We say a spectrum E is **smashing** if, for all spectra X, $L_E X \cong X \wedge L_E \mathbb{S}$.

Example 1.4.9. Let SG be any Moore spectrum, first considered in Example 1.2.4. Then it makes sense to localize a spectrum X with respect to SG so as to yield a spectrum $L_{SG}X$. An important example is localization at the Moore spectrum $\mathbb{S}/p$, also called ***p*-completion**. Let's expand on what it means for a spectrum X to be E-local, where $E = \mathbb{S}/p$. It means that, for any E-acyclic spectrum A, we have $[A, X]_* = 0$. In turn, a spectrum A is acyclic if $A \wedge \mathbb{S}/p = 0$. As taking cofibres commutes with smash products, this is equivalent to saying that the cofibre of the map $A \wedge \mathbb{S} \xrightarrow{\mathrm{Id} \wedge (\times p)} A \wedge \mathbb{S}$ is trivial. Put in short, A being acyclic means that the multiplication-by-p map $A \xrightarrow{\times p} A$ is an isomorphism in $\mathrm{Ho}(\mathsf{Sp})$.

Another special case is when $G = \mathbb{Z}_{(p)}$. In that case, localization at SG is called ***p*-localization**. Rather than writing the rather clumsy $L_{S\mathbb{Z}_{(p)}}X$, one often sees the notation $X_{(p)}$ instead. Without giving a proof, we state the following result, referring the interested reader to [1, Prop. 2.4], and the references mentioned therein, for more information.

Lemma 1.4.10. Localization at the Moore spectrum $S\mathbb{Z}_{(p)}$ is smashing, and for any spectrum X, we have $\pi_* L_{S\mathbb{Z}_{(p)}} X \cong \pi_*(X) \otimes \mathbb{Z}_{(p)}$. □

Example 1.4.11. A particularly important example is the localization of complex K-theory. We denote the p-localized complex K-theory spectrum by $\mathrm{KU}_{(p)}$. In turn, any spectrum X can be localized at $\mathrm{KU}_{(p)}$, and in fact studying the $\mathrm{KU}_{(p)}$-local stable homotopy category will be the main goal of this thesis.

We know from Example 1.3.6 that KU admits a complex orientation, corresponding to some ring map $\mathrm{MU} \to \mathrm{KU}$. We may compose this with the p-localization map to obtain a complex orientation on $\mathrm{KU}_{(p)}$, whose formal group law is essentially the same except taken over the ring $\pi_*^{\mathrm{S}}(\mathrm{KU}_{(p)}) \cong \pi_*^{\mathrm{S}}(\mathrm{KU}) \otimes \mathbb{Z}_{(p)} \cong \mathbb{Z}_{(p)}[u^{\pm 1}]$.

Chapter 2

Hopf algebroids and algebraic stacks

The goal of this chapter is to introduce Hopf algebroids and algebraic stacks. Briefly speaking, Hopf algebroids are groupoid objects in the category of affine schemes, and algebraic stacks are groupoid-valued sheaves on the category of affine schemes satisfying a descent condition. We spend some time introducing both notions, before investigating the relation between them.

2.1 Groupoid objects and Hopf algebroids

Recall first that a groupoid is just a category whose morphisms are invertible. In particular, a one-object groupoid is just a group when looking at the set of morphisms. Given a category $\mathcal{C}$ admitting finite fibre products, a **groupoid object** in $\mathcal{C}$ consists of a pair of objects (O, M) in $\mathcal{C}$ — the O stands for 'objects' and the M stands for 'morphisms' — along with five morphisms

$$\begin{aligned} \mathrm{srce}, \mathrm{trgt} \colon & M \to O\,, \\ \mathrm{unit} \colon & O \to M\,, \\ \mathrm{inv} \colon & M \to M\,, \\ \mathrm{comp} \colon & M \times_{O,\mathrm{srce},\mathrm{trgt}} M \to M\,. \end{aligned}$$

We ask for the morphisms to satisfy the following axioms. Rather than remembering them, we urge the reader to understand why the axioms are as they are.

$$\begin{gathered} \mathrm{srce} \circ \mathrm{unit} = \mathrm{trgt} \circ \mathrm{unit} = \mathrm{Id}_O\,, \\ \mathrm{srce} \circ \mathrm{comp} = \mathrm{srce} \circ \mathrm{pr}_1 \qquad \text{and} \qquad \mathrm{trgt} \circ \mathrm{comp} = \mathrm{trgt} \circ \mathrm{pr}_2\,, \\ \mathrm{comp} \circ (\mathrm{Id}_M \times \mathrm{comp}) = \mathrm{comp} \circ (\mathrm{comp} \times \mathrm{Id}_M)\,, \\ \mathrm{comp} \circ (\mathrm{unit} \circ \mathrm{srce}, \mathrm{Id}_M) = \mathrm{comp} \circ (\mathrm{Id}_M, \mathrm{unit} \circ \mathrm{trgt}) = \mathrm{Id}_M\,, \\ \mathrm{inv} \circ \mathrm{inv} = \mathrm{Id}_R\,, \qquad \mathrm{srce} \circ \mathrm{inv} = \mathrm{trgt}\,, \qquad \mathrm{trgt} \circ \mathrm{inv} = \mathrm{srce}\,, \\ \mathrm{comp} \circ (\mathrm{Id}_R, \mathrm{inv}) = \mathrm{unit} \circ \mathrm{srce}\,, \qquad \mathrm{comp} \circ (\mathrm{inv}, \mathrm{Id}_R) = \mathrm{unit} \circ \mathrm{trgt}\,. \end{gathered}$$

Example 2.1.1. If O is the final object of our category $\mathcal{C}$, so that the source and target maps are the trivial ones, then M is a group object in the category of $\mathcal{C}$. Conversely, any group object can be turned into a groupoid object in this way.

Lemma 2.1.2. A pair (O, M) is a groupoid object in $\mathcal{C}$ if and only if their functors of points give rise to a (strict) functor $\mathcal{C}^{\mathrm{op}} \to \mathsf{Grpd}$. □

Proof: If (O, M) is a groupoid object, then for any object X, $\mathrm{Hom}(X, O)$ and $\mathrm{Hom}(X, M)$ form the object and morphism set of a groupoid. For any morphism $X \to X'$, we get a functor of groupoids in the opposite direction. This association is strictly commutative. Conversely, if the functors of points of O and M can be lifted to a strict functor $\mathcal{C}^{\mathrm{op}} \to \mathsf{Grpd}$, then the Yoneda Lemma implies that this structure comes from morphisms on O and M that automatically turn them into a groupoid object. ■

A **(commutative) Hopf algebroid** is nothing but a groupoid object in the category of affine schemes. The definition has an obvious extension to the relative context, where we instead ask for a Hopf algebroid to be a groupoid object in the category of affine R-schemes for some base ring R.

Example 2.1.3. The most important example for us is the pair $(\operatorname{Spec} L, \operatorname{Spec} W)$, where L and W are the rings introduced in Section A.2. We verified in that section that, given an affine scheme $\operatorname{Spec} A$, $\mathrm{Hom}(\operatorname{Spec} A, \operatorname{Spec} L)$ and $\mathrm{Hom}(\operatorname{Spec} A, \operatorname{Spec} W)$ form the object and morphism set of a groupoid. Moreover, the operations defining this groupoid are natural in the sense that, for any morphism $\operatorname{Spec} A \to \operatorname{Spec} A'$, the corresponding map on Hom-sets is a functor on groupoids. Lemma 2.1.2 then implies that the pair (L, W) forms a Hopf algebroid.

The following technical restriction will often be imposed. A Hopf algebroid is said to be **flat** if either the source or the target, or equivalently both of them, is a flat morphism of schemes (or a flat ring map, if you prefer). We point out that the unit maps are then automatically *faithfully* flat, as they have an inverse defined by the unit map.

Let $(\operatorname{Spec} A, \operatorname{Spec} \Gamma)$ be a Hopf algebroid. Then a **left comodule** over our Hopf algebroid is an A-module M, along with a Γ-coaction in the form of an A-module morphism $\psi \colon M \to \Gamma \otimes_A M$, satisfying the property that the diagrams

$$\begin{array}{ccc} M & \xrightarrow{\sim} & A \otimes_A M \\ & {\scriptstyle\psi}\searrow \quad \nearrow{\scriptstyle \mathrm{unit}^\sharp \otimes \mathrm{Id}_M} & \\ & \Gamma \otimes_A M & \end{array} \qquad \text{and} \qquad \begin{array}{ccc} M & \xrightarrow{\psi} & \Gamma \otimes_A N \\ {\scriptstyle\psi}\downarrow & & \downarrow{\scriptstyle \mathrm{comp}^\sharp \otimes \mathrm{Id}_M} \\ \Gamma \otimes_A M & \xrightarrow[\psi \otimes \mathrm{Id}_M]{} & \Gamma \otimes_A \Gamma \otimes_A M \end{array}$$

should both commute. Notice that they encode dualized versions of unitality and the action property. A **right comodule** is defined analogously, but this notion is equivalent. We will therefore always stick to left comodules, and simply refer to them as **comodules** over $(\operatorname{Spec} A, \operatorname{Spec} \Gamma)$, or even just as Γ-comodules. There is an obvious notion of a **morphism of Γ-comodules**: it is a morphism of A-modules that is compatible with the Γ-coaction in the way that one would expect. This yields a category Γ–Comod.

Every Γ-comodule has an underlying A-module structure. Conversely, given an A-module M, there is a natural way to associate a Γ-comodule to it: take the tensor product $\Gamma \otimes_A M$, and notice that it has a natural Γ-coaction induced by the comultiplication map on Γ. This defines a functor A–Mod $\to$ Γ–Comod, which one can show is a right adjoint to the forgetful functor.

Lemma 2.1.4. The category of comodules over a *flat* Hopf algebroid is an abelian category. □

Proof sketch: Even without flatness, the category of comodules is an additive category, so it remains to be shown that we have kernels and cokernels, and they behave as desired. We give the construction of the kernels, indicating where flatness is needed, leaving the dual construction of cokernels and the remaining verifications to the reader.

Take a morphism $f\colon M \to M'$ of comodules over $(\operatorname{Spec} A, \operatorname{Spec} \Gamma)$. Start out with the kernel K of the underlying morphism of A-modules. We endow K with a coaction $K \to \Gamma \otimes_A K$ as follows. We have a commutative diagram of A-modules

$$\begin{array}{ccccc} K & \longrightarrow & M & \xrightarrow{\;f\;} & M' \\ \downarrow & & \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \psi'} \\ \Gamma \otimes_A K & \longrightarrow & \Gamma \otimes_A M & \xrightarrow[\mathrm{Id}_\Gamma \otimes f]{} & \Gamma \otimes_A K' \end{array}$$

The top row is left-exact by definition, and hence by flatness of Γ, so is the bottom row. This means that $\Gamma \otimes_A K$ is isomorphic to the kernel of the map $\mathrm{Id}_\Gamma \otimes f$. The universal property of kernels now forces the dashed arrow to exist. This dashed arrow satisfies the axioms of a coaction, thus turning K into the desired Γ-comodule. ■

The category of Γ-comodules has enough injectives. We sketch how the reader can prove this, if he wishes to do so. The category of A-modules is known to have enough injectives (see [11, Tag 01DD]). Let M now be a Γ-comodule, and write M_A for the underlying A-module of M. Take a monomorphism from the A-module M_A into some injective A-module I. By the aforementioned cofree adjunction, we have a morphism of Γ-comodules $M \to \Gamma \otimes_A I$. A direct verification shows that this morphism is also a monomorphism, and that $\Gamma \otimes_A I$ is injective.

Since we have enough injectives, we have a well-defined notion of an Ext-functor in the category of Γ-comodules, and, more generally, we can define right-derived functors of any left-exact functor.

Example 2.1.5. Under suitable conditions, ring spectra admit the structure of a Hopf algebroid, and hence give rise to an algebraic stack via the correspondence introduced in the previous section. Our next aim in this section is to briefly expand on this, following [13, Ch. 17]. Our story begins with a commutative ring spectrum E, and assume that E is evenly graded — that is, the ring $\pi_*(E)$ only has evenly graded homogeneous elements. This ensures that the ring $\pi_*(E)$ is commutative rather than just anti-commutative. Without this assumption, we'd end up with a non-commutative generalization of Hopf algebroids. While still interesting, this will not fit within the framework introduced in Section 2.2 and beyond (or, at least, not easily), hence we impose the condition.

Under an additional hypothesis which we mention in a moment, the pair of commutative rings $(\pi_*E, \widetilde{E}_*E)$ is going to form a Hopf algebroid. Looking back at the definition of a Hopf algebroid, the reader should convince himself that we need the following maps. A multiplication on $\widetilde{E}_*E$ turning it into a commutative ring; source and target maps $\pi_*E \to \widetilde{E}_*E$; an inversion map $\widetilde{E}_*E \to \widetilde{E}_*E$; a unit map $\widetilde{E}_*E \to \pi_*E$; a composition map $\widetilde{E}_*E \to \widetilde{E}_*E \otimes_{\pi_*E} \widetilde{E}_*E$. All but the very last one have a straightforward definition.

The multiplication $\widetilde{E}_*E \otimes_{\mathbb{Z}} \widetilde{E}_*E$ takes two maps $f\colon \Sigma^k\mathbb{S} \to E \wedge E$ and $g\colon \Sigma^\ell\mathbb{S} \to E \wedge E$, smashes them into a map $f \wedge g\colon \Sigma^{k+\ell}\mathbb{S} \to E \wedge E \wedge E \wedge E$, then applies the multiplication $\mu\colon E \wedge E \to E$ to the first and third component, and to the second and fourth component. The source and target maps $\pi_*E \to \widetilde{E}_*E$ takes a map $f\colon \Sigma^k\mathbb{S} \to E$, and smashes it either from the left or the right with the unit map $\eta\colon \mathbb{S} \to E$ that is part of the monoid structure of E. The inversion map $\widetilde{E}_*E \to \widetilde{E}_*E$ takes a map $f\colon \Sigma^k\mathbb{S} \to E \wedge E$, and swaps the two components in the wedge product. The unit map $\widetilde{E}_*E \to \pi_*E$ composes a class in $\widetilde{E}_*E$ with the multiplication μ.

What about the composition map? We impose the additional constraint that $\widetilde{E}_*E$ is flat as a module over π_*E, where we can choose either one of the two module structures induced by the source and target map. Such ring spectra are called **flat**. In that case, [13, Thm. 13.75] tells us that $\widetilde{E}_*(E \wedge E) \cong \widetilde{E}_*E \otimes \widetilde{E}_*E$. The composition map is now easy to define: start with a class $f\colon \Sigma^k\mathbb{S} \to E \wedge E$. Now smash this class with itself to find a map $\Sigma^{2k}\mathbb{S} \to E \wedge E \wedge E \wedge E$. Apply μ to the second and third component of this wedge product. Finally, apply the isomorphism mentioned above to find the two maps $\Sigma^k\mathbb{S} \to E \wedge E$.

In the next sections, we will talk about algebraic stacks, and then discuss exactly how Hopf algebroids are related to them. For more algebraic theory on Hopf algebroids, I also refer the reader to [9, Appendix A1].

2.2 Stacks over sites

In this section we recall some basic definitions about sites, sheaves, and stacks over sites. We will assume that the reader has seen the material before, and will therefore be rather concise here. For more details I refer the reader to [11, 00UZ] and [14, Section 2.3] for a discussion on sites and sheaves, and to [14, Ch. 3 and 4] for an introduction to stacks.

Let $\mathcal{C}$ be a category admitting finite fibre products. A **Grothendieck topology** on $\mathcal{C}$ is a set $\mathrm{Cov}(\mathcal{C})$ of families of morphisms $\{U_i \to U\}$ with fixed targets, called **coverings** of $\mathcal{C}$ or of U, such that the following axioms are satisfied.

- If $V \to U$ is an isomorphism then $\{V \to U\}$ is in $\mathrm{Cov}(\mathcal{C})$;
- if $\{U_i \to U\}$ forms a cover of U, and for each i, $\{V_{ij} \to U_i\}$ forms a cover of U_i, then the set of morphisms $\{V_{ij} \to U\}$ obtained by taking all compositions covers U;
- if $\{U_i \to U\}$ is a covering of U, and $V \to U$ is some morphism, then $\{U_i \times_U V \to V\}$ forms a covering of V.

A category equipped with a Grothendieck topology is called a **site**.

Let $\mathcal{C}$ be a site. Recall that a (Set-valued) presheaf on $\mathcal{C}$ is defined to be nothing but a functor from $\mathcal{C}^{\mathrm{op}}$ to Set. A presheaf $\mathscr{F}$ is called a **sheaf** if for every covering $\{U_i \to U\}$ in $\mathcal{C}$ the diagram

$$\mathscr{F}(U) \longrightarrow \prod_{i\in I} \mathscr{F}(U_i) \rightrightarrows \prod_{i,j\in I} \mathscr{F}(U_{i_0} \times_U U_{i_1})$$

is an equalizer diagram. Here the two right arrows are the two obvious projection maps. In exactly the same way, one can define sheaves of rings, sheaves of modules, sheaves of algebras, and so on. A **morphism of sheaves** is defined in the same way as a morphism of presheaves.

For every object X, the functor $h_X = \mathrm{Hom}(\,\cdot\,, X)\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Set}$ is a presheaf. In general, it need *not* be a sheaf. A site for which all representable functors are sheaves is also called **subcanonical**. The sites that we are interested in are all subcanonical.

Example 2.2.1. There are many examples of Grothendieck topologies on the category Sch/S of schemes over a fixed base scheme S. We give two of them. The most well-known one is the **Zariski site**, where we declare a collection of morphisms to be a covering if the underlying maps of topological spaces form an open covering with respect to the Zariski topology. We sometimes denote this site by $(\mathsf{Sch}/S)_{\mathrm{Zar}}$.

Next we define the fpqc topology, following [14, Section 2.3.2]. Let $\pi\colon X \to Y$ be a morphism of schemes. We call it **fpqc** if it is faithfully flat, and in addition either of the two following equivalent properties are satisfied.

- Every quasi-compact open subset of Y is the image of a quasi-compact open subset of X;
- there exists an affine open covering $\{\mathrm{Spec}\, R_i\}$ of Y such that each $\mathrm{Spec}\, R_i$ is the image of a quasi-compact open subset of X.

Notice in particular that, if Y is affine and X is quasi-compact, an fpqc morphism is just a faithfully flat morphism. We now define the **fpqc topology** on Sch/S to be the topology in which the coverings $\{U_i \to U\}$ are the collections of morphisms such that the induced morphism $\bigsqcup_i U_i \to U$

is fpqc. Denote the resulting site by $(\mathsf{Sch}/S)_{\text{fpqc}}$. The fpqc site is subcanonical. This is a non-trivial result, and is sometimes called **descent of morphisms**. A proof can be found in [11, Tag 022H].

If we want, we can impose a finiteness constraint on the fpqc topology. That is, we can choose to declare $\{U_i \to U\}_{i\in I}$ to be an fpqc covering if the map $\bigsqcup_i U_i \to U$ is fpqc, and *in addition* the indexing set I is finite. Whether we do this or not won't make a difference anywhere until all the way in Lemma 4.2.6, where finiteness is needed and will be imposed from that point on.

Having defined a sheaf over an arbitrary site, we turn to stacks now. Essentially, a stack is almost the same as a sheaf, except it takes values in categories rather than in sets. There are two well-known ways of formalizing this, namely via fibred categories and via pseudo-functors. We briefly consider the definitions, but we assume that the reader is familiar with both approaches, and we will swap between one and the other wherever it is convenient. More details can be found in [14].

Given a category $\mathcal{C}$, a **(contravariant) pseudo-functor** F on $\mathcal{C}$ consists of the following data.

- For every object U in $\mathcal{C}$, a category $F(U)$;
- for every morphism $f\colon U \to V$, a functor $f^*\colon F(V) \to F(U)$;
- for every object U, a natural isomorphism of functors $\varepsilon_U\colon \mathrm{Id}_U^* \xrightarrow{\sim} \mathrm{Id}_{F(U)}$;
- for any two composable morphisms $f\colon U \to V$ and $g\colon V \to W$, a natural isomorphism $\alpha_{f,g}\colon f^* \circ g^* \xrightarrow{\sim} (g \circ f)^*$.

We ask for the natural isomorphisms α's to satisfy the cocycle conditions, and to be compatible with the ε's. More precisely, for composable arrows $U \xrightarrow{f} V \xrightarrow{g} W \xrightarrow{h} X$ and $\eta \in F(X)$, we ask for

$$\begin{array}{ccc} f^*g^*h^*\eta & \xrightarrow{\alpha_{f,g}(h^*\eta)} & (g\circ f)^*h^*\eta \\ \downarrow{\scriptstyle f^*\alpha_{g,h}(\eta)} & & \downarrow{\scriptstyle \alpha_{g\circ f,h}(\eta)} \\ f^*(h\circ g)^*\eta & \xrightarrow[\alpha_{f,h\circ g}(\eta)]{} & (h\circ g\circ f)^*\eta \end{array}$$

to commute, and for all morphisms $f\colon U \to V$ and elements $\eta \in F(V)$, we ask for the equalities $\alpha_{\mathrm{Id}_U,f}(\eta) = \varepsilon_U(f^*\eta)$ and $\alpha_{f,\mathrm{Id}_V}(\eta) = f^*\varepsilon_V(\eta)$.

Fix a category $\mathcal{C}$. By a **category over** $\mathcal{C}$ we mean a category $\mathcal{F}$ together with a functor $p_{\mathcal{F}}\colon \mathcal{F} \to \mathcal{C}$. Rather than writing $p_{\mathcal{F}}(f) = c$ for objects f in $\mathcal{C}$ and c in $\mathcal{C}$, we use the shorthand $f \mapsto c$. A morphism $\phi\colon f \to g$ of $\mathcal{F}$ will be called a **Cartesian morphism** if for any solid diagram

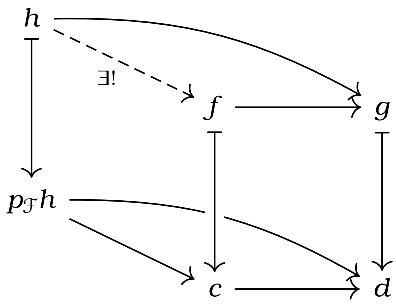

the dashed arrow can be made to exist and commute with everything else; in particular, the image of the dashed arrow under $p_{\mathcal{F}}$ should be the arrow $p_{\mathcal{F}}h \to c$ that we started out with. Moreover, the dashed arrow should be unique. We also say that f is a **pullback** of g along c.

A **fibred category** is now a category $\mathcal{F}$ over $\mathcal{C}$ such that, given a morphism $c \to d$ in $\mathcal{C}$ and an object g such that $g \mapsto d$, one can always find a Cartesian morphism $f \to g$ making the diagram

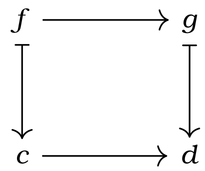

commute. In other words, we can always pull back objects of $\mathcal{F}$ along any morphism of $\mathcal{C}$. Given a category $\mathcal{F}$ fibred over $\mathcal{C}$. Take an object c in $\mathcal{C}$. The **fibre** $\mathcal{F}(c)$ is the subcategory of $\mathcal{C}$ whose objects are objects of $\mathcal{F}$ that are sent to c under $p_{\mathcal{F}}$, and whose arrows are those arrows of $\mathcal{F}$ that are sent to Id_c under $p_{\mathcal{F}}$.

Lemma 2.2.2. Let $\mathcal{F}$ be a fibred category over $\mathcal{C}$. If all the fibres $\mathcal{F}(c)$ are groupoids rather than arbitrary categories, or for short, if $\mathcal{F}$ is **fibred in groupoids**, then Cartesian morphisms are unique in the sense that, for any morphism in $\mathcal{C}$ there exists a unique Cartesian morphism in $\mathcal{C}$ lying above it. Conversely, if Cartesian morphisms are unique in the aforementioned sense, then the $\mathcal{F}(c)$ are groupoids. □

A **cleavage** of a fibred category $\mathcal{F} \to \mathcal{C}$ is a class of Cartesian morphism in $\mathcal{F}$ such that for each arrow $c \to d$ in $\mathcal{C}$ and for each object $\eta \in \mathcal{F}(d)$ there is a unique Cartesian morphism in said cleavage which provides a lift. By the axiom of choice, every fibred category has a cleavage. As the above lemma tells us, if the fibred category is fibred in groupoids, then there is only one cleavage, and the notion becomes irrelevant.

Lemma 2.2.3. Given a fibred category $\mathcal{F}$ over $\mathcal{C}$, with a fixed choice of cleavage, there exists a pseudo-functor on $\mathcal{C}$ which sends an object c to the fibre $\mathcal{F}(c)$. Conversely, every pseudo-functor gives rise to a fibred category with a cleavage. □

A **morphism of fibred categories** $F\colon \mathcal{F} \to \mathcal{G}$ is a functor F from $\mathcal{F}$ to $\mathcal{G}$ which commutes (*not* just up to natural isomorphism) with the projections $p_{\mathcal{F}}$ and $p_{\mathcal{G}}$, and which sends Cartesian morphism to Cartesian morphism. Similarly one can define morphisms of pseudo-functors, but we will omit this.

From this point on we assume that $\mathcal{C}$ is equipped with a Grothendieck topology. Let $\mathcal{F}$ be a fibred category over a site $\mathcal{C}$, and fix a cleavage for $\mathcal{F}$. Consider a covering $\mathcal{U} = \{\sigma_i \colon U_i \to U\}$ of an object U in $\mathcal{C}$. Denote by U_{ij} the object $U_i \times_U U_j$, and similarly write U_{ijk} instead of $U_i \times_U U_j \times_U U_k$. For fixed i and j, we use the following notation to denote the relevant maps:

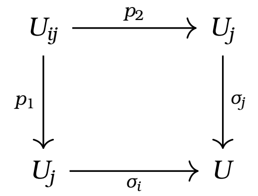

Notice that the notation is not in one-to-one correspondence with the maps, as we use the notation p_1 and p_2 regardless of the choice of i and j. Context should make it clear which map we mean. Similarly, for fixed i, j and k, we shall write p_{12}, p_{13}, and p_{23} for the maps in the diagram

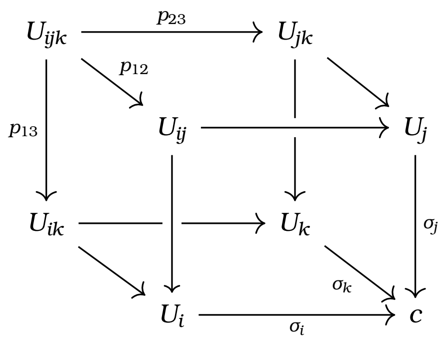

We now have the language to make the aforementioned notion of 'local data' precise. We define **descent data** at U to consist of the following. For all i, a choice of object ξ_i in the fibre $\mathcal{F}(U_i)$; for all i and j, a choice of isomorphism ϕ_{ij} from the pullback $p_1^*\xi_j$ to the pullback $p_2^*\xi_i$ — recall here that the pullbacks are determined by our fixed choice of cleavage. The isomorphisms must moreover satisfy the following cocycle conditions: for all i, j and k, we have $p_{13}^*\phi_{ik} = p_{12}^*\phi_{ij} \circ p_{23}^*\phi_{jk}$. A **morphism of descent data** $(\xi_i, \phi_{ij}) \to (\eta_i, \psi_{ij})$ to be a collection of arrows $a_i \colon \xi_i \to \eta_i$ in $\mathcal{F}(U_i)$, such that for every pair i, j, the diagram

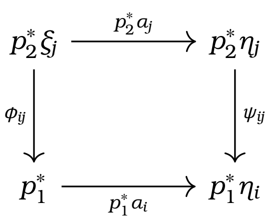

commutes. This turns the collection of descent data at an object c into a category, which we denote by $\mathcal{F}(\mathcal{U})$.

For each object $\xi \in \mathcal{F}(U)$, there is a natural choice of descent data for ξ, simply by letting the objects ξ in $\mathcal{F}(U_i)$ be $\sigma_i^*\xi_i$. The desired properties are then in fact trivially satisfied. Similarly if we have a morphism $\xi \to \eta$ in $\mathcal{F}(U)$, then we get arrows on the canonical choices of descent data. This yields a functor $\mathcal{F}(U) \to \mathcal{F}(\mathcal{U})$. With these definitions in place, we say $\mathcal{F}$ is a **stack** over $\mathcal{C}$ if the functor $\mathcal{F}(U) \to \mathcal{F}(\mathcal{U})$ is an equivalence of categories. We point out that the notion of a stack is well-defined in the sense that it does not depend on the choice of cleavage.

Example 2.2.4. In the case that $\mathcal{C}$ is subcanonical, a representable functor h_X is a sheaf, so that it may also be viewed as a stack, which we denote by $\mathcal{C}/X$ or even simply by X if no confusion can arise, and which we refer to as a **representable stack**.

There are two general constructions that we will need later on: stackification and 2-fibre products. They are really just stack-theoretic analogues of sheafification and ordinary fibre products, respectively, so no new ideas are going to come up until the end of this section.

We first deal with stackification. We just state the general claim, referring the reader to [11, Tag 02ZN] for more details. Given a fibred category $\mathcal{F}$ over a site $\mathcal{C}$, there is always a stack $\mathcal{F}'$ over $\mathcal{C}$, along with a morphism of fibred categories $\mathcal{F} \to \mathcal{F}'$, such that two properties hold.

- The first property is like a descent condition for objects. For any element c in the fibre $\mathcal{F}'(U)$, there exists a covering $\{h_i \colon U_i \to U\}$ such that $h_i^* c$ is in the essential image of the functor $\mathcal{F} \to \mathcal{F}'$.
- The second property is a descent condition for morphisms. Take two objects c and c' in a fibre $\mathcal{F}(U)$. For any object $h \colon V \to U$, we can consider $\mathrm{Hom}_{\mathcal{F}}(h^* c, h^* c')$, which we may interpret as a presheaf over $\mathcal{C}/U$. The morphism $\mathcal{F} \to \mathcal{F}'$ yield natural maps $\mathrm{Hom}_{\mathcal{F}}(h^* c, h^* c') \to \mathrm{Hom}_{\mathcal{F}'}(h^* c, h^* c')$, and we require this to be precisely the sheafification of our presheaf.

Next, we deal with (2-)fibre products. $\mathcal{F} \times_{\mathcal{H}} \mathcal{G}$. Consider two morphisms of fibred categories $f \colon \mathcal{F} \to \mathcal{H}$ and $g \colon \mathcal{G} \to \mathcal{H}$ over a fixed category $\mathcal{C}$. We aim to construct the **2-fibre product** $\mathcal{F} \times_{\mathcal{H}} \mathcal{G}$. As fibred categories form a 2-category, it is natural to relax any strict commutativity requirements, and replace them by commutativity up to some natural isomorphism. The description goes as follows. Its objects are all quadruples (U, x, y, F) with $U \in \mathcal{C}$, $x \in \mathcal{F}(U)$, $y \in \mathcal{G}(U)$, and with F an isomorphism $f(x) \to g(y)$ in $\mathcal{H}(U)$. The morphisms $(U, x, y, F) \to (U', x', y', F')$ are given by a morphism $\alpha \colon x \to x'$, a morphism $\beta \colon y \to y'$, so that $p_{\mathcal{F}}(\alpha) = p_{\mathcal{G}}(\beta)$, and so that

$$\begin{array}{ccc} f(x) & \xrightarrow{F} & g(y) \\ {\scriptstyle f(\alpha)}\downarrow & & \downarrow{\scriptstyle g(\beta)} \\ f(x') & \xrightarrow[F']{} & g(y') \end{array}$$

commutes in $\mathcal{H}(U)$. The resulting category satisfies the universal property that defines a 2-limit, but we omit these technicalities.

Lemma 2.2.5. Let the notation be as above. If $\mathcal{F}$, $\mathcal{G}$ and $\mathcal{H}$ are fibred in groupoids, then so is the 2-fibre product $\mathcal{F} \times_{\mathcal{H}} \mathcal{G}$. □

Proof: A morphism in a fibre $(\mathcal{F} \times_{\mathcal{H}} \mathcal{G})(U)$ consists of a morphism α between objects in $\mathcal{F}(U)$, and a morphism β between objects in $\mathcal{G}(U)$, satisfying some properties. As $\mathcal{F}(U)$ and $\mathcal{G}(U)$ are groupoids, α and β will be invertible. This gives us an obvious choice for an inverse for the morphism we started out with; the verification that it is indeed an inverse is straightforward. ■

Assume now that $\mathcal{C}$ is a site, and suppose that $\mathcal{F}$, $\mathcal{G}$ and $\mathcal{H}$ are stacks over $\mathcal{C}$ rather than just fibred categories. The above lemma tells us that the fibre product is again fibred in groupoids. If the fibre product does not satisfy the descent condition that defines a stack, we might need to apply a stackification procedure to get the appropriate 2-limit in the 2-category of stacks over $\mathcal{C}$. But as it turns out, this is not needed.

Lemma 2.2.6. Let $\mathcal{C}$ be a site, and let $f\colon \mathcal{F} \to \mathcal{H}$ and $g\colon \mathcal{G} \to \mathcal{H}$ be two morphisms of stacks over $\mathcal{C}$. Then the 2-fibre product $\mathcal{F} \times_{\mathcal{H}} \mathcal{G}$ as defined above is again a stack. □

Proof sketch: The proof is straightforward but tedious. Take a covering $\mathcal{U} = \{h_i\colon U_i \to U\}$ in $\mathcal{C}$, and start with a descent datum in $\mathcal{F} \times_{\mathcal{H}} \mathcal{G}$ relative to $\mathcal{U}$. When writing out the definition, find that this descent datum gives rise to descent data for $\mathcal{F}$ and $\mathcal{G}$. As these are assumed to be stacks, these descent data give rise to unique (up to isomorphism) elements $x \in \mathcal{F}(U)$ and $y \in \mathcal{G}(U)$ that are compatible with the descent data.

Locally, there exists from $f(x)$ to $g(y)$, and since $\mathcal{H}$ too is a stack, this glues to a unique isomorphism $f(x) \to g(y)$. The elements x and y, along with this isomorphism, comprise the data needed to define the desired element in $(\mathcal{F} \times_{\mathcal{H}} \mathcal{G})(U)$ that corresponds to the descent datum we started with. ■

2.3 Algebraic stacks

We are now ready to introduce algebraic stacks. Much of our presentation will be based on various sections in [11], and one can find much more information there if needed.

Let $\mathcal{C}$ be a subcanonical site, and let $f\colon \mathcal{F} \to \mathcal{G}$ be a morphism of stacks over $\mathcal{C}$. We say f is **(relatively) representable** if, for all representable stacks $\mathcal{C}/X$ and for all morphisms of stacks $\mathcal{C}/X \to \mathcal{G}$, the fibre product $\mathcal{F} \times_{\mathcal{G}} \mathcal{C}/X$ is again representable, and, as automatically follows by the Yoneda lemma, the morphism $\mathcal{F} \times_{\mathcal{G}} \mathcal{C}/X \to \mathcal{C}/X$ comes from a morphism between the representing objects. More generally, if T is a property that certain stacks over $\mathcal{C}$ satisfy, we can say f is representable by T if for all morphisms $\mathcal{C}/X \to \mathcal{G}$, the fibre product $\mathcal{F} \times_{\mathcal{G}} \mathcal{C}/X$ is of type T.

We can use relative representability to define properties for morphisms of stacks. If P is a property that certain morphisms in $\mathcal{C}$ satisfy and which is preserved under base change (so as to prevent pathologies), a morphism $\mathcal{F} \to \mathcal{G}$ of stacks can be said to satisfy P if it is relatively representable, and every base-changed morphism $\mathcal{F} \times_{\mathcal{G}} \mathcal{C}/X$ satisfies P.

We are now ready for the most important definitions in this sections. For us, an **algebraic stack** will be a stack fibred in groupoids over $\mathsf{Sch}_{\mathrm{fpqc}}$. An **Adams stack** is an algebraic stack such that the following two properties hold.

- There exists an affine scheme X, along with an fpqc morphism $X \to \mathcal{X}$, also called an **atlas** of $\mathcal{X}$ (that is, the map $X \to \mathcal{X}$ is relatively representable by fpqc morphisms of schemes);
- the diagonal morphism $\mathcal{X} \to \mathcal{X} \times \mathcal{X}$ is affine.

A **morphism of algebraic stacks** and a **morphism of Adams stacks** are defined in the same way as a morphism of general stacks.

A few remarks are in order. First, our definition of algebraic stack is very different from the one that can be found in the literature. What we call an Adams stack is much closer to what algebraic geometers would call an algebraic stack, except that we do not impose any finiteness condition on either the diagonal or the atlas.

Second, from this point on, all stacks that we will be interested in are, at least, algebraic stacks. We will therefore be a little bit sloppy, and write ‘stack’ when it should really be ‘algebraic stack’.

Finally, we have defined our stacks over Sch, but in practice it suffices to consider the stack over the subcategory Aff of affine schemes. Indeed, as stacks satisfy the descent condition, we need only describe the data locally. We will make use of this quite often, and implicitly identify the fibred category defining an algebraic stack with the subcategory over Aff.

Lemma 2.3.1. Let $\mathcal{F}$ be an algebraic stack. Then for any ring R, elements of the fibre $\mathcal{F}(\operatorname{Spec} R)$ correspond to morphisms of stacks $\operatorname{Spec} R \to \mathcal{F}$. □

In practice, we won't distinguish between an element of $\mathcal{F}(\operatorname{Spec} R)$ and its corresponding morphism $\operatorname{Spec} R \to \mathcal{F}$. It is perhaps helpful to note that, if $\mathcal{F}$ is a representable stack, the correspondence is precisely the one you'd expect.

Proof sketch of Lemma 2.3.1: Start with an element $f \in \mathcal{F}(\operatorname{Spec} R)$. We describe the corresponding morphism of stacks $\overline{f}\colon \operatorname{Spec} R \to \mathcal{F}$. To do so, for any scheme X we need to describe a functor $f_X\colon h_X(R) \to \mathcal{F}(X)$ and for any morphism $X \to X'$, there should be a natural transformation $f_X \to f_{X'}$. We describe f_X as follows. Given $a \in h_X(R)$ corresponding to a map of schemes $a\colon X \to \operatorname{Spec} R$, let $f_X(a)$ be the element $a^*(f)$ in $\mathcal{F}$, where a^* is the (unique) Cartesian morphism in $\mathcal{F}$ over a. Conversely, given a morphism $\overline{f}\colon \operatorname{Spec} R \to \mathcal{F}$, our element $f \in \mathcal{F}(\operatorname{Spec} R)$ can be retrieved by evaluating $\overline{f}$ at the element of $h_R(R)$ representing the identity map. ■

It is about time that we talk about examples of algebraic stacks. There are many examples of stacks that we do not have the time to consider, such as the moduli stack of elliptic curves or Picard stacks. The reader can find many examples in [11, Ch. 89] and elsewhere.

Example 2.3.2. Quotient stacks are an important family of examples of algebraic stacks. We follow [11, Tag 04UI]. Let (U, R) be a groupoid in the category of schemes over a base scheme S. Define the **quotient stack** $[U/R]$ to be the stackification of the pseudo-functor $\mathsf{Sch}/S \to \mathsf{Grpd}$ sending an S-scheme X to the groupoid $(\mathrm{Hom}(X, U), \mathrm{Hom}(X, R))$.

There are three important special cases that we will care about. The first example is when U and R are affine schemes over the base scheme $S = \operatorname{Spec} \mathbb{Z}$, thus forming a Hopf algebroid. The resulting quotient stack turns out to be an Adams stack. We discuss this in detail in Section 2.5.

The second special case goes as follows. Let G be a group S-scheme acting on an S-scheme X. The pair $(X, G \times X)$ has the structure of a groupoid object in the category of schemes (but it might be difficult to see what the maps should be — see [11, Tag 0444]). The resulting quotient stack is denoted $[X/G]$. In [11, Tag 04UV], it is shown that this stack admits the more familiar description involving G-torsors over X.

The third special case is really a subcase of the situation above. Suppose X is the scheme S, and the action of G on S is trivial. The resulting quotient stack $[S/G]$ is also called the **classifying stack** and is denoted by $\mathrm{B}G$. Unravelling the definition, it may simply be interpreted as the stack of G-torsors over S.

2.4 Quasi-coherent sheaves and substacks

In this section we define quasi-coherent sheaves over algebraic stacks. There are several equivalent ways to approach this, and we pick a particular approach that will simplify future considerations. We begin by defining a particular example of an algebraic stack.

Example 2.4.1. Let $\mathcal{C}$ be the category Sch of schemes endowed with the Zariski topology (see Example 2.2.1), which takes as its coverings the topological open coverings of the underlying spaces. We take $\mathcal{F}$ to be the category QCoh of all pairs $(X, \mathscr{F})$ where X is a scheme and $\mathscr{F}$ is a quasi-coherent sheaf over X. The functor $\mathsf{QCoh} \to \mathsf{Sch}$ is given by the obvious projection map. One can verify that this construction defines a stack over the Zariski site $\mathsf{Sch}_{\mathrm{Zar}}$ fibred in groupoids. Remarkably, the category QCoh defines a stack not only over $\mathsf{Sch}_{\mathrm{Zar}}$, but even over $\mathsf{Sch}_{\mathrm{fpqc}}$. This is called **descent of quasi-coherent sheaves**. A proof can be found in [11, Tag 023R] or [14, Thm. 4.23].

Let $\mathcal{F}$ be an algebraic stack. A **quasi-coherent sheaf** (or just **sheaf** for short if quasi-coherence is clear from context) on $\mathcal{F}$ is just a morphism $\mathcal{F} \to \mathsf{QCoh}$. The reader can check that, in the case that $\mathcal{F}$ is a representable stack, this coincides with the usual definition. We also point out that, by descent, it suffices to describe the morphism $\mathcal{F} \to \mathsf{QCoh}$ as stacks over Aff. A **morphism of quasi-coherent sheaves** is simply a natural transformation between the two functors $\mathcal{F} \to \mathsf{QCoh}$ defining the sheaves. This allows us to define a category $\mathsf{QCoh}(\mathcal{F})$ of quasi-coherent sheaves over our algebraic stack. As with schemes, we have the following result.

Lemma 2.4.2. For any algebraic stack $\mathcal{F}$, the category $\mathsf{QCoh}(\mathcal{F})$ of quasi-coherent sheaves on $\mathcal{F}$ forms an abelian category. □

Proof: The idea is reasonably straightforward. Given a natural transformation f between two functors $\mathscr{G}, \mathscr{H} \colon \mathcal{F} \to \mathsf{QCoh}$ we can define a kernel $\operatorname{Ker} f \colon \mathcal{F} \to \mathsf{QCoh}$ of this natural transformation by sending an object c above a scheme X to the kernel of the map of sheaves $f(c) \colon \mathscr{G}(c) \to \mathscr{H}(c)$. Notice that, in the case where $\mathcal{F}$ is a representable stack, this definition is sensible because pullbacks commute with taking kernels. Defining cokernels works in the same way. We leave it to the most enthusiastic of readers to verify the remaining tedious nonsense. ■

In a similar way, one can define **coherent sheaves** over stacks to be morphisms to the stack Coh of coherent sheaves. The above proof carries over to show that the subcategory $\mathsf{Coh}(\mathcal{F})$ of coherent sheaves again forms an abelian category. For the rest of this section, we will focus on quasi-coherent sheaves, but much of the discussion also applies to coherent sheaves.

Quasi-coherent sheaves can be tensored, pulled back and, in some cases, pushed forward. Pullbacks are defined quite easily. If $f \colon \mathcal{F} \to \mathcal{G}$ is a morphism of algebraic stacks, then a morphism $\mathcal{G} \to \mathsf{QCoh}$ can be composed with f to yield a sheaf on $\mathcal{F}$, which we refer to as the **pullback**. The definition of a **tensor product** of two sheaves $\mathcal{F} \to \mathsf{QCoh}$ is equally straightforward.

Now what about pushforwards? We do not expect them to exist always, as even in the case of schemes, pushforwards of quasi-coherent modules need not be quasi-coherent anymore. It suffices for us to describe pushforwards under the additional assumption that f is *affine* — that is, f is relatively representable by affine morphisms of schemes. We begin with a quasi-coherent sheaf $\mathcal{F} \to \mathsf{QCoh}$. We wish to describe a morphism $\mathcal{G} \to \mathsf{QCoh}$. It suffices to describe this morphism affine-locally. So take a morphism $\operatorname{Spec} R \to \mathcal{G}$. We wish to functorially associate an R-module M_R to this morphism. As f is affine, the morphism $\mathcal{F} \times_{\mathscr{G}} \operatorname{Spec} R \to \operatorname{Spec} R$ is affine, hence the pullback is in fact an affine scheme, say $\operatorname{Spec} S$. Associated to the morphism $\operatorname{Spec} S \to \mathcal{F}$ is an S-module M_S. We now define M_R to be the pushforward of M_S along the map $\operatorname{Spec} S \to \operatorname{Spec} R$; that is, it is the module whose underlying set is M_S, but whose R-module structure is induced by the ring map $R \to S$. This description gives rise to a morphism $\mathcal{G} \to \mathsf{QCoh}$, which we call the **pushforward sheaf**.

If quasi-coherent sheaves on algebraic stacks make sense, then in particular, so do quasi-coherent *ideal* sheaves, and hence also closed substacks. The definition is precisely as one would expect. Let $\mathcal{F}$ be a stack. A quasi-coherent **sheaf of ideals** on $\mathcal{F}$ is a quasi-coherent sheaf

$\mathscr{I} : \mathcal{F} \to \mathsf{QCoh}$ sending every object in $\mathcal{F}$ above X to a sheaf of ideals on X. Given a quasi-coherent sheaf $\mathscr{I}$ on $\mathcal{F}$, we define the **closed substack** $V(\mathscr{I})$ of $\mathcal{F}$ to be the full subcategory of $\mathcal{F}$ consisting only of those objects of $\mathcal{F}$ that get sent to the zero ideal by the morphism $\mathcal{F} \to \mathsf{QCoh}$. By descent of sheaves this indeed defines a stack. The inclusion of the closed substack into $\mathcal{F}$ is relatively representable by closed immersions, and in fact this characterizes closed substacks.

Let $\mathscr{I}$ be a quasi-coherent sheaf of ideals on $\mathcal{F}$. We would like to define what it means for $\mathscr{I}$ to be principal. We have to be somewhat careful here, as it is possible for an ideal sheaf on an affine scheme to be locally principal, but not globally so. (See [11, Tag 0CBZ] for an example.) We shall say $\mathscr{I}$ is **principal** if there exists an fpqc covering $\operatorname{Spec} A \to \mathcal{F}$ from an affine scheme such that the pullback of $\mathscr{I}$ to $\operatorname{Spec} A$ is cut out by a single element $f \in A$. Of course, it is not guaranteed that an fpqc covering $\operatorname{Spec} A \to \mathcal{F}$ exists in the first place, but in all situations where we need principal ideal sheaves, this will be the case.

Let $\mathcal{F}$ be a stack, with quasi-coherent ideal sheaf $\mathscr{I}$. We wish to define the open complement of the closed substack $V(\mathscr{I})$ in some way. Here we should be a bit careful: it is tempting to just take the complement of the subcategory $V(\mathscr{I})$ defined within the fibred category $\mathcal{F}$, but this is false. Indeed, the functor of points h_U of an open complement U of a closed subscheme $Z \subseteq X$ is not the complement of the functor of points h_Z within h_X. To motivate the correct definition, we investigate what the functor h_U should be. The following result is elementary and added for completeness.

Lemma 2.4.3. Take an affine scheme $\operatorname{Spec} R$, and let I be an ideal of R defining a closed subscheme $\operatorname{Spec} R/I$. We claim that the morphisms $\operatorname{Spec} A \to \operatorname{Spec} R$ factoring through the open complement of $\operatorname{Spec} R/I$ are precisely those ring maps $f \colon R \to A$ such that the ideal generated by the image $f(I)$ is all of A. □

Proof: Let's first suppose that the ring map f satisfies the property that $f(I)$ generates A. Take a point of $\operatorname{Spec} A$ corresponding to a prime ideal $\mathfrak{p}$ of A. If the image of this point were in $\operatorname{Spec} R/I$, it would imply that $f^{-1}(\mathfrak{p})$ would contain I, which in turn would tell us that $f(I) \subseteq f(f^{-1}(\mathfrak{p})) \subseteq \mathfrak{p}$ — a contradiction. Now assume that the ideal generated by $f(I)$ is not all of A. An argument invoking Zorn's lemma shows $f(I)$ is contained within some maximal ideal, say $\mathfrak{m}$. The pre-image $f^{-1}(\mathfrak{m})$ is an ideal of R containing I, hence the map of spectra corresponding to f cannot factor through the open complement. ■

We use the above lemma to motivate the definition of an **open complement** of a closed substack $V(\mathscr{I})$ of some stack $\mathcal{F}$. For simplicity, view our stacks as stacks fibred over Aff. Define the open complement as the full subcategory of $\mathcal{F}$ consisting of those maps $\operatorname{Spec} R \to \mathcal{F}$ such that the R-ideal we get by applying the morphism $\mathscr{I} : \mathcal{F} \to \mathsf{QCoh}$ is all of R.

Having defined closed and open substacks, there's one more flavour of substack that we'd like to consider. Recall that we can define the formal completion of a closed subscheme as a certain formal scheme that, intuitively, consists of the closed subscheme along with information about the neighbourhood of the subscheme within the ambient scheme. This construction has a generalization to algebraic stacks. We quickly review the classical construction before writing down our generalization to stacks.

Take an affine scheme $\operatorname{Spec} R$, with a closed subscheme $Z = Z_1$ defined by an ideal I of R. Let Z_2 be the closed subscheme defined by the ideal I^2. Its underlying points are the same of Z, but the nilpotence of the quotient R/I^2 yields some infinitesimal information about the neighbourhood of Z within $\operatorname{Spec} R$. We go on and define Z_3 by the ideal I^3. The ring R/I^3 contains yet more nilpotence, which may be thought of as yet more information about the neighbourhood. Continuing in this way, the formal completion may be thought of as the limit of this procedure.

There are two different ways of making this precise. First, we may define the **completion** of R to be the limit $\widehat{R}$ over the sequence of quotient maps $\cdots \to R/I^3 \to R/I^2 \to R/I$, and we may consider the spectrum $\operatorname{Spec} \widehat{R}$ of this completion. Alternatively, we may consider the functors of points $h_{Z_i} \colon \mathsf{Ring} \to \mathsf{Set}$ and define the **formal completion** to be the colimit along the sequence of functors $h_{Z_1} \to h_{Z_2} \to \cdots$. We denote the resulting functor of points by $\operatorname{Spf} \widehat{R}$. By [11, Tag 0AI2], this functor is an fpqc sheaf, so that we may fit formal schemes within our framework of algebraic stacks.

The second construction is more amenable to a generalization to algebraic stacks than the first one. This is because it is not clear at all how to view the completion construction of a ring from a functorial perspective. So let $\mathcal{F}$ be an algebraic stack with closed substack $V(\mathscr{I})$ defined by an ideal sheaf $\mathscr{I}$. Then the **formal completion** $\widehat{\mathcal{F}}$ is the full subcategory of $\mathcal{F}$ consisting of those objects that get sent to a locally nilpotent ideal by the morphism $\mathcal{F} \to \mathsf{QCoh}$. By [11, Tag 0BPF], this indeed defines an algebraic stack.

2.5 Flat Hopf algebroids and Adams stacks

Recall that Hopf algebroids are groupoid objects in the category of affine schemes, and so they give rise to a functor $\mathsf{Ring} \to \mathsf{Grpd}$, which is a fortiori a pseudo-functor. Similarly, stacks over the category of schemes may be defined as pseudo-functors $\mathsf{Ring} \to \mathsf{Grpd}$ satisfying a descent condition, and algebraic stacks are then required to satisfy some additional properties. By comparing the definitions, it seems reasonable to expect the two notions to be comparable. This is indeed the case, and we'll expand upon this now.

Let $\mathcal{M}$ be an Adams stack, and let $\operatorname{Spec} A \to \mathcal{M}$ be an fpqc covering of $\mathcal{M}$. Consider the 2-fibre product $\operatorname{Spec} A \times_{\mathcal{M}} \operatorname{Spec} A$. This is again an affine scheme. One way to see this is by noting that it also arises as the 2-fibre product of the diagram

$$\begin{array}{ccc} \operatorname{Spec} A \times_{\mathcal{M}} \operatorname{Spec} A & \longrightarrow & \operatorname{Spec} A \times \operatorname{Spec} A \\ \downarrow & & \downarrow \\ \mathcal{M} & \longrightarrow & \mathcal{M} \times \mathcal{M} \end{array}$$

and recalling that the diagonal map on $\mathcal{M}$ was affine by definition. Let us write $\operatorname{Spec}\Gamma$ for the fibre product for suggestive reasons. The pair $(\operatorname{Spec} A, \operatorname{Spec}\Gamma)$ forms a flat Hopf algebroid. The source and target maps $\operatorname{Spec}\Gamma \to \operatorname{Spec} A$ are the structure maps of the fibre product, and these maps are flat as flatness of the map $\operatorname{Spec} A \to \mathcal{M}$ is preserved by base change. The unit map $\operatorname{Spec} A \to \operatorname{Spec}\Gamma$ is the diagonal map. The inversion map $\operatorname{Spec}\Gamma \to \operatorname{Spec}\Gamma$ is obtained by swapping the components of the product $\operatorname{Spec} A \times_{\mathcal{M}} \operatorname{Spec} A$.

Finally, we have the composition $\operatorname{Spec}\Gamma \times_A \operatorname{Spec}\Gamma \to \operatorname{Spec}\Gamma$. It can be found as follows. We begin by writing down the diagram

$$\begin{array}{ccccc} \operatorname{Spec}\Gamma \times_A \operatorname{Spec}\Gamma & \longrightarrow & \operatorname{Spec}\Gamma & \longrightarrow & \operatorname{Spec} A \\ \downarrow & & \downarrow & & \\ \operatorname{Spec}\Gamma & \longrightarrow & \operatorname{Spec} A & & \\ \downarrow & & & & \\ \operatorname{Spec} A & & & & \end{array}$$

The top horizontal maps and the left vertical maps from $\operatorname{Spec}\Gamma \times_A \operatorname{Spec}\Gamma$ to $\operatorname{Spec} A$ naturally yield a map to $\operatorname{Spec}\Gamma$, which is the one we declare to be the composition. It remains to be shown that all the axioms defining a Hopf algebroid are satisfied. All these verifications are entirely formal, and we omit them.

Let us now start out with a flat Hopf algebroid $(\operatorname{Spec} A, \operatorname{Spec}\Gamma)$. As mentioned before, it gives rise to a pseudo-functor, and hence to a category fibred over Aff. We define $\mathcal{M}_{(A,\Gamma)}$ to be the algebraic stack obtained by stackifying our fibred category with respect to the fpqc topology. As it turns out, this is an Adams stack, in precisely the way one would expect. There is an obvious map $\operatorname{Spec} A \to \mathcal{M}_{(A,\Gamma)}$, and by flatness over our Hopf algebroid, this becomes an fpqc covering; moreover, taking the fibre product over the diagram $\operatorname{Spec} A \to \mathcal{M}_{(A,\Gamma)} \leftarrow \operatorname{Spec} A$ yields $\operatorname{Spec}\Gamma$.

In fact, one can use this approach to prove something stronger. The category of Adams stacks with a fixed choice of fpqc atlas has the structure of a 2-category, as does the category of flat Hopf algebroids. The association above defines an equivalence of 2-categories, where equivalence should be interpreted in an appropriate 2-categorical sense. This can be found in [8, Section 3].

Example 2.5.1. Recall from Example 2.1.5 that certain ring spectra, which we called flat ring spectra, give rise to a Hopf algebroid. That is, if E is a flat ring spectrum, then $(\pi_* E, \widetilde{E}_* E)$ has the structure of a Hopf algebroid. We repeat here that, in order for our constructions to fit within the framework of algebraic geometry, our rings need to be commutative, so the ring spectrum also needs to be evenly graded. If this is the case, then our correspondence tells us that E gives rise to an Adams stack, which we will denote by $\mathcal{M}_E$.

Comodules over Hopf algebroids have a natural interpretation in our stack-theoretic language. They are precisely quasi-coherent sheaves over the corresponding algebraic stacks. The proof is really just an application of fpqc descent of quasi-coherent sheaves (see Lemma 4.2.9 for details). For the sake of completeness, we'll sketch how to pass from one realm to the other.

Start with a quasi-coherent sheaf $\mathscr{F}$ on an Adams stack $\mathcal{M}_{(A,\Gamma)}$. The pullback of $\mathscr{F}$ along the atlas $\operatorname{Spec} A \to \mathcal{M}_{(A,\Gamma)}$ gives rise to an A-module. This A-module is the underlying module structure of our comodule, and we can find the coaction by considering the pullback of $\mathscr{F}$ along the two maps $\operatorname{Spec} \Gamma \to \mathcal{M}_{(A,\Gamma)}$. Conversely, if we start out with a comodule M and we'd like to find $\mathscr{F}$, the underlying A-module tells us what the pullback of $\mathscr{F}$ along the map $\operatorname{Spec} A \to \mathcal{M}_{(A,\Gamma)}$ should be, while the comodule structure gives us the desired descent data on pullbacks and triple overlaps. For this, too, more details can be found at [8, Section 3].

Chapter 3

The moduli stack of formal groups

The moduli stack of formal groups is the algebraic stack classifying formal group laws and isomorphisms between them. We start out in this chapter by giving several constructions of this moduli stack of formal groups, before proceeding to study some of its properties. We'll find, among other things, that the height of formal group laws gives rise to a filtration by substacks of this moduli stack, and the various layers of this filtration will end up being easier in nature than the entire stack.

3.1 Constructions of the stack of formal groups

The moduli stack of formal groups can be constructed in several ways, all of which will end up being equivalent. We spend a bit of time giving the various definitions.

We recall from Section A.2 that, for any ring R, the formal group laws over R correspond to ring maps $\mathrm{Hom}(L, R)$, and the isomorphisms between any two formal group laws over R correspond to ring maps $\mathrm{Hom}(W, R)$. This turns the pair $(\mathrm{Spec}\, L, \mathrm{Spec}\, W)$ into a Hopf algebroid. Via the correspondence in Section 2.5, this gives rise to an Adams stack, denoted $\mathcal{M}_{\mathrm{FG}}$, and called the **moduli stack of formal groups**. More explicitly, the moduli stack of formal groups is the stack associated to the pseudofunctor $\mathsf{Ring} \to \mathsf{Grpd}$ sending a ring R to the groupoid of formal group laws over R, along with their isomorphisms.

The stackification in this definition is not redundant. Let R be a ring with non-trivial Picard group, and take a non-trivial line bundle $\mathscr{L}$ on $\mathrm{Spec}\, R$ — for concreteness, take R to be a Dedekind domain and $\mathscr{L}$ a fractional ideal. Let $\{U_i\}$ be a trivializating cover. Assign, to each U_i, the additive formal group law, and glue this data using the gluing data of $\mathscr{L}$. The resulting assignment does not come from a single formal group law over R. One way to check this directly is via [5, Prop. 7 of Lecture 11].

The stack of formal groups can also be constructed as a certain quotient stack. Start with the Lazard ring L. Write $G = \operatorname{Spec} \mathbb{Z}[a_0^{\pm 1}, a_1, \ldots]$. Then, for any ring R, $\operatorname{Hom}(\operatorname{Spec} R, G)$ is naturally isomorphic to $\{h(t) \in R[\![t]\!] : h'(0) \in R^*\}$, which has a group structure by composition of power series, thus turning G into an affine group scheme. For any R, $G(R)$ acts on $L(R) = \mathrm{FGL}(R)$ via the map $(h(t), f(x,y)) \mapsto h^{-1}(f(h(x), h(y)))$. This yields an action of the group scheme G on $\operatorname{Spec} L$, and we may define the quotient stack $[\operatorname{Spec} L/G]$ as we did in Example 2.3.2.

There is a variation of the stack of formal groups that is worth mentioning. Recall that an isomorphism of formal group laws $h\colon f(x,y) \to g(x,y)$ is defined by a certain formal power series $h(t)$, where $h'(0)$ is necessarily a unit. If it equals 1, we said that h is a strict. We write W^{s} for the ring representing strict isomorphisms. The pair (L, W^{s}) is another Hopf algebroid whose associated stack is denoted $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$. There's an obvious 'inclusion' from $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ into $\mathcal{M}_{\mathrm{FG}}$.

The following slightly different perspective on the above definition is perhaps more enlightening to some. Recall from Example 2.1.5 that we associated Hopf algebroids, and hence Adams stacks, to flat ring spectra. Also recall from Theorem 1.3.9 that $\pi_*\mathrm{MU}$ is isomorphic to the Lazard ring. Perhaps, then, $(\pi_*\mathrm{MU}, \widetilde{\mathrm{MU}}_*\mathrm{MU})$ forms a Hopf algebroid that is similar to some stack of formal groups. This is indeed the case: by [13, Thm. 17.16], $\widetilde{\mathrm{MU}}_*\mathrm{MU} \cong L[a_2, a_3, \ldots]$, and there are indeed structure maps on the pair $(\pi_*\mathrm{MU}, \widetilde{\mathrm{MU}}_*\mathrm{MU})$ that turn it into the same Hopf algebroid as the one considered above. Thus, $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ is precisely the stack associated to the spectrum MU. For any space X, $\widetilde{\mathrm{MU}}_*(X)$ is a comodule over the Hopf algebroid, so that X gives rise to a quasi-coherent sheaf over $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$, which we denote $\mathscr{F}_X^{\mathrm{s}}$.

The non-strict version also comes from a homology theory, and in fact, it's almost the same as MU. Let us define MUP to be the infinite wedge sum $\bigvee_{n\in\mathbb{Z}} \Sigma^{2n}\mathrm{MU}$. Notice that the zero-th homotopy group is $\pi_0^{\mathrm{S}}(\mathrm{MUP}) = \bigoplus_{n\in\mathbb{Z}} \pi_{2n}^{\mathrm{S}}(\mathrm{MU}) = \pi_*^{\mathrm{S}}(\mathrm{MU}) = L$, where the last step uses the fact that MU is evenly graded. Similarly, we have

$$\begin{aligned}
\widetilde{\mathrm{MUP}}_0(\mathrm{MUP}) &= [\mathbb{S}, \mathrm{MUP} \wedge \mathrm{MUP}] \\
&= \left[\mathbb{S}, \bigvee_{n\in\mathbb{Z}} \bigvee_{m\in\mathbb{Z}} \Sigma^{2n+2m}(\mathrm{MU} \wedge \mathrm{MU})\right] \\
&= \bigoplus_{n\in\mathbb{Z}} \widetilde{\mathrm{MU}}_*(\mathrm{MU}) \\
&= \widetilde{\mathrm{MU}}_*(\mathrm{MU})[a_1^{\pm 1}] \\
&= L[a_1^{\pm 1}, a_2, a_3, \ldots] = W
\end{aligned}$$

This turns the zero-th graded part of $\pi_*^{\mathrm{S}}(\mathrm{MUP})$ and of $\widetilde{\mathrm{MUP}}_*(\mathrm{MUP})$ into an *ungraded* Hopf algebroid that is precisely (L, W). The algebraic stack associated to it must be $\mathcal{M}_{\mathrm{FG}}$. Moreover, for any space X, in the same way that $\mathrm{MUP}_*(X)$ would be a comodule over $(\pi_*^{\mathrm{S}}(\mathrm{MUP}), \widetilde{\mathrm{MUP}}_*(\mathrm{MUP}))$, so is the zero-th graded part a comodule over $(\pi_0^{\mathrm{S}}(\mathrm{MUP}), \widetilde{\mathrm{MUP}}_0(\mathrm{MUP})) = (L, W)$. In this way, every space gives rise to a quasi-coherent sheaf over $\mathcal{M}_{\mathrm{FG}}$, denoted $\mathscr{F}_X$.

The above discussion tells us that the difference between $\mathcal{M}_{\mathrm{FG}}$ and $\mathcal{M}^{\mathrm{s}}_{\mathrm{FG}}$ is essentially just a grading issue — something which, in hindsight, is already seen from the definitions. One can always associate, to a given evenly graded Hopf algebroid (A_*, Γ_*), an ungraded Hopf algebroid $(A, \Gamma[u^{\pm 1}])$, where the grading is implicitly being hidden within the variable u. As a special case, applying this to (L, W^{s}), viewed as a graded Hopf algebroid, we get the ungraded Hopf algebroid (L, W). At no point will the grading issue cause significant issues, but it is nonetheless profoundly annoying at times.

Finally, a word of caution. It is tempting to mentally picture $\mathcal{M}_{\mathrm{FG}}$ as being very similar to $\operatorname{Spec} L$; indeed, any formal group law over a ring R yields both a map $\operatorname{Spec} R \to \operatorname{Spec} L$ and a map $\operatorname{Spec} R \to \mathcal{M}_{\mathrm{FG}}$. This is a misleading way of thinking. Take two formal group laws, one over a ring R and the other over R'. Then the pullback along the maps $\operatorname{Spec} R \to \operatorname{Spec} L \leftarrow \operatorname{Spec} R'$ is given by $\operatorname{Spec}(R \otimes_L R')$; but if we replace $\operatorname{Spec} L$ by $\mathcal{M}_{\mathrm{FG}}$, we get something completely different.

Lemma 3.1.1. Take two formal group laws over R and R'. The pullback along the maps $\operatorname{Spec} R \to \mathcal{M}_{\mathrm{FG}} \leftarrow \operatorname{Spec} R'$ is given by $\operatorname{Spec}(R \otimes_L W \otimes_L R')$. Here we interpret W as a bimodule over L via the source and target maps $L \to W$. □

Proof: Let's view the pullback $\operatorname{Spec} R \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} R'$ as a fibred category. We unravel the definition of a 2-fibre product. An object of this pullback above a scheme $\operatorname{Spec} A$ (or, if you wish, a morphism $\operatorname{Spec} A \to \operatorname{Spec} R \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} R'$) should consist of a morphism $\operatorname{Spec} A \to \operatorname{Spec} R$, a morphism $\operatorname{Spec} A \to \operatorname{Spec} R'$, and an isomorphism between the two resulting formal group laws over A. Essentially by definition, this is just a morphism $\operatorname{Spec} A \to \operatorname{Spec}(R \otimes_L W \otimes_L R')$. This would prove the result, but we need to check if there aren't any non-trivial automorphisms in $\operatorname{Spec} R \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} R'$. Looking at the definition again, any such morphism should involve non-trivial morphisms in the fibred categories defining $\operatorname{Spec} R$ and $\operatorname{Spec} R'$, but as $\operatorname{Spec} R$ and $\operatorname{Spec} R'$ are discrete, there are none. ■

The moduli stack of formal groups is not a nice algebraic stack by algebro-geometric standards. It does not have a presentation by a scheme locally of finite type, it is not separated, and it does not admit a smooth atlas — at least, not a priori so. It is, however, the limit of 'nice' algebraic stacks, namely by algebraic stacks of truncated formal group laws. These truncated formal group laws carry the rather unappealing name of 'k-buds', where k refers to the point at which we truncate.

More precisely, a **k-bud** is a polynomial $f(x, y)$ in $R[\![x, y]\!]/(x, y)^{k+1}$ such that, modulo $(x, y)^{k+1}$, $f(x, y)$ satisfies the usual associativity, unitality, and commutativity properties. An **isomorphism of k-buds** $h\colon f(x, y) \to g(x, y)$ is described by an expression $h(t) \in R[\![t]\!]/(t^{k+1})$ such that, modulo (t^{k+1}), we have the usual identity $f(h(x), h(y)) = h(g(x, y))$. They assemble into a stack, called the **moduli stack of k-buds**, which we denote by $\mathcal{M}_{\mathrm{FG}}\langle k\rangle$. Results analogous to Theorem A.2.2 and Lemma A.2.3 carry over to this context, which shows that $\mathcal{M}_{\mathrm{FG}}\langle k\rangle$ is the

Adams stack associated to a Hopf algebroid $(L\langle k\rangle, W\langle k\rangle)$, where $L\langle k\rangle = \operatorname{Spec}\mathbb{Z}[x_1, \ldots, x_{k-1}]$ and $W\langle k\rangle = L\langle k\rangle[t_1^{\pm 1}, t_2, \ldots, t_k]$.

In the same way that formal group laws are a 'limit' of k-buds, so does the moduli stack of formal groups turn out to be a 'limit' of the various stacks of k-buds. More precisely, we have the following result, which we can find as [4, Thm. 3.20]:

Lemma 3.1.2. The natural maps $\mathcal{M}_{\mathrm{FG}} \to \mathcal{M}_{\mathrm{FG}}\langle k\rangle$, obtained by truncating formal group laws, give rise to an equivalence of stacks $\mathcal{M}_{\mathrm{FG}} \to \operatorname{holim} \mathcal{M}_{\mathrm{FG}}\langle k\rangle$. □

In the next sections, we take a look at the height filtration of $\mathcal{M}_{\mathrm{FG}}$, and at the geometric properties of the various layers of the filtration. The definition of height carries over to k-buds (although being of a certain height n requires the k in 'k-bud' to be sufficiently large). Consequently, virtually everything we do in the next section can also be done in the context of k-buds. For now, we'll omit the details of this, but it will become crucial in Chapter 4. We'll come back to this in Section 4.3.

3.2 The height filtration

In Section A.3 we introduced an invariant associated to formal group laws called the height. By restricting attention to formal group laws of a fixed height, we obtain substacks of the moduli stack of formal groups, which we can study one by one. This will be the most important aim of this section. The contents are based on [4, Ch. 5 and 6].

Throughout this section, we work at a fixed prime p. So, when we say "$\mathcal{M}_{\mathrm{FG}}$", we really mean $\mathcal{M}_{\mathrm{FG}} \times \operatorname{Spec}\mathbb{Z}_{(p)}$. And when we say "height" — a term which implicitly invokes a choice of a prime — we choose the same p. Also, we remark right away that our discussion also works for the stack $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ with strict isomorphisms, but we omit this as it involves nothing new.

For every positive integer n, we denote by $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ the moduli stack of formal groups of height $\geq n$, localized at p. That is, we take the functor sending a ring R to the groupoid of formal group laws of height $\geq n$ along with their isomorphisms, and stackify. As with $\mathcal{M}_{\mathrm{FG}}$, the reader can verify that we can describe $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ as an Adams stack associated to a Hopf algebroid. It is the Hopf algebroid (A, Γ), where A is the ring $L_{(p)}/(p, v_1, \ldots, v_{n-1})$, and Γ is the ring $A[t_0^{\pm 1}, t_1, \ldots]$.

Lemma 3.2.1. For every positive integer n, the algebraic stack $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ is a closed substack of the stack $\mathcal{M}_{\mathrm{FG}}$ of formal groups. □

Proof: It suffices to verify that $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ is a closed substack of $\mathcal{M}_{\mathrm{FG}}^{\geq n-1}$, where $\mathcal{M}_{\mathrm{FG}}^{\geq 0}$ is just $\mathcal{M}_{\mathrm{FG}}$. In order to verify this, we need to show that $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ is defined by a quasi-coherent ideal sheaf over $\mathcal{M}_{\mathrm{FG}}^{\geq n-1}$. By descent of sheaves, it suffices to describe the ideal sheaf as descent data on the fpqc covering $\operatorname{Spec} L_{(p)}/(p, v_1, \ldots, v_{n-1}) \to \mathcal{M}_{\mathrm{FG}}^{\geq n}$. We take the ideal $(v_n) \in L_{(p)}/(p, v_1, \ldots, v_{n-1})$, which may be pulled back to $W_{(p)}$ along the two structure maps $\mathrm{srce}, \mathrm{trgt}\colon L_{(p)} \to W_{(p)}$. It remains to be shown that these two pullbacks admit a $W_{(p)}$-module isomorphism satisfying cocycle conditions. The isomorphism is obtained by descending the structure map $\mathrm{inv}\colon W_{(p)} \to W_{(p)}$ to a map $W_{(p)}/(p, v_1, \dots, v_{n-1}) \to W_{(p)}/(p, v_1, \dots, v_{n-1})$, which can be done thanks to Lemma A.3.3. ■

The closed substacks we have just defined turn out to be fundamental to the structure of the stack of formal groups. In fact, this so-called height filtration is the unique filtration on $\mathcal{M}_{\mathrm{FG}}$, as is evidenced by the following result, found in [4, Thm. 5.13].

Theorem 3.2.2. The substacks $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ are reduced. Moreover, any reduced closed substack of $\mathcal{M}_{\mathrm{FG}}$ is either $\mathcal{M}_{\mathrm{FG}}$ itself or $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ for some n. □

Let us write $\mathcal{M}_{\mathrm{FG}}^{<n}$ for the open complement of $\mathcal{M}_{\mathrm{FG}}^{\geq n}$ inside $\mathcal{M}_{\mathrm{FG}}$. In the language of Section A.3, it is precisely the moduli stack of formal groups of height $< n$. As emphasized in the definition of open complements, introduced in Section 2.4, this does not mean we take the complementary subcategory of $\mathcal{M}_{\mathrm{FG}}^{\geq n}$.

Lemma 3.2.3. For all $n > 0$, the open substack $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ admits an fpqc covering by the affine scheme $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \dots, v_n^{\pm 1}]$. □

Proof: We define the map $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \dots, v_n^{\pm 1}] \to \mathcal{M}_{\mathrm{FG}}$ to be the map corresponding to the formal group law over $\mathbb{Z}_{(p)}[v_1, \dots, v_n^{\pm 1}]$ that is defined by the map $L_{(p)} \cong \mathbb{Z}_{(p)}[t_1, t_2, \dots] \to \mathbb{Z}_{(p)}[v_1, \dots, v_n^{\pm 1}]$ sending t_{p^i-1} to v_i for $i = 1, \dots, n$ and sending the other t_j to 0 — but see the discussion above Lemma 3.3.11, where we 'revisit' this definition. By Corollary A.3.10, we see that this formal group law is indeed of height $< n+1$. Both the definition we pick and the revisited definition are flat and affine. The former property follows from Theorem 3.3.4; as for the latter, it is in fact true that any map from an affine scheme is affine; indeed, if $\operatorname{Spec} A \to \mathcal{M}_{\mathrm{FG}}$ is any morphism, the pullback $\operatorname{Spec} R \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} A$ is precisely the pullback in the diagram

$$\begin{array}{ccc}
\operatorname{Spec} R \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} A & \longrightarrow & \operatorname{Spec} R \times \operatorname{Spec} A \\
\downarrow & & \downarrow \\
\mathcal{M}_{\mathrm{FG}} & \xrightarrow[\Delta]{} & \mathcal{M}_{\mathrm{FG}} \times \mathcal{M}_{\mathrm{FG}}
\end{array}$$

Being flat and affine, we need only check that the map is also surjective, which we do now. As surjectivity may be checked stalk-locally, it suffices to consider the pullback of our potential covering along any morphism $\operatorname{Spec} k \to \mathcal{M}_{\mathrm{FG}}^{<n+1}$ from the spectrum of a field, corresponding to a formal group law $f(x, y)$ of height $< n+1$ over k. We next point out that the result is trivial if $p = 0$, for then all primes are invertible in k so that Lemma A.3.6 applies. Thus we may assume k is of characteristic p. Let's write d for the height of our formal group law — this is well-defined as we are working over a field.

Once we show that there exists some morphism $\operatorname{Spec} A \to \operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}] \times_{\mathcal{M}_{FG}^{<n+1}} \operatorname{Spec} k$, we should be done, because this would imply that the pullback is non-empty and hence the map to $\operatorname{Spec} k$ is surjective on the underlying spaces. We take A to be the separable closure $\overline{k}$ of k. By Theorem A.3.12, the pullback of our formal group law $f(x, y)$ to $\overline{k}$ is isomorphic to any other formal group law of height d over $\overline{k}$. We pick a particularly suitable formal group law: take the formal group law $g(x, y)$ corresponding to the map $L_{(p)} \cong \mathbb{Z}_{(p)}[t_1, t_2, \ldots] \to \overline{k}$ sending t_{p^d-1} and t_{p^n-1} to 1 and sending all other t_i to 0. Using Corollary A.3.10 again, we see that $g(x, y)$ is of height d, and that moreover $v_n(g)$ is invertible. Essentially by construction, the map $L_{(p)} \to \overline{k}$ factors through $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}]$. Thus we find an induced map $\operatorname{Spec} \overline{k} \to \operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}] \times_{\mathcal{M}_{FG}^{<n+1}} \operatorname{Spec} k$, which proves the result. ■

The above result omits the open substack $\mathcal{M}_{FG}^{<1}$, but this one turns out to be particularly simple in nature. Looking at the definition, $\mathcal{M}_{FG}^{<1}$ classifies the formal group laws for which v_0, i.e. p, is invertible. As we are already working over $\mathbb{Z}_{(p)}$, this tells us that $\mathcal{M}_{FG}^{<1}$ is really just $\mathcal{M}_{FG} \times \operatorname{Spec} \mathbb{Q}$. As established in Lemma A.3.6, formal group laws over $\mathbb{Q}$-algebras are well understood, and this allows us to come up with the following result.

Lemma 3.2.4. The stack $\mathcal{M}_{FG}^{<1}$ is isomorphic to $B\mathbb{G}_m \times \operatorname{Spec} \mathbb{Q}$, where $\mathbb{G}_m$ denotes $\operatorname{Spec} \mathbb{Z}[x^{\pm 1}]$ with the usual group structure. □

Proof: A proof can be found in [4, Cor. 3.21]. We give an alternative proof here (although it boils down to the same idea). Let R be a $\mathbb{Q}$-algebra. We investigate the structure of the groupoid of formal group laws over R. Its objects will correspond to the formal group laws over R, and for any two formal group laws, the morphisms are parametrized by the unit group R^*, so that composition of any two morphisms correspond to multiplication of the unit group.

We first consider the automorphism group of the additive formal group law over R (where the group structure is composition). Suppose h is an automorphism, so that $h(x + y) = h(x) + h(y)$. Write $h(t) = h_1 t + h_2 t^2 + \cdots$, expand both sides, and compare the coefficients. One easily finds that $h_2 = h_3 = \cdots = 0$, so that $h(t) = ut$ for some unit $u \in R^*$.

Now take any formal group law $f(x, y)$ over R. In the proof of Lemma A.3.6, we saw that we can construct an isomorphism $\log_f$ between $f(x, y)$ and the additive formal group law, so that $\log_f(f(x, y)) = \log_f(x) + \log_f(y)$. For any unit u, $u \cdot \log_f$ again defines an isomorphism. We claim that there are no others. To do this, take two isomorphisms h_1 and h_2 between $f(x, y)$ and the additive formal group, and consider the composition $h_1 \circ h_2^{-1}$. This composition defines an automorphism of the additive formal group, which by the previous paragraph must be of the form ut with u a unit.

Next, consider the automorphism group of any formal group law $f(x, y)$ over R. Take an automorphism $h\colon f(x, y) \to f(x, y)$, and compose it with $\log_f$. By the previous paragraph, the result must be $u \cdot \log_f$ for some unit u; that is, $\log_f(h(t)) = u \cdot \log_f(t)$. Let $\exp_f(t)$ be the compositional inverse of $\log_f(t)$. Then $h(t)$ is determined to be $\exp_f(u \cdot \log_f(t))$. This establishes a bijection between $\mathrm{Aut}(f(x, y))$ and R^*, and one can check directly that it respects the group structures.

Finally, take two formal group laws $f(x, y)$ and $g(x, y)$, and consider an isomorphism from $f(x, y)$ to $g(x, y)$. For any two such isomorphisms, we can again compose one with the inverse of the other, and the previous paragraph tells us that the result will be parametrized by R^*. This allows us to conclude that the isomorphisms from $f(x, y)$ to $g(x, y)$ are parametrized by R^* as well. If $h(x, y)$ is a third formal group law, then isomorphisms $f(x, y) \to g(x, y)$ and $g(x, y) \to h(x, y)$, corresponding to two units, can be composed, and a direct verification shows that their composition corresponds to the product of the two units.

In order to prove that $\mathcal{M}_{\mathrm{FG}}^{<1}$ and $\mathrm{B}\mathbb{G}_m \times \operatorname{Spec}\mathbb{Q}$ are equivalent, it suffices to show that the prestacks we constructed to define these two stacks (namely, in Section 3.1 and Example 2.3.2) are already equivalent before stackifying. We define a functor from the prestack defining $\mathcal{M}_{\mathrm{FG}}^{<1}$ to the prestack defining $\mathrm{B}\mathbb{G}_m \times \operatorname{Spec}\mathbb{Q}$ by sending a formal group law over a $\mathbb{Q}$-algebra R to the unique object in the fibre $(\mathrm{B}\mathbb{G}_m \times \operatorname{Spec}\mathbb{Q})(\operatorname{Spec} R)$, while sending an isomorphism of formal group laws over R corresponding to a unit $u \in R^*$ to the morphism in the fibre $(\mathrm{B}\mathbb{G}_m \times \operatorname{Spec}\mathbb{Q})(\operatorname{Spec} R)$ corresponding to the same unit. This association defines a morphism of prestacks. It is easily seen to be essentially surjective, while fully faithfulness comes from our analysis of the groupoid $\mathcal{M}_{\mathrm{FG}}^{<1}(\operatorname{Spec} R)$. Thus we have an equivalence of prestacks, which remains an equivalence after stackifying. ∎

In a similar way, we define $\mathcal{M}_{\mathrm{FG}}^{n}$ for the open complement of $\mathcal{M}_{\mathrm{FG}}^{\geq n+1}$ within $\mathcal{M}_{\mathrm{FG}}^{\geq n}$. As $\mathcal{M}_{\mathrm{FG}}^{\geq n+1}$ is defined by the vanishing of the principal ideal sheaf (v_n) inside $\mathcal{M}_{\mathrm{FG}}^{\geq n}$, it follows that $\mathcal{M}_{\mathrm{FG}}^{n}$ is the moduli stack of formal groups of height exactly n. For the sake of giving different perspectives where we can, it is also the Adams stack associated to the Hopf algebroid (A, Γ) where A is the ring $(L_{(p)}/(p, v_1, \ldots, v_{n-1}))[v_n^{-1}]$, and Γ is the ring $A[t_0^{\pm 1}, t_1, \ldots]$.

The geometry of the substacks $\mathcal{M}_{\mathrm{FG}}^{n}$ are understood quite well geometrically, and we'll briefly go over the main result, although we won't need them in any serious way. We follow [4, Section 5.3].

To start off, we need the following construction. Take the formal group laws $f(x, y)$ and $g(x, y)$ over a ring R, corresponding to a map $\operatorname{Spec} R \to \mathcal{M}_{\mathrm{FG}} \times \mathcal{M}_{\mathrm{FG}}$, and write $\mathrm{Iso}_R(f(x, y), g(x, y))$ for the 2-fibre product of this map along the diagonal $\mathcal{M}_{\mathrm{FG}} \to \mathcal{M}_{\mathrm{FG}} \times \mathcal{M}_{\mathrm{FG}}$, omitting the reference to R if context is clear. As $\mathcal{M}_{\mathrm{FG}}$ is an Adams stack, the diagonal is affine, hence $\mathrm{Iso}(f(x, y), g(x, y))$ is an affine scheme. For any affine scheme $\operatorname{Spec} A$, the A-valued points of $\mathrm{Iso}(f(x, y), g(x, y))$ are in correspondence with a map $\pi\colon \operatorname{Spec} A \to \operatorname{Spec} R$, along with an isomorphism of formal group laws $\varphi\colon \pi^* f(x, y) \to \pi^* g(x, y)$. The following result is [4, Thm. 5.23] and may be regarded as a geometric interpretation of Theorem A.3.12.

Theorem 3.2.5. Let $f(x, y)$ and $g(x, y)$ be two formal group laws of height n over R. Then the map $\operatorname{Iso}(f(x, y), g(x, y)) \to \operatorname{Spec} R$ is surjective and pro-étale. □

The scheme $\operatorname{Iso}(f(x, y), g(x, y))$ admits an action by $\operatorname{Iso}(f(x, y), f(x, y))$ (and, by symmetry, also by $\operatorname{Iso}(g(x, y), g(x, y))$ but this isn't going to give us anything new), and this turns $\operatorname{Iso}(f(x, y), g(x, y))$ into an $\operatorname{Iso}(f(x, y), f(x, y))$-torsor. We may now state the following result, found as [4, Thm. 5.30], which, as promised, describes the geometry of $\mathcal{M}^n_{\mathrm{FG}}$.

Theorem 3.2.6. The stack $\mathcal{M}^n_{\mathrm{FG}}$ has a single geometric point, represented by any height-n formal group law $f(x, y)$ over $\mathbb{F}_p$. Fixing $f(x, y)$, the map $\mathcal{M}^n_{\mathrm{FG}} \to B\operatorname{Iso}(f(x, y), f(x, y))$ sending a formal group law $g(x, y)$ over R to the torsor $\operatorname{Iso}(f(x, y), g(x, y))$ is an equivalence of algebraic stacks. □

3.3 The Landweber exact functor theorem

Recall from Section 1.3 that every complex orientation on a ring spectrum E corresponds to a map of ring spectra $\mathrm{MU} \to E$, which gives rise to a map of coefficient rings $\mathrm{MU}_* \to E_*$ that by Theorem 1.3.9 in turn yields a formal group law over E_*. We now ask the following converse question. Given some graded ring E_*, and a graded formal group law $f(x, y)$ over E_* classified by a graded ring map $\mathrm{MU}_* \to E_*$, is there a spectrum E with $\pi_*(E) = E_*$, endowed with a complex orientation $\mathrm{MU} \to E$ that gives rise to the formal group $f(x, y)$? This is what we turn to in this section. Much of the contents in this section can be found back in [5, Lectures 15–17] and [7, Sections 4.1 and 4.7]. As the word 'Spec' will be used, we'll need to assume that E_* is an evenly graded ring so that it becomes commutative rather than just anti-commutative.

Here's what will turn out to be the winning recipe. Given a graded ring E_*, and a formal group law classified by a graded map $L \cong \mathrm{MU}_* \to E_*$, we take the functor $X \mapsto \mathrm{MU}_*(X) \otimes_L E_*$ from spectra to graded abelian groups. In decent cases, this will turn out to be a homology theory, represented by some spectrum E, which has strong potential to satisfy the desired properties. We turn to the question when our functor defines a homology theory. To check that the functor defines a homology theory, we should verify the Eilenberg–Steenrod axioms. The only non-trivial part is the following. Given a cofibre sequence of spaces $A \to X \to C$, will we end up with an exact sequence upon applying $\mathrm{MU}_*(\,\cdot\,) \otimes_L E_*$?

If E_* is flat over L, then the answer is obviously yes. It is often said, however, that there are not many interesting flat modules over L. To illustrate this, suppose R is a flat L-module, where we forget about gradings for a moment. Let a be an element of L, and consider the map $L \to L$ sending x to ax. This is an L-module homomorphism, and it is injective because L doesn't have zero-divisors. Now tensor this with the identity map on R to find the map $\mu_a \colon R \to R$ sending $r \in R$ to ar. If R is a flat L-module, this should still be an injection. This indicates that M must be 'very large', so to speak, to ensure flatness over L. To make this even more precise, if R is actually an L-algebra, we can apply the various μ_a to $1 \in R$, which shows that $L \to R$ must be an injection.

As it turns out, the functor is already a homology theory under much weaker circumstances. Evidence of this is already given by the fact that we do not need to consider arbitrary modules over L, but rather only those of the form $\mathrm{MU}_*(X)$. In fact, we find the following result.

Theorem 3.3.1 (Landweber Exact Functor Theorem). Consider a graded formal group law over a graded ring E_*, corresponding to a ring map $L \to E_*$. If the corresponding map $\operatorname{Spec} E_* \to \mathcal{M}_{\mathrm{FG}}$ is a flat morphism of algebraic stacks, then the functor $X \mapsto \mathrm{MU}_*(X) \otimes_L E_*$ is a homology theory. □

We remark right away that there's also a 2-periodic version of this theorem, which we talk about above Example 3.3.5. Before proving this result, it is worth spending some time on why it answers to our prayers. First, we mention that this indeed a weaker condition on E_* than the previous one. Indeed, if $\operatorname{Spec} E_* \to \operatorname{Spec} L$ is a flat morphism, then so is the resulting map $\operatorname{Spec} E_* \to \mathcal{M}_{\mathrm{FG}}$, as the canonical map $\operatorname{Spec} L \to \mathcal{M}_{\mathrm{FG}}$ is flat as well. We saw this in Section 2.5.

The homology theory $X \mapsto \mathrm{MU}_*(X) \otimes_L E_*$ is representable by some unique spectrum which we denote by E. Moreover, thanks to Theorem 1.1.3, any map $E_* \to E'_*$ lifts to a map of spectra. We would like to then conclude that the association sending E_* to E is functorial. But we need to be careful: due to the existence of phantom maps, the lifts may perhaps not be unique. As it turns out, however, in certain special cases, the existence of phantom maps is ruled out. In [5, Prop. 10 of Lecture 17] we find that there are no non-trivial phantom maps between evenly graded spectra. This applies to us because E is evenly graded. Indeed, notice that

$$\pi_*^{\mathrm{S}}(E) = [\mathbb{S}, \Sigma^{-*}\mathbb{S} \wedge E] = E_0(\Sigma^{-*}\mathbb{S}) = E_*(\mathbb{S}) = \mathrm{MU}_*(\mathbb{S}) \otimes_L E_* \,.$$

By Theorem 1.2.7, MU is evenly graded hence the above equalities tell us that E is, too.

There are multiplicative structures on the homology $X \mapsto \mathrm{MU}_*(X) \otimes_L E_*$, and thanks to the uniqueness of lifts discussed above, we know that this turns E into another ring spectrum; moreover, it follows that the map of spectrum $\mathrm{MU} \to E$ defined by the unique lift of the obvious natural transformation $\mathrm{MU}_*(X) \to \mathrm{MU}_*(X) \otimes_L E_*$ must be in fact be a map of ring spectra. By Theorem 1.3.8 it follows that E is complex orientable, and that the formal group law associated to E is precisely the one we started out with.

It is about time we look at the proof of Theorem 3.3.1. We start with a technical lemma that already does most of the work.

Lemma 3.3.2. Take a formal group law over a commutative ring R, yielding a morphism $f\colon \operatorname{Spec} R \to \mathcal{M}_{\mathrm{FG}}$. Let E be an R-module that is flat over $\mathcal{M}_{\mathrm{FG}}$. Then the functor sending a quasi-coherent sheaf M over $\mathcal{M}_{\mathrm{FG}}$ to $f^*M \otimes_R E$ is an exact functor $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}) \to R\text{-}\mathsf{Mod}$. □

More generally, the claim holds even if f is a morphism that does not come from a formal group law over R. Let us also mention the meaning of E being flat over $\mathcal{M}_{\mathrm{FG}}$. Recall that we defined a pushforward of quasi-coherent sheaves along any affine morphism of algebraic stacks. The map f is affine — see the first paragraph in the proof of Lemma 3.2.3. The pushforward f_* of quasi-coherent sheaves therefore makes sense, and when we say E is flat over $\mathcal{M}_{\mathrm{FG}}$, we really mean that f_*E is flat over $\mathcal{M}_{\mathrm{FG}}$, which in turn means that, for any morphism $q\colon \operatorname{Spec} A \to \mathcal{M}_{\mathrm{FG}}$, the pullback of f_*E along q is flat as an A-module.

Proof of Lemma 3.3.2: We pull our map f back along the obvious map $\operatorname{Spec} L \to \mathcal{M}_{\mathrm{FG}}$ to find a 2-pullback diagram

$$\begin{array}{ccc} \operatorname{Spec} R[b_0^{\pm 1}, b_1, b_2, \ldots] & \xrightarrow{f'} & \operatorname{Spec} L \\ {\scriptstyle p'}\downarrow & & \downarrow{\scriptstyle p} \\ \operatorname{Spec} R & \xrightarrow[f]{} & \mathcal{M}_{\mathrm{FG}} \end{array}$$

The map p is faithfully flat hence so is p'. Let us look at the composed functor $M \mapsto f^*M \otimes_R E \mapsto (p')^*(f^*M \otimes_R E)$. It turns out that it suffices to show that this composed functor is left-exact. Indeed, the claim below implies that this tells us that the first half of the composition is also left-exact, and this is enough for our purposes, as right-exactness is automatic.

Lemma 3.3.3. Take two additive functors $F\colon \mathcal{A} \to \mathcal{B}$ and $G\colon \mathcal{B} \to \mathcal{C}$ between abelian categories and assume that G is faithful and exact. Then F is left-exact if $G \circ F$ is left-exact. □

Proof of Lemma 3.3.3: To check this, left-exactness is equivalent to preservation of kernels, hence for any morphism f in $\mathcal{A}$, we have $G \circ F(\operatorname{Ker} f) = \operatorname{Ker}(G \circ F(f)) = G(\operatorname{Ker} F(f))$. Thus, the natural map $F(\operatorname{Ker} f) \to \operatorname{Ker} F(f)$ becomes an isomorphism after applying G. We need only verify now that $F(\operatorname{Ker} f) \to \operatorname{Ker} F(f)$ was already an isomorphism before applying G. This is indeed the case, due to the faithfulness and exactness assumption on G. ■

We continue our main proof. We notice that

$$\begin{aligned} (p')^*(f^*M \otimes_R E) &\cong (f \circ p')^*M \otimes_{R[b_0^{\pm 1}, b_1, \ldots]} (p')^*E && \text{pullbacks commute with } \otimes \\ &\cong (p \circ f')^*M \otimes_{R[b_0^{\pm 1}, b_1, \ldots]} (p')^*E && \text{commutativity} \\ &\cong p^*M \otimes_L f'_*(p')^*E && \text{definition of pullback} \\ &\cong p^*M \otimes_L p^*f_*E && \text{flat base change [11, Tag 02KH]} \end{aligned}$$

By flatness of p, the pullback functor p^*M is exact. By flatness of E as a module over $\mathcal{M}_{\mathrm{FG}}$, p^*f_*E is flat as a module over L. It follows that the functor $M \mapsto p^*M \otimes_L p^*f_*E$ is exact, which proves the result. ■

Proof of Theorem 3.3.1: We claim that the functor $X \mapsto E_*(X) = \mathrm{MU}_*(X) \otimes_{\mathrm{MU}_*} E_*$ is a homology theory. By looking at the Eilenberg–Steenrod axioms for reduced homology, it suffices to show the following. Given an inclusion $i\colon A \hookrightarrow X$ of pointed spaces, with mapping cone $X \to C(i)$, the resulting morphism $\widetilde{E}_*(A) \to \widetilde{E}_*(X) \to \widetilde{E}_*(C(i))$ is an exact sequence. We know that $\widetilde{\mathrm{MU}}_*(A) \to \widetilde{\mathrm{MU}}_*(X) \to \widetilde{\mathrm{MU}}_*(C(i))$ is an exact sequence, and in fact, it is exact in the category of graded MU-comodules. In Section 3.1 we saw how graded comodules correspond to quasi-coherent sheaves over $\mathcal{M}_{\mathrm{FG}}$, hence we have a resulting exact sequence $\mathscr{F}_A \to \mathscr{F}_X \to \mathscr{F}_{C(i)}$ in $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}})$.

At this point, we apply Lemma 3.3.2 in the special case where f is the map $\operatorname{Spec} \mathrm{MU}_* \to \mathcal{M}_{\mathrm{FG}}$, and E is the graded MU_*-module E_*, which is flat over $\mathcal{M}_{\mathrm{FG}}$ thanks to the flatness assumption on the morphism $\operatorname{Spec} E_* \to \mathcal{M}_{\mathrm{FG}}$. We find that the functor $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}) \to \mathrm{MU}_*\text{-}\mathsf{Mod}$ sending M to $f^*M \otimes_{\mathrm{MU}_*} E_*$ will be exact. Apply this exact functor to the exact sequence $\mathscr{F}_A \to \mathscr{F}_X \to \mathscr{F}_{C(i)}$, and, if you wish, compose with forgetful functors to Ab (which are exact as well) and the desired result follows. ∎

There's another well-known theorem, intimately connected with the one above, that carries the same name. Let us take a formal group law over a ring R, corresponding to a ring map $L \to R$. Let v_n be the p^n-th coefficient of the p-series of the formal group law over R. To formulate the original criterion ensuring flatness of R over $\mathcal{M}_{\mathrm{FG}}$, we recall the following notion from commutative algebra. Given a commutative ring R, and an R-module M, an M-**regular sequence** is a sequence of elements $r_1, r_2, \ldots$ in R such that r_i is a non-zero-divisor of $M/(r_1, \ldots, r_{i-1})M$ for all $i \geq 1$. We remark here that the element 0 is never a zero-divisor.

Theorem 3.3.4 (Landweber Exact Functor Theorem). A module M over the Lazard ring L is flat over $\mathcal{M}_{\mathrm{FG}}$ if and only if, for every prime p, the sequence $v_0, v_1, \ldots$ in L is an M-regular sequence. □

The proof of this statement can be found in [5, Lecture 16]. We'll see applications of this later in this section. But before we go on I'd like to briefly talk about the grading issue that we also mentioned in Section 3.1. Rather than working with $\mathcal{M}_{\mathrm{FG}}$, we could've worked with $\mathcal{M}^{\mathrm{s}}_{\mathrm{FG}}$ as well, and all results would still hold.

Also, if R_* is an evenly graded ring, then let us write $R[u^{\pm 1}]$ for the associated ungraded ring, perhaps interpreted as a graded ring concentrated in degree 0 if needed. If there's a flat map $\operatorname{Spec} R_* \to \mathcal{M}_{\mathrm{FG}}$ classifying a formal group law over R, then there's also a map $\operatorname{Spec} R[u^{\pm 1}] \to \mathcal{M}_{\mathrm{FG}}$ and surely this map will also be flat. Lemma 3.3.2 dictates we get another homology theory, namely $X \mapsto \mathrm{MU}_*(X) \otimes_L R[u^{\pm 1}]$, which is exactly $X \mapsto \mathrm{MUP}_*(X) \otimes_L R$. Thus there's also a 2-periodic version of Landweber's exact functor theorem.

Example 3.3.5. Recall from Example 1.2.5 the definition of the complex K-theory spectrum KU. We have $\pi_*^{\mathrm{S}}(\mathrm{KU}) \cong \mathbb{Z}[u^{\pm 1}]$, where u is an element of degree 2. In Example 1.3.6 we learned that KU admits a complex orientation, whose formal group law is given by $f(x, y) = x + y + uxy$. A simple calculation shows that, for any prime p, we have $v_0 = p$, $v_1 = u^{p-1}$, and $v_i = 0$ for all

$i \geq 2$. This clearly defines a regular sequence for all p, hence the map $\operatorname{Spec} \pi_*^{\mathrm{S}}(\mathrm{KU}) \to \mathcal{M}_{\mathrm{FG}}$ is flat. Theorem 3.3.1 now tells us that the functor $X \mapsto \mathrm{MU}_*(X) \otimes_L \pi_*^{\mathrm{S}}(\mathrm{KU})$ is a homology theory.

It is reasonable to believe that the spectrum arising in the above example is related to the complex K-theory spectrum that we started with, and indeed this is the case. By Theorem 1.3.8, the complex orientation on KU yields a map of ring spectra $\mathrm{MU} \to \mathrm{KU}$; for all spaces X, this in turn gives us a morphism $\mathrm{MU}_*(X) \otimes_L \pi_*^{\mathrm{S}}(\mathrm{KU}) \to \mathrm{KU}_*(X)$. This morphism clearly induces an isomorphism on the coefficient groups, hence by Theorem 1.1.1 the morphism is an isomorphism of spectra in $\mathrm{Ho}(\mathsf{Sp})$.

Example 3.3.6. Consider the inclusion of the open substack $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ into $\mathcal{M}_{\mathrm{FG}}$. By virtue of Lemma 3.2.3 this substack admits an fpqc atlas from $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}]$. The composition of the atlas with the open inclusion is flat (since all open inclusions are flat), and hence by Theorem 3.3.1 we find a homology theory $X \mapsto \mathrm{MU}_*(X) \otimes_L \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}]$, called **Johnson–Wilson theory** or **(uncompleted) Morava E-theory**, and denoted $E(n)$. By the discussion on grading above, there are also a 2-periodic versions of these spectra, defined by the functors $X \mapsto \mathrm{MUP}_*(X) \otimes_L \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}]$. They carry the same names and notations, and the literature often does not distinguish between the two versions. We'll denote the periodic version by $E(n)^{\mathrm{P}}$.

Lemma 3.3.7. The spectrum $E(n)$ is a flat ring spectrum, and the Adams stack associated to it, in the sense of Example 2.5.1, is precisely $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$. Similarly, the Adams stack associated to the ungraded flat Hopf algebroid $\left(\pi_0^{\mathrm{S}}(E(n)^{\mathrm{P}}), \widetilde{E(n)}_0^{\mathrm{P}}(E(n)^{\mathrm{P}})\right)$ is $\mathcal{M}_{\mathrm{FG}}^{<n+1}$. □

Proof: The second part of the lemma is proved in essentially the same way as the first part, so we only consider the first part. We start out with the claim that we have a pullback diagram

$$\begin{array}{ccc} \operatorname{Spec} \widetilde{E(n)}_* E(n) & \longrightarrow & \operatorname{Spec} E(n)_* \\ \downarrow & & \downarrow \\ \operatorname{Spec} E(n)_* & \longrightarrow & \mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1} \end{array}$$

To prove our claim, we first point out that by Lemma A.3.7, we may replace the $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ by $\mathcal{M}_{\mathrm{FG}}$. By Lemma 3.1.1, the pullback is naturally isomorphic to $\operatorname{Spec} E(n)_* \otimes_L W^{\mathrm{s}} \otimes_L E(n)_*$. At this point, we remark that

$$\begin{aligned} \widetilde{E(n)}_* E(n) &= \mathrm{MU}_*(E(n)) \otimes_L E(n)_* \\ &= [\mathbb{S}, \mathrm{MU} \wedge E(n)]_* \otimes_L E(n)_* \\ &= \widetilde{E(n)}_*(\mathrm{MU}) \otimes_L E(n)_* \\ &= \mathrm{MU}_*(\mathrm{MU}) \otimes_L E(n)_* \otimes_L E(n)_* \\ &= E(n)_* \otimes_L W^{\mathrm{s}} \otimes_L E(n)_* \end{aligned}$$

which shows what we wanted. It remains to be verified that the projection maps are precisely the source and target maps $E(n)_* \to E(n)_* E(n)$ as described in Example 2.1.5. From the correspondence defined in Section 2.5 the desired result would then follow. This verification is a matter of bookkeeping, and will be omitted. ■

Under the correspondence between the 2-category of Adams stacks and the 2-category of flat Hopf algebroids (see Section 2.5 and the reference to [8] therein), the inclusion of $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$ into $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ corresponds to the obvious morphism of Hopf algebroids $(\mathrm{MU}_*, \widetilde{\mathrm{MU}}_*\mathrm{MU}) \to (E(n)_*, \widetilde{E(n)}_*E(n))$. If X is a space, the associated sheaf $\mathscr{F}_X^{\mathrm{s}}$ on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ can then be pulled back to a sheaf on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$. The comodule associated to this pullback is precisely $\widetilde{E(n)}_*(X)$. It follows that a spectrum A is $E(n)$-acyclic if and only if the restriction of $\mathscr{F}_A^{\mathrm{s}}$ to $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$ vanishes.

Suppose now that X is another spectrum, whose sheaf on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ is supported on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$. Then for any $E(n)$-acyclic spectrum A, one does not expect any non-trivial maps $A \to X$, hence we may expect X to be $E(n)$-local. Conversely, suppose that the sheaf $\mathscr{F}_X^{\mathrm{s}}$ of X over $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ is *not* supported on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$. Then one can also argue why we do not expect X to be $E(n)$-local, although it's a bit awkward to explain why at this point. Taking this for granted for now, what this tells us is that we expect $E(n)$-localization to somehow correspond to restriction to the open substack $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$, and this can indeed be made precise. We'll never need it, but the reader can find it back in [5, Thm. 1 of Lecture 22].

Theorem 3.3.8. Localization at $E(n)$ is smashing. That is, if X is a spectrum, then $L_{E(n)}X \cong E(n) \wedge X$. As a consequence, the sheaf on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s}}$ associated to $L_{E(n)}X$ is the sheaf associated to the comodule $\widetilde{\mathrm{MU}}_*(\mathrm{MU} \wedge X) \otimes_{\pi_*\mathrm{MU}} E(n)_*$ over $(\pi_*^{\mathrm{s}}(\mathrm{MU}), \widetilde{\mathrm{MU}}_*\mathrm{MU})$. □

Of course, the open substack $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ admits many other fpqc coverings, and the one we chose doesn't appear to be special in any particular way. Every fpqc covering gives rise to a spectrum in the same way. While they may not be isomorphic, we do have the following comparison.

Lemma 3.3.9. Any two spectra defined by applying Landweber's exact functor theorem to an fpqc covering of $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ are Bousfield equivalent. □

Proof: For any two spectra defined this way, we claim that acyclicity is equivalent. In the discussion above we argued that a spectrum A is $E(n)$-acyclic if and only if $\mathscr{F}_A^{\mathrm{s}}$ restricts to zero on $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$. This statement relied only on the fact that the stack of $E(n)$ is $\mathcal{M}_{\mathrm{FG}}^{\mathrm{s},<n+1}$, as followed from Lemma 3.3.7. At no point in that proof did the choice of fpqc covering matter; thus, the same discussion is valid for any other spectrum obtained from a covering of $\mathcal{M}_{\mathrm{FG}}^{<n+1}$. ■

Consider now the special case where $n = 1$. We have $\pi_*^{\mathrm{S}}E(1) \cong \mathbb{Z}_{(p)}[v_1^{\pm 1}]$, which is very similar to the graded ring $\pi_*^{\mathrm{S}}(\mathrm{KU}_{(p)}) \cong \mathbb{Z}_{(p)}[u^{\pm 1}]$ that we considered in Example 1.4.11. We point out, however, that they are not isomorphic as graded rings: the element u has degree 2, whereas v_1 has degree $2p-2$. Nonetheless, there is more to be said.

Fix an isomorphism $L_{(p)} \cong \mathbb{Z}_{(p)}[t_1, t_2, \ldots]$ such that $t_{p^i-1} = v_i$. The formal group law associated to p-local complex K-theory is the map $L_{(p)} \to \mathbb{Z}_{(p)}[u^{\pm 1}]$ sending t_{p-1} to u^{p-1}, sending the other t_{p^i-1} to zero, and doing mysterious things to the remaining t_i. What about the formal group law associated to $E(1)$? Looking back at Lemma 3.2.3, it must be the map $L_{(p)} \to \mathbb{Z}_{(p)}[v_1^{\pm 1}]$ sending t_{p-1} to v_1 and sending all the other t_i to zero.

It sounds like the two constructions cannot be compared. But we can reason our way out of this. At no point does the proof of Lemma 3.2.3 depend on the precise map $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}] \to \mathcal{M}_{\mathrm{FG}}^{<n+1}$, so long as the corresponding formal group law is of height $< n+1$. It is therefore entirely irrelevant what the map $L_{(p)} \to \mathbb{Z}_{(p)}[v_1, \ldots, v_n^{\pm 1}]$ does to those t_j for $j \neq p^i - 1$: it always yields an fpqc cover, and hence by Landweber's exact functor theorem, a spectrum that deserves to be called $E(1)$; moreover, by Lemma 3.3.9, no matter what choice we make for the covering, the resulting $E(1)$'s are Bousfield equivalent.

So we need to ask ourselves whether we can choose the formal group law over $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$ in such a way that the pullback along the map $\mathbb{Z}_{(p)}[v_1^{\pm 1}] \to \mathbb{Z}_{(p)}[u^{\pm 1}]$ is the formal group law $f(x, y) = x + y + uxy$ of complex K-theory. The coefficient u not being in the image of this ring map, and the answer therefore seems to be negative. But we can fix this. As it turns out, we are free to replace the formal group law of $\mathrm{KU}_{(p)}$ by an isomorphic one.

Lemma 3.3.10. Let $f(x, y)$ and $g(x, y)$ be two Landweber exact formal group laws over a ring R. Then the two spectra given rise to via Landweber's exact functor theorem are homotopy-equivalent. □

Proof: The formal group law $f(x, y)$ corresponds to a map $L_{(p)} \to R$, which we may write as $L_{(p)} \xrightarrow{\mathrm{srce}} W_{(p)} \to R$; the formal group law $g(x, y)$ corresponds to the map $L_{(p)} \xrightarrow{\mathrm{trgt}} W_{(p)} \to R$. The homology theory E_f of $f(x, y)$ is defined by $(E_f)_*(X) = \mathrm{MU}_*(X) \otimes_{L,\mathrm{srce}} W \otimes_W E_*$, while that of $g(x, y)$ is defined by $(E_g)_*(X) = \mathrm{MU}_*(X) \otimes_{L,\mathrm{trgt}} W \otimes_W E_*$. Define the map $(E_f)_*(X) \to (E_g)_*(X)$ sending a simple tensor $m \otimes w \otimes e$ to $m \otimes \mathrm{inv}(w) \otimes e$. This defines an isomorphism of groups which by Theorem 1.1.1 lifts to an isomorphism of spectra. ■

We replace the formal group law $f(x, y)$ of $\mathrm{KU}_{(p)}$ by its so-called p-typification. This is something we haven't talked about, so we'll just state, without proof, some definitions and facts. More on this can be found in [9, Section A2.1].

- Certain formal group laws deserve to be called p-typical.
- Every formal group law over a p-local ring is canonically isomorphic to a p-typical one, called its p-typification.
- The p-typical formal group laws are represented by a p-typical analogue V of the Lazard ring, which is non-canonically isomorphic to $\mathbb{Z}_{(p)}[v_1, v_2, \ldots]$. The v_i's in this notation coincide with the v_i's in the expression of the p-series *up to* $(p, v_1, \ldots, v_{i-1})$.

The reader can find more about this in [9, Section A2.1]. In any case, we replace the formal group law of $\mathbb{Z}_{(p)}[u^{\pm 1}]$ with its p-typification, and it should correspond to a certain map from the ring $V \cong \mathbb{Z}_{(p)}[v_1, v_2, \ldots]$ to $\mathbb{Z}_{(p)}[u^{\pm 1}]$ sending v_1 to u^{p-1} and sending the other v_i to 0. Having gone through this, we now know what formal group law we should pick to define $E(1)$: the map $V \to \mathbb{Z}_{(p)}[u^{\pm 1}]$ *does* factor through $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$, and the factored map corresponds to a p-typical formal group law over $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$ which satisfies the hypotheses needed for the map $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1^{\pm 1}]$ to be fpqc.

Lemma 3.3.11. We have a splitting $\mathrm{KU}_{(p)} \cong \bigvee_{i=0}^{p-2} \Sigma^{2i} E(1)$, and consequently, $\mathrm{KU}_{(p)}$ and $E(1)$ are Bousfield equivalent. □

The Bousfield equivalence can also already be seen from Lemma 3.3.9, as the formal group law of $\mathrm{KU}_{(p)}$ defines an fpqc covering of $\mathcal{M}_{\mathrm{FG}}^{<2}$ as well.

Proof: On the one hand, $(\mathrm{KU}_{(p)})_*(X) = \mathrm{MU}_*(X) \otimes_L \mathbb{Z}_{(p)}[u^{\pm 1}]$, whereas on the other hand, $E(1)_*(X) = \mathrm{MU}_*(X) \otimes_L \mathbb{Z}_{(p)}[v_1^{\pm 1}]$. By the discussion above, the L-algebra structure on $\mathbb{Z}_{(p)}[u^{\pm 1}]$ factors through the L-algebra structure on $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$, in the sense that the diagram

$$\begin{array}{ccc} L & \longrightarrow & \mathbb{Z}_{(p)}[u^{\pm 1}] \\ & \searrow & \nearrow_{v_1 \mapsto u^{p-1}} \\ & \mathbb{Z}_{(p)}[v_1^{\pm 1}] & \end{array}$$

commutes. Consequently, we may write

$$\begin{aligned} (\mathrm{KU}_{(p)})_*(X) &= \mathrm{MU}_*(X) \otimes_L \mathbb{Z}_{(p)}[u^{\pm 1}] \\ &= \mathrm{MU}_*(X) \otimes_L \mathbb{Z}_{(p)}[v_1^{\pm 1}] \otimes_{\mathbb{Z}_{(p)}[v_1^{\pm 1}]} \mathbb{Z}_{(p)}[u^{\pm 1}] \\ &= \mathrm{MU}_*(X) \otimes_L \bigoplus_{i=0}^{p-2} u^i \cdot \mathbb{Z}_{(p)}[u^{\pm 1}] \\ &= \bigoplus_{i=0}^{p-2} \mathrm{MU}_*(X) \otimes_L u^i \cdot \mathbb{Z}_{(p)}[u^{\pm 1}] \\ &= \bigoplus_{i=0}^{p-2} \mathrm{MU}_{*-2i}(X) \otimes_L \mathbb{Z}_{(p)}[u^{\pm 1}] \end{aligned}$$

As $\mathrm{MU}_{*-2i}(X)$ is just $\Sigma^{2i}\mathrm{MU}_*(X)$, it follows that the natural map of spectra $\mathrm{KU}_{(p)} \to \bigvee_{i=0}^{p-2} \Sigma^{2i} E(1)$ induces an isomorphism on homology, which by Theorem 1.1.1 is sufficient to prove that the map is an isomorphism of spectra. With this equality in place, it is easy to see that a spectrum is $\mathrm{KU}_{(p)}$-acyclic if and only if it is $E(1)$-acyclic, which proves the stated consequence. ■

We find that the $\mathrm{KU}_{(p)}$-local stable homotopy category coincides with the $E(1)$-local stable homotopy category. We will make heavy use of this in Chapter 4.

3.4 Lubin–Tate deformation theory

The algebraic stack $\mathcal{M}_{FG}^{n}$ may be viewed as a closed substack of $\mathcal{M}_{FG}^{<n}$, and consequently, we may define the formal completion of $\mathcal{M}_{FG}^{n}$ within $\mathcal{M}_{FG}^{<n}$, which we conveniently denote by $\widehat{\mathcal{M}}_{FG}^{n}$. In the language of Section A.3, this algebraic stack classifies the formal group laws that are *almost* of height n. This substack, too, has interesting geometric properties, and in fact we devote this entire section to it. Our discussion is based on [4, Ch. 7] but we aim to give some additional or alternative proofs. Throughout this section, we work localized at the prime p.

Let $f(x, y)$ be a formal group law over a field k, and let A be a local Artinian ring whose residue field isomorphic to k, say via an isomorphism $\varphi\colon A/\mathfrak{m} \to k$. We define a **deformation** of $f(x, y)$ over A to be a formal group law $f_A(x, y)$ over A such that $\varphi^*\pi^* f_A(x, y) = f(x, y)$. Here π denotes the projection map $A \to A/\mathfrak{m}$. An **isomorphism of deformations** $f_A(x, y) \to f'_A(x, y)$ is an isomorphism $h\colon f_A(x, y) \to f'_A(x, y)$ such that $\varphi^*\pi^* h(t) = t$.

Write $\mathsf{Art}(k)$ for the category whose objects are pairs (A, φ), where A is a local Artin rings, and where φ is an isomorphism of fields $A/\mathfrak{m} \to k$; and whose morphisms $(A, \varphi) \to (A', \varphi')$ are local ring homomorphisms whose induced maps $A/\mathfrak{m} \to A'/\mathfrak{m}'$ on residue fields are compatible with φ and φ'. We define a functor $\mathrm{Def}\colon \mathsf{Art}(k) \to \mathsf{Set}$ by sending an Artin ring A to the collection of isomorphism classes of deformations of $f(x, y)$ over A.

There are several alternative definitions that one can work with. For instance, one can ask for the isomorphism $\varphi\colon A/\mathfrak{m} \to k$ not to be part of the category of Artin rings, but to instead be part of the data defining a deformation. One can work with local Artinian k-algebras rather than rings, or with complete local rings rather than Artinian ones. When defining deformations, one can also ask for $\varphi^*\pi^* f_A(x, y)$ to be isomorphic to $f(x, y)$ rather than being equal to it, and perhaps require the isomorphism to be another part of the data.

In practice, these choices make little difference to the outcome, which is that, if $f(x, y)$ is of height exactly n, the deformation functor will be **pro-representable**, which means that it is a colimit of a small filtered diagram of representable functors. When working with k-algebras rather than rings, one can expect the representing pro-object to be different, but because the notion of deformation does not depend on the presence of the algebra structure, it is easy to argue what the representing pro-k-algebra will be if you know what the representing pro-Artin ring is. That said, the choices we have made were made consciously so, as we'll explain in a second. We first formally state the promised outcome.

Theorem 3.4.1 (Lubin–Tate Theorem). Let k be a perfect field of characteristic p. If $f(x, y)$ is a formal group law over k of height exactly n, where n is a finite positive integer, then the functor Def is pro-representable. In fact, the functor is represented by the complete local ring $R(k, f) = W(k)[\![t_1, \ldots, t_{n-1}]\!]$, and there exists a formal group law $f_{\mathrm{univ}}(x, y)$ over $R(k, f)$ such that the deformation corresponding to a local ring map $R(k, f) \to A$ is precisely the pullback of $f_{\text{univ}}(x, y)$ along this map. □

Our goal is to prove this theorem using Schlessinger's representation theorem in formal deformation theory. We swiftly recall the contents of this theorem. To start off with, there is the following well-known result.

Theorem 3.4.2. Let $\mathcal{C}$ be an essentially small category with finite limits. Then a functor $F\colon \mathcal{C} \to \mathsf{Set}$ is pro-representable if and only if it preserves finite limits. □

Proof: I learned the proof from [15]. The idea goes like this. Given $F\colon \mathcal{C} \to \mathsf{Set}$, one may define a category $\mathrm{el}(F)$ consisting of all elements in $F(c)$ as c ranges over $\mathcal{C}$. The functor F is tautologically the colimit over a certain functor $\mathrm{el}(F) \to \mathrm{Funct}(\mathcal{C}, \mathsf{Set})$, and the properties of $\mathcal{C}$ tell us that the diagram over which we take our colimit is essentially small and filtered. ■

This theorem tells us that, in order to prove Theorem 3.4.1, it suffices to verify that the deformation functor preserves finite limits, which in turn follows once we show that it preserves pullbacks and the terminal object. At this point we remark that our choice of category of Artin rings becomes relevant: had we taken the isomorphism $\varphi\colon A/\mathfrak{m} \to k$ of residue fields as part of the definition of a deformation rather than part of the category of Artin rings, our deformation functor would not have preserved terminal objects (unless k has no non-trivial automorphisms).

Schlessinger's representation theorem essentially says that, if $\mathcal{C}$ is the category $\mathsf{Art}(k)$, it suffices to check certain weaker conditions. The way we state the theorem isn't the strongest possible version, but it suffices for us.

Theorem 3.4.3 (Schlessinger's Representation Theorem). Let $\mathcal{C}$ be the category $\mathsf{Art}(k)$, where k is perfect and of characteristic p. Then in order for $F\colon \mathsf{Art}(k) \to \mathsf{Set}$ to be pro-representable, it suffices to verify the following hypotheses.

- F preserves terminal objects;
- whenever we have a pullback diagram in $\mathsf{Art}(k)$ of the form

$$\begin{array}{ccc} A' \times_A A'' & \longrightarrow & A'' \\ \downarrow & & \downarrow \\ A' & \longrightarrow & A \end{array}$$

  where the map $A' \to A$ is *surjective*, the natural map $F(A' \times_A A'') \to F(A') \times_{F(A)} F(A'')$ is a bijection;
- the set $F(k[\varepsilon]/(\varepsilon^2))$, when endowed with the k-vector space structure we clarify below, is finite-dimensional.

One can always choose the pro-object to be (R, φ), where R is a complete local ring R such that $R/\mathfrak{m}^n$ is in $\mathsf{Art}(k)$ for all n, and φ is a fixed isomorphism $R/\mathfrak{m} \to k$. If moreover F preserves surjections, then R is $W(k)[[v_1, \ldots, v_d]]$, where d is the dimension of $F(k[\varepsilon]/(\varepsilon^2))$ and $W(k)$ is the ring of Witt vectors of k. □

About the vector space structure. The ring $k[\varepsilon]/(\varepsilon^2)$ is a k-vector space object in the category $\mathsf{Art}(k)$. That is, it admits an addition map $k[\varepsilon]/(\varepsilon^2) \times_k k[\varepsilon]/(\varepsilon^2) \to k[\varepsilon]/(\varepsilon^2)$, defined by sending $(\varepsilon, 0)$ and $(0, \varepsilon)$ to ε, and for every $v \in k$, it admits an action map $k[\varepsilon]/(\varepsilon^2) \to k[\varepsilon]/(\varepsilon^2)$ defined by sending ε to $v\varepsilon$. As F preserves products, it preserves the vector space structure, and a vector space object in Set just so happens to be, well, a vector space. It is this vector space whose dimension we put bounds on in the statement above.

What's up with the appearence of the ring of Witt vectors? As it turns out, it has to do with the following property satisfied by $W(k)$, which holds only when k is perfect. For any complete local ring R, there exists a unique dashed arrow making the solid diagram

$$\begin{array}{ccc} W(k) & \dashrightarrow & R \\ \downarrow & & \downarrow \\ k & \longrightarrow & R/\mathfrak{m} \end{array}$$

commute. Rings with such lifting properties are also called **Cohen rings**. For all fields of characteristic p, even non-perfect ones, Cohen rings exist. Considering that this is the only point where we need the fact that k is perfect, perhaps Theorem 3.4.1 also holds for non-perfect fields, the ring of Witt vectors being replaced by a Cohen ring. The author is not aware of a reference for the facts stated above.

Proof sketch of Theorem 3.4.3: Every statement except the very last sentence is the main result of [10]. In his notation, we set Λ to be $W(k)$, where we remark that the universal property of the Witt ring always induces a $W(k)$-algebra structure on a given local Artinian ring. We thus need only show the last part. Let $F\colon \mathsf{Art}(k) \to \mathsf{Set}$ and suppose F preserves surjections. For any vector v in $F(k[\varepsilon]/(\varepsilon^2))$, we can find a lift $\overline{v}$ to $F(W(k)[[t]])$ along the surjection $W(k)[[t]] \to k[\varepsilon]/(\varepsilon^2)$ sending t to ε. The reader may object that $W(k)[[t]]$ isn't local Artinian, but it's the filtered limit of local Artinian rings (namely, $W(k)[[t]]/\mathfrak{m}^n$) but we can just extend the definition of F to complete local rings by preserving limits.

Pick a k-vector space basis $\{v_1, \ldots, v_d\}$ of $F(k[\varepsilon]/(\varepsilon^2))$. As F preserves surjections, we may lift them to elements $\{\overline{v}_1, \ldots, \overline{v}_n\}$ of $F(W(k)[[t]])$. In turn, lift these elements to a single element $\overline{v}$ of $F(W(k)[[t_1, \ldots, t_d]])$ along the surjections $W(k)[[t_1, \ldots, t_d]] \to \prod_{i=1}^{d} W(k)[[t_i]] \xrightarrow{p_j} W(k)[[t_j]]$ for $j = 1, \ldots, d$. Our claim is that, for any local Artin ring A, there's a bijection between $\operatorname{Hom}(W(k)[[t_1, \ldots, t_d]], A)$ and $F(A)$, where the bijection is defined by pushing forward $\overline{v}$ along whatever local ring map.

The proof of our claim proceeds by induction on the length of the Artin rings — a technique the reader will also find back in [10]. Note that this suffices because local Artin rings are of finite length. If A is a local Artin ring of length 1, then A is isomorphic to k, so that $F(A)$ must be terminal, i.e. a one-element set. The same is obviously true for $\mathrm{Hom}\,(W(k)[\![t_1, \ldots, t_d]\!], A)$. This settles the induction basis. Suppose now that the result is true for all local Artin rings of length smaller than n, and let A be a length-n local Artin ring. Pick a non-zero element x in A that is annihilated by $\mathfrak{m}_A$. (Such an element exists. The chain $\mathfrak{m}_A \supseteq \mathfrak{m}_A^2 \supseteq \cdots$ stabilizes as A is Artinian, and the stable ideal is 0 by Nakayama's Lemma. Let n be the largest integer for which $\mathfrak{m}_A^n \neq 0$, and take x to be a non-zero element in $\mathfrak{m}_A^n$.) As F preserves pullbacks, we have a pullback diagram

$$\begin{array}{ccc} F(A \times_{A/(x)} A) & \longrightarrow & F(A) \\ \downarrow & & \downarrow \\ F(A) & \longrightarrow & F(A/(x)) \end{array}$$

The top-left ring, $A \times_{A/(x)} A$, is isomorphic to $k[\varepsilon]/(\varepsilon^2) \times_k A$. The isomorphism is described as follows. Given $(a, a') \in A \times_{A/(x)} A$, write $a - a' = bx$. Send it to $(a_0 - b\varepsilon, a)$, where a_0 is the residue of a mod $\mathfrak{m}_A$. As its inverse, a pair $(f_0 + f_1\varepsilon, a)$ in $k[\varepsilon]/(\varepsilon^2) \times_k A$ should get sent to $(a, a + f_1 x)$. Having established the isomorphism, we may write down a pullback diagram

$$\begin{array}{ccc} F(k[\varepsilon]/(\varepsilon^2)) \times F(A) & \longrightarrow & F(A) \\ \downarrow & & \downarrow \\ F(A) & \longrightarrow & F(A/(x)) \end{array}$$

At this point, we recall that $k[\varepsilon]/(\varepsilon^2)$ was a group object in $\mathsf{Art}(k)$. The top horizontal arrow may be interpreted as a group action of the group $F(k[\varepsilon]/(\varepsilon^2))$ on the set $F(A)$, whereas the left vertical arrow is a projection. The fact that this diagram is a pullback determines what $F(A)$ should be; indeed, for every element y in $F(A/(x))$, the pre-image of y in $F(A)$, say $f^{-1}(y)$ is a pseudo-torsor over $F(k[\varepsilon]/(\varepsilon^2))$, and in fact it's a torsor due by surjectivity. (Preservation of surjectivity is critical here!) The set $f^{-1}(y)$ is therefore isomorphic to $F(k[\varepsilon]/(\varepsilon^2))$ as an $F(k[\varepsilon]/(\varepsilon^2))$-set, and as this is true for all y, $F(A)$ is just $F(A/(x)) \times F(k[\varepsilon]/(\varepsilon^2))$.

We may now repeat the entire discussion, replacing F with $\mathrm{Hom}\,(W(k)[\![t_1, \ldots, t_d]\!], \cdot\,)$. When repeating the discussion we need to invoke the universal property of $W(k)$ to ensure that surjections are preserved. We conclude that, in an analogous way, $\mathrm{Hom}\,(W(k)[\![t_1, \ldots, t_d]\!], A)$ is determined by $\mathrm{Hom}\,(W(k)[\![t_1, \ldots, t_d]\!], k[\varepsilon]/(\varepsilon^2))$ and $\mathrm{Hom}\,(W(k)[\![t_1, \ldots, t_d]\!], A/(x))$. Now invoke the induction hypothesis to conclude that $F(A) \cong \mathrm{Hom}\,(W(k)[\![t_1, \ldots, t_d]\!], A)$, as desired. ■

Having proved this result, the hardest work is already done for us, and in order to prove Theorem 3.4.1 it suffices to verify that our functor satisfies the hypotheses of Schlessinger's theorem. We'll divide up the work in several steps.

Proof of Theorem 3.4.1: **Step 1.** The preservation of the terminal object. This one is easily seen: as the isomorphism of k was chosen *not* to be part of the data of the deformations, there cannot be any non-trivial deformations over the pair (k, Id).

Step 2. We show that $F(k[\varepsilon]/(\varepsilon^2))$ is finite-dimensional. To this end, we first construct a map $H\colon \mathrm{Def}\,(k[\varepsilon]/(\varepsilon^2)) \to k^{n-1}$, which we'll show to be a well-defined bijection in the next substeps.

Step 2.1. Take a deformation class, and represent it by a particular deformation $f'(x, y)$ over $k[\varepsilon]/(\varepsilon^2)$, corresponding to a map $L_{(p)} \to k[\varepsilon]/(\varepsilon^2)$. Thanks to Corollary A.3.10 and because $f(x, y)$ is of height n, we know that the map $L_{(p)} \to k$ representing $f(x, y)$ can be chosen to correspond to a map $\mathbb{Z}_{(p)}[t_1, t_2, \ldots] \to k$ sending t_{p^k-1} to 0 for $k = 1, \ldots, n-1$. It follows that the map $L_{(p)} \to k[\varepsilon]/(\varepsilon^2)$ representing $f'(x, y)$ will send t_{p^k-1} to $c_i\varepsilon$, where c_i are elements in k. We take $(c_1, \ldots, c_{n-1})$ to be the image of H.

Step 2.2. We verify that H is well-defined. Suppose $f'(x, y)$ and $f''(x, y)$ are two isomorphic deformations of $f(x, y)$ over $k[\varepsilon]/(\varepsilon^2)$, related by an isomorphism h. The power series $h(t)$ representing the isomorphism must be of the form

$$h(t) = (1 + h_1\varepsilon)t + h_2\varepsilon t^2 + h_3\varepsilon t^3 + \cdots .$$

Write $[p]'(t)$ and $[p]''(t)$ for the p-series of our two deformations. They are related by $h \circ [p]'(t) = [p]'' \circ h(t)$. Since $f(x, y)$ has height n, Lemma A.3.1 dictates that the images of $[p]'(t)$ and $[p]''(t)$ must be of the form

$$[p]'(t) = a_2\varepsilon t^2 + a_3\varepsilon t^3 + \cdots + a_{p^n-1}\varepsilon t^{p^n-1} + v_n t^{p^n} + \cdots ,$$

and similarly for $[p]''(t)$. At this point we simply plug the expressions for $h(t)$, $[p]'(t)$ and $[p]''(t)$ into the relation $h \circ [p]'(t) = [p]'' \circ h(t)$ to find that, for $k = 1, \ldots, n-1$, the p^k-th coefficients of $[p]'(t)$ and $[p]''(t)$ coincide.

Step 2.3. The map H is a homomorphism. Given two deformations over $k[\varepsilon]/(\varepsilon^2)$, represent them by two formal group laws over $k[\varepsilon]/(\varepsilon^2)$ corresponding to two maps $L_{(p)} \to k[\varepsilon]/(\varepsilon^2)$, or equivalently, to a single map $L_{(p)} \to k[\varepsilon]/(\varepsilon^2) \times_k k[\varepsilon]/(\varepsilon^2)$. Compose this with the addition map. Essentially by definition, this just results in the addition of the images of the $t_i \in L_{(p)}$.

Step 2.4. H is surjective. This is easy to see: $L_{(p)}$ is isomorphic to a free polynomial algebra $\mathbb{Z}_{(p)}[t_1, t_2, \ldots]$, so we can choose to lift t_{p^k-1} to a map $L_{(p)} \to k[\varepsilon]/(\varepsilon^2)$ in whatever way we want.

Step 2.5. We show that the kernel of H is trivial, so that it is injective. It suffices to show that a deformation $f'(x, y)$ of $f(x, y)$ over $k[\varepsilon]/(\varepsilon^2)$ of height exactly n must be isomorphic to the trivial deformation. By Theorem 3.2.5 (whose application strictly requires the variables c_i to be zero), the map of affine schemes $\mathrm{Iso}_{k[\varepsilon]/(\varepsilon^2)}(f(x, y), f'(x, y)) \to \mathrm{Spec}\, k[\varepsilon]/(\varepsilon^2)$ is pro-étale, so that in particular it is formally smooth; this allows us to construct the dashed lift in the diagram

$$\begin{array}{ccccc}
\operatorname{Spec} k & \hookrightarrow & \operatorname{Iso}_{k[\varepsilon]/(\varepsilon^2)}\big(f(x,y), f'(x,y)\big) & \longrightarrow & \mathcal{M}_{\mathrm{FG}} \\
\downarrow & \nearrow & \downarrow & & \downarrow \Delta \\
\operatorname{Spec} k[\varepsilon]/(\varepsilon^2) & \xrightarrow[\mathrm{Id}]{} & \operatorname{Spec} k[\varepsilon]/(\varepsilon^2) & \xrightarrow[(f(x,y), f'(x,y))]{} & \mathcal{M}_{\mathrm{FG}} \times \mathcal{M}_{\mathrm{FG}}
\end{array}$$

which yields the desired result.

Step 3. We show that the pullback along a surjection of local Artinian algebras remains preserved by our deformation functor. So let's say we have a diagram

$$\begin{array}{ccc}
A' \times_A A'' & \xrightarrow[p_2]{} & A'' \\
\downarrow p_1 & & \downarrow q_2 \\
A' & \xrightarrow[q_1]{} & A
\end{array}$$

of local Artinian rings, along with fixed maps from the residue fields to k, where all maps still commute. Assume moreover that the map $q_1 \colon A' \to A$ is surjective. We construct a map $P\colon \operatorname{Def}(A' \times_A A'') \to \operatorname{Def}(A') \times_{\operatorname{Def}(A)} \operatorname{Def}(A'')$, which we'll verify to be a bijection in a moment.

Step 3.1. The definition of P is simple. A class in $\operatorname{Def}(A' \times_A A'')$ can always be represented by a formal group law $f(x,y)$ over $A' \times_A A''$, which we can then pull back to formal group laws $p_1^* f(x,y)$ and $p_2^* f(x,y)$ over A' and A'', and then further to a formal group law over A. Consider the classes of these formal group laws, and let this be the image of our starting deformation under P.

Step 3.2. The fact that P is well-defined is easy to see. Indeed, if we take two different representatives of the same class in $\operatorname{Def}(A' \times_A A'')$, then the isomorphism between them just gets pulled back between isomorphisms in A', A and A''.

At this point, it seems reasonable to construct an inverse of P as follows. Start with deformations over A', A, and A'', and represent them by formal group laws, which should correspond to ring maps from $L_{(p)}$. The universal property of the pullback get us a map $L_{(p)} \to A' \times_A A''$, whose class is the desired deformation over $A' \times_A A''$. This doesn't work, because the maps from $L_{(p)}$ to A', A and A'' only commute *up to isomorphism*. We'll have to be more careful.

Step 3.3. We first show that P is a surjection. Start with an element $([\gamma'], [\gamma''])$ in $\operatorname{Def}(A') \times_{\operatorname{Def}(A)} \operatorname{Def}(A'')$. Represent the class $[\gamma'']$ by a formal group law $\gamma'' \colon L_{(p)} \to A''$ over A''. Represent $[\gamma]$ by the pullback of γ'' along q_2, or eqivalently by the composition $L_{(p)} \xrightarrow{\gamma''} A'' \xrightarrow{q_2} A$. Finally, represent $[\gamma']$ by choosing a lift of the map $L_{(p)} \to A$ along the surjection q_1 — this can be done because $L_{(p)}$ is a free polynomial ring. Essentially by construction we can now define a map $L_{(p)} \to A' \times_A A''$, and the class of this formal group law is the desired deformation.

Step 3.4. We show that P is also an injection. Suppose we have two formal group laws $f_1(x,y)$ and $f_2(x,y)$ over $A' \times_A A''$ defining isomorphic deformations when pulled back to A', A and A''.

Represent the isomorphism by maps $W_{(p)} \to A''$, $W_{(p)} \to A$, and $W_{(p)} \to A'$. It may seem as though the maps need not be compatible with q_1 and q_2, even though the sources and targets are. But this turns out not to be the case, as the following claim shows.

Lemma 3.4.4. Let $f(x, y)$ be a formal group law of height n, and let $f_A(x, y)$ and $f'_A(x, y)$ be two deformations of $f(x, y)$ to a local Artin ring A with residue field isomorphic to k. There exists at most one isomorphism $h \colon f_A(x, y) \to f'_A(x, y)$ defining an isomorphism of deformations. □

Proof of Lemma 3.4.4: We have a pullback diagram

$$\begin{array}{ccccc} \mathrm{Iso}_k\left(f(x,y), f(x,y)\right) & \longrightarrow & \mathrm{Iso}_A\left(f_A(x,y), f'_A(x,y)\right) & \longrightarrow & \mathcal{M}_{FG} \\ \downarrow & & \downarrow & & \downarrow{\scriptstyle \Delta} \\ \mathrm{Spec}\, k & \longrightarrow & \mathrm{Spec}\, A & \longrightarrow & \mathcal{M}_{FG} \times \mathcal{M}_{FG} \end{array}$$

By Theorem 3.2.5, the left-most vertical map is pro-étale, and hence in particular formally unramified. As A is a local Artinian ring, the map $\mathrm{Spec}\, k \to \mathrm{Spec}\, A$ represents its only point. As the property of a morphism being formally unramified can be checked fibrewise, it follows that $\mathrm{Iso}_A\left(f_A(x, y), f'_A(x, y)\right) \to \mathrm{Spec}\, A$ is formally unramified as well. The ring A being local Artinian, the maximal ideal of A will be nilpotent so that $\mathrm{Spec}\, k \to \mathrm{Spec}\, A$ qualifies as an infinitesimal thickening, and therefore there can exist at most one dashed arrow in the diagram

$$\begin{array}{ccc} \mathrm{Spec}\, k & \longrightarrow & \mathrm{Iso}_A\left(f_A(x,y), f'_A(x,y)\right) \\ \downarrow & \nearrow & \downarrow \\ \mathrm{Spec}\, A & \xrightarrow[\mathrm{Id}]{} & \mathrm{Spec}\, A \end{array}$$

This proves the claim. ■

Step 4. At this point we know that the deformation functor is pro-representable by some complete local ring $R(k, f)$. If we check that the deformation functor preserves surjections, then we know what the complete local ring will be. This step is easy. Given a surjection $A' \to A$ and a deformation over A corresponding to a map $L_{(p)} \to A$, one can use the fact that $L_{(p)}$ is free to choose an appropriate lift to a deformation over A', thus proving the result.

Step 5. We explain how to see that the deformation functor is described by a universal deformation over $R(k, f) = W(k)[\![v_1, \ldots, v_{n-1}]\!]$. Consider the map $L_{(p)} \to k$ classifying $f(x, y)$. Choose an isomorphism $L_{(p)} \cong \mathbb{Z}_{(p)}[t_1, t_2, \ldots]$ such that the image of t_{p^k-1} is precisely $v_k(f)$, for all k. Define the formal group law $f_{\mathrm{univ}}(x, y)$ over $R(k, f)$ to be any lift of $\mathbb{Z}_{(p)}[t_1, t_2, \ldots] \to k$ to $W(k)[\![v_1, \ldots, v_{n-1}]\!]$ along the canonical map $W(k)[\![v_1, \ldots, v_{n-1}]\!] \to k$, so long as this lift sends t_{p^k-1} to the generator $v_i \in W(k)[\![v_1, \ldots, v_{n-1}]\!]$, for $k = 1, \ldots, n - 1$. For any local ring map $W(k)[\![v_1, \ldots, v_{n-1}]\!] \to A$, pulling back $f_{\mathrm{univ}}(x, y)$ to A yields a deformation over A. This yields a natural transformation from $\mathrm{Hom}\left(W(k)[\![v_1, \ldots, v_{n-1}]\!], \cdot\right)$ to Def. As both maps preserve surjections, and the transformation is an isomorphism for $A = k[\varepsilon]/(\varepsilon^2)$, the inductive argument in the proof of Theorem 3.4.3 carries over to prove that the transformation is an isomorphism for all A. ∎

Aside that we won't need: The universal formal group law over $R(k,f)$ gives rise to a map $\operatorname{Spec} R(k,f) \to \mathcal{M}^n_{\mathrm{FG}}$. This map satisfies the hypotheses of Theorem 3.3.4, so that Theorem 3.3.1 *almost* gives us a homology theory. I say "almost" because the map $g\colon L \to R(k,f)$ is not graded. There's a cheap fix for this. Consider the bigger ring $R(k,f)[u^{\pm 1}]$, where we declare u to have degree 2, and extend our map g to a graded map $L \to R(k,f)[u^{\pm 1}]$, sending a homogeneous element x of L not to $g(x)$, but to $g(x) \cdot u^{|x|}$. We may now safely construct our homology theory $X \mapsto \mathrm{MU}_*(X) \otimes_{\pi_*\mathrm{MU}} R(k,f)[u^{\pm 1}]$, denoted E_n (but also $E(n)$ sometimes, e.g. in [5, Rmk. 9 of Lecture 21]) and called **(completed) Morava E-theory**. We remark that, despite its notation, the construction implicitly depends on a choice of perfect field k and a formal group law of height n over k.

Let A be a local Artin ring with residue field k. As the elements of A not in k will be nilpotent, any local ring map $R(k,f) \to A$ inducing an isomorphism on residue fields will factor through $R(k,f)/\mathfrak{m}^N$ for some N. Consequently, the local ring maps $R(k,f) \to A$ correspond to maps $\operatorname{Spec} A \to \operatorname{Spf} R(k,f)$. From this point of view, a map $\operatorname{Spec} A \to \operatorname{Spf} R(k,f)$ makes sense even if A is not a local Artin ring. So take *any* ring A, and any map $\operatorname{Spec} A \to \operatorname{Spf} R(k,f)$, which, now by definition, corresponds to a ring map $R(k,f)/\mathfrak{m}^N \to A$. The universal formal group law on $R(k,f)$ can be pulled back to $R(k,f)/\mathfrak{m}^N$, and then further to A. The elements $v_0, \ldots, v_{n-1}$ of this formal group law over A will be nilpotent, and the element v_n will be a unit. This yields a morphism $\operatorname{Spec} A \to \widehat{\mathcal{M}}^n_{\mathrm{FG}}$, and as this is valid for all A, we find a morphism $\pi\colon \operatorname{Spf} R(k,f) \to \widehat{\mathcal{M}}^n_{\mathrm{FG}}$ of fibred categories over Aff.

Theorem 3.4.5. The morphism $\pi\colon \operatorname{Spf} R(k,f) \to \widehat{\mathcal{M}}^n_{\mathrm{FG}}$ is fpqc. □

The above result is based on [4, Prop. 7.11], but we fill in a small gap by showing that π is representable by schemes rather than by formal schemes.

Proof: We need to show that π is representable by schemes, so that we can make sense of flatness, surjectivity, and quasi-compactness of π. This will prove the result. To prove representability, pick a morphism $\operatorname{Spec} A \to \widehat{\mathcal{M}}^n_{\mathrm{FG}}$, and localize if needed to ensure that the morphisms is represented by a formal group law $g(x,y)$ over A that is almost of height n. We will attempt to study the 2-fibre product $\operatorname{Spf} R(k,f) \times_{\widehat{\mathcal{M}}^n_{\mathrm{FG}}} \operatorname{Spec} A$.

By definition, a morphism from $\operatorname{Spec} B$ into this product should consist of a morphism $\pi_2\colon \operatorname{Spec} B \to \operatorname{Spec} A$, a morphism $\pi_1\colon \operatorname{Spec} B \to \operatorname{Spf} R(k,f)$, and finally an isomorphism between $\pi_1^* f_{\mathrm{univ}}(x,y)$ and $\pi_2^* g(x,y)$. By Lemma A.3.8, $\pi_1^* f_{\mathrm{univ}}$ is almost of height n, and in fact there exists some N such that the coefficients $v_0, \ldots, v_{n-1}$ of $\pi_1^* f_{\mathrm{univ}}$ satisfy $v_0^N = v_1^N = \cdots = v_{n-1}^N$. In turn, this tells us that there exists a yet larger N' such that the map π_1 will always factor through $\operatorname{Spec} R(k,f)/\mathfrak{m}^{N'}$. In fact, this integer N' does not depend on our initial choice of B, but only on our choice of morphism $\operatorname{Spec} A \to \widehat{\mathcal{M}}^n_{\mathrm{FG}}$. It follows that we may as well consider the fibre product $\operatorname{Spec} R(k,f)/\mathfrak{m}^{N'} \times_{\widehat{\mathcal{M}}^n_{\mathrm{FG}}} \operatorname{Spec} A$. By the first part of Lemma A.3.8, this is in fact isomorphic to $\operatorname{Spec} R(k,f)/\mathfrak{m}^{N'} \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} A$. But by Lemma 3.1.1 this is just the affine scheme $\operatorname{Spec}(R(k,f)/\mathfrak{m}^{N'} \otimes_L W \otimes_L A)$.

We now show that π is flat. We keep the notation as above. Consider the diagram

$$\begin{array}{ccccc}
\operatorname{Spf} R(k,f) \times_{\widehat{\mathcal{M}}^n_{\mathrm{FG}}} \operatorname{Spec} A & \xrightarrow{(*)} & \operatorname{Spec} R(k,f) \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} A & \longrightarrow & \operatorname{Spec} A \\
\downarrow & & \downarrow & & \downarrow \\
\operatorname{Spec} R(k,f)/\mathfrak{m}^{N'} & \longrightarrow & \operatorname{Spec} R(k,f) & \longrightarrow & \mathcal{M}_{\mathrm{FG}}
\end{array}$$

Note that we are talking about the spectrum of $R(k,f)$ rather than the formal spectrum. By our discussion above, the product in the top left corner is in fact the 2-fibre product of the outer diagram. We claim that the map $(*)$ is an isomorphism. This would in fact prove the result; indeed, the map $\operatorname{Spec} R(k,f) \to \mathcal{M}_{\mathrm{FG}}$ is flat by Theorem 3.3.4 hence so is the lifted map. To prove the claim, consider any map $\operatorname{Spec} B \to \operatorname{Spec} R(k,f) \times_{\mathcal{M}_{\mathrm{FG}}} \operatorname{Spec} A$. By definition, this should be a map $\operatorname{Spec} B \to \operatorname{Spec} R(k,f)$, a map $\operatorname{Spec} B \to \operatorname{Spec} A$, and an isomorphism between the two resulting formal group laws. Invoking Lemma A.3.8 again, it follows that the map $\operatorname{Spec} B \to \operatorname{Spec} R(k,f)$ must factor through $\operatorname{Spec} R(k,f)/\mathfrak{m}^{N'}$.

We now check that π is quasi-compact and surjective. Quasi-compactness is immediate as it is affine. As for surjectivity, notice that surjectivity may be checked stalk-locally, so that we need only check that the map π' in the 2-pullback diagram

$$\begin{array}{ccc}
\operatorname{Spec} K \times_{\widehat{\mathcal{M}}^n_{\mathrm{FG}}} \operatorname{Spf} R(k,f) & \xrightarrow{\pi'} & \operatorname{Spec} K \\
\downarrow & & \downarrow q \\
\operatorname{Spf} R(k,f) & \xrightarrow[\pi]{} & \widehat{\mathcal{M}}^n_{\mathrm{FG}}
\end{array}$$

is surjective, and in fact, after localizing if needed, we may assume that the morphism q represents a formal group law over K. Now, as K is a field, it has no zero-divisors, so any formal group law almost of height n, is in fact of exactly height n. The morphism q thus factors through $\mathcal{M}^n_{\mathrm{FG}}$. Now, we assert that we have a pullback diagram

$$\begin{array}{ccc}
\operatorname{Spec} k & \longrightarrow & \mathcal{M}^n_{\mathrm{FG}} \\
\downarrow & & \downarrow \\
\operatorname{Spf} R(k,f) & \xrightarrow[\pi]{} & \widehat{\mathcal{M}}^n_{\mathrm{FG}}
\end{array}$$

Before we verify this assertion, we explain how this will prove the theorem. In Theorem 3.2.6 we learned that $\mathcal{M}_{\mathrm{FG}}^n$ has a single geometric point so that the top horizontal morphism must be surjective. Thus, if we base-change further to the map π', we should still have surjectivity.

To prove our assertion, simply notice that any morphism $\operatorname{Spec} B \to \operatorname{Spf} R(k, f) \times_{\widehat{\mathcal{M}}_{\mathrm{FG}}^n} \mathcal{M}_{\mathrm{FG}}^n$ should consist of a map $\pi_1 \colon \operatorname{Spec} B \to \operatorname{Spf} R(k, f)$ and a formal group law over B of height exactly n along with an isomorphism from $\pi_1^* f_{\mathrm{univ}}$ to this formal group law. The pullback $\pi_1^* f_{\mathrm{univ}}$ must then also be of height exactly n. It follows that B will factor through $\operatorname{Spec} R(k, f)/\mathfrak{m}$, which is precisely $\operatorname{Spec} k$. This shows the desired result. ■

In the next chapter, we will be interested in sheaves over $\widehat{\mathcal{M}}_{\mathrm{FG}}^n$. Thanks to fpqc descent of quasi-coherent sheaves (Example 2.4.1) the existence of the fpqc morphism π will help us in understanding them better.

Chapter 4

The $\mathrm{KU}_{(p)}$-local stable homotopy category

Having introduced the necessary language from homotopy theory and algebraic geometry, and having investigated the structure of the moduli stack of formal groups, we finally reap the fruits of our work by applying it to the study of the $\mathrm{KU}_{(p)}$-local stable homotopy category. This category was first systematically studied by Bousfield in the 1970s. We give an overview of some of his results in the first section. After that we will see how we can use the work of the previous chapters to approach Bousfield's results from a new perspective.

4.1 Overview of Bousfield's results

In 1979, Alridge Bousfield published a paper, [2], in which he investigated the algebraic structure of the $\mathrm{KU}_{(p)}$-local stable homotopy category, which we briefly came across in Example 1.4.11. The goal of this section is to briefly summarize some of the results of his paper. As we proved in Lemma 3.3.11, $E(1)$ and $\mathrm{KU}_{(p)}$ are Bousfield equivalent so that their localized homotopy categories coincide. Bousfield was already aware of this when he wrote his paper, and he chose to consider both points of view. We will only summarize the parts involving $E(1)$, as this is the perspective we can revisit in the next section.

The 'starting point' of Bousfield's paper is the following well-known result, which is discussed in Section 8, and about which more can be found in [9, Section 2.2] and [13, Ch. 19].

Theorem 4.1.1 ($E(1)$-Adams Spectral Sequence). For any two spectra X and Y, there exists a strongly convergent spectral sequence

$$E_2^{s,t} \cong \mathrm{Ext}^{s,t}_{E(1)\text{-}\mathsf{Comod}}\left(E(1)_*(X), E(1)_*(Y)\right) \Rightarrow [L_{E(1)}X, L_{E(1)}Y]_* \,.$$

Here $\mathrm{Ext}^{s,t}$ is the t-th graded part of the s-th graded Ext-group, and $E(1)$–Comod is the category of *graded*(!) comodules over the Hopf algebroid $(\pi_* E(1), \widetilde{E(1)}_* E(1))$. □

The above theorem tells us that, on a vague level, we understand the $E(1)$-local homotopy category (and hence the $\mathrm{KU}_{(p)}$-local homotopy category) as soon as we understand the graded Ext groups $\mathrm{Ext}^{s,t}_{E(1)\text{–Comod}}(E(1)_*(X), E(1)_*(Y))$ of all $E(1)$-local spectra X and Y.

As it turns out, *if the prime p is odd*, then the second page of the spectral sequence has a very simple structure. In Section 7 of his paper, Bousfield shows that, for odd primes p, the category of $E(1)$-comodules has homological dimension at most 2. To show this, Bousfield chooses to work with a different category, which he denotes by $\mathcal{B}(p)_*$, and which he later, in Section 10, proves to be equivalent to the category of comodules. The construction of $\mathcal{B}(p)_*$, first given in Section 3 and revisited in Section 5, is rather complicated, but their definitions are motivated by the formal properties of the stable Adams operations — something which Bousfield elaborates on in Sections 2 and 4.

Let's expand on the construction of the category $\mathcal{B}(p)_*$. Write Γ^n for the quotient of the unit group $(\mathbb{Z}/p^{n+1})^*$ by its (unique) subgroup of order $p-1$. Note that there's a canonical map $\mathbb{Z}^*_{(p)} \to (\mathbb{Z}/p^{n+1})^*$, which we may compose with the quotient map to Γ^n. This may be different from the canonical map $\mathbb{Z}^*_{(p)} \to (\mathbb{Z}/p^n)^*$.

Let M be a module over the group ring $\mathbb{Z}_{(p)}[\mathbb{Z}^*_{(p)}]$. We may view it as a $\mathbb{Z}_{(p)}$-module along with operations $\psi^k \colon M \to M$ for all $k \in \mathbb{Z}^*_{(p)}$. We define $\mathcal{B}(p)_f$ to be the full subcategory of $\mathbb{Z}_{(p)}[\mathbb{Z}^*_{(p)}]$-Mod consisting of those M for which the following holds.

- Viewed as a module over $\mathbb{Z}_{(p)}$, M is finitely generated;
- for all $m \geq 1$, the action of $\mathbb{Z}^*_{(p)}$ on $M/p^m M$ factors through the map $\mathbb{Z}^*_{(p)} \to \Gamma^N$ for some large enough N;
- the vector space $M \otimes_{\mathbb{Z}^*_{(p)}} \mathbb{Q}$ admits an eigenspace decomposition into $\bigoplus_{j\in\mathbb{Z}} W_j$ such that, for $w \in W_j$, and $k \in \mathbb{Z}^*_{(p)}$, we have $(\psi^k \otimes \mathrm{Id})(w) = k^{j(p-1)} w$, where we invoke the $\mathbb{Z}_{(p)}$-module structure on the right-hand side.

Bousfield argues that $\mathcal{B}(p)_f$ is closed under finite sums, taking subobjects and taking quotients, so that it is in fact an abelian category. He goes on to define $\mathcal{B}(p)$ to be the full subcategory of the category of $\mathbb{Z}_{(p)}[\mathbb{Z}^*_{(p)}]$-modules consisting of those M such that, for all $x \in M$, the $\mathbb{Z}_{(p)}[\mathbb{Z}^*_{(p)}]$-submodule (x) lies in $\mathcal{B}(p)_f$. This, too, is an abelian category.

We are now ready for the definition of $\mathcal{B}(p)_*$. If M is a $\mathbb{Z}_{(p)}[\mathbb{Z}^*_{(p)}]$-module, then for all $i \in \mathbb{Z}$ we may write $T^i M$ for the $\mathbb{Z}_{(p)}[\mathbb{Z}^*_{(p)}]$-module whose underlying $\mathbb{Z}_{(p)}$-module is M, but where $\psi^k \colon T^i M \to T^i M$ is $k^i \cdot \psi^k \colon M \to M$. If M lies in $\mathcal{B}(p)$, then so does $T^i M$. We define $\mathcal{B}(p)_*$ to be the abelian category whose objects are collections $\{M_n\}_{n\in\mathbb{Z}}$ of objects M_n in $\mathcal{B}(p)$ along with isomorphisms between $T^{p-1} M_n$ and M_{n+2p-2} for all n (the isomorphisms being part of the data); the morphisms are defined as you would expect. For future reference, we attach a number to the claim that $\mathcal{B}(p)_*$ is equivalent to the comodule category.

Theorem 4.1.2. The category $\mathcal{B}(p)_*$ as constructed above is equivalent to the category of graded $E(1)$-comodules. □

We now look at how Bousfield computes the homological dimension of this category. The idea is to first deal with $\mathcal{B}(p)$, and then (quite easily) extend to $\mathcal{B}(p)_*$. Starting in Section 6, Bousfield constructs a functor $\mathcal{U}\colon \mathbb{Z}_{(p)}\text{-}\mathsf{Mod} \to \mathcal{B}(p)$. We remark that there's also a functor $\mathbb{Z}_{(p)}[v^{\pm 1}] \to \mathcal{B}(p)_*$ carrying the same notation. Now, we won't go over the construction, but it is useful to mention that is is the right adjoint of the forgetful functor in the other direction. This allows us to prove the following result, found as Prop. 7.3.

Lemma 4.1.3. For G a $\mathbb{Z}_{(p)}$-module and L an object in $\mathcal{B}(p)$, we have $\operatorname{Ext}^s_{\mathcal{B}(p)}(L, \mathcal{U}(G)) \cong \operatorname{Ext}^s_{\mathbb{Z}_{(p)}}(L, G)$. □

Proof: Bousfield relies on a characterization of injective objects to prove the result. This isn't needed: as $\mathcal{U}$ is a right adjoint, it is left-exact and preserves injectives, so we may apply $\mathcal{U}$ to an injective resolution of G to obtain an injective resolution of $\mathcal{U}(G)$. Together with the adjunction isomorphism this yields the result. ■

To compute $\operatorname{Ext}^s_{\mathcal{B}(p)}$ for more general objects in $\mathcal{B}(p)$, Bousfield constructs, for any object M in $\mathcal{B}(p)$, a short exact sequence

$$0 \longrightarrow M \longrightarrow \mathcal{U}(M) \longrightarrow \mathcal{U}(M) \longrightarrow M \otimes \mathbb{Q} \longrightarrow 0$$

where $\mathcal{U}$ is being applied to the underlying $\mathbb{Z}_{(p)}$-module structure of M. This allows him to relate $\operatorname{Ext}^s_{\mathcal{B}(p)}(L, M)$ to $\operatorname{Ext}^s_{\mathcal{B}(p)}(L, \mathcal{U}(M))$, the latter being known thanks to the above lemma, via a spectral sequence. Having done this, he concludes what he aimed to show, namely that $\operatorname{Ext}^s_{\mathcal{B}(p)}(L, M)$ vanishes for $s > 2$.

It should be remarked that the result fails for $p = 2$, although Bousfield does not point out why this is the case. The key issue is that the group Γ^n which we defined in our definition of $\mathcal{B}(p)_*$ is no longer cyclic if $p = 2$; because of this, the simplified construction of $\mathcal{B}(p)_*$ as outlined in Section 5 no longer holds.

The low homological dimension causes the Adams spectral sequence to degenerate at the third page, and it is this degeneration that allows Bousfield to introduce an invariant that is capable of distinguishing $E(1)$-local spectra, and which thus 'governs' the entire structure of the $E(1)$-local stable homotopy category. To be more precise, in Section 8 Bousfield defines an invariant k_X associated to every $E(1)$-local spectrum X such that two $E(1)$-local spectra X and Y are isomorphic (in the homotopy category) if and only if $E(1)_*(X) \cong E(1)_*(Y)$ and $k_X = k_Y$. This invariant k_X turns out to be a particular element of $\operatorname{Ext}^{2,1}_{E(1)\text{-}\mathsf{Comod}}(E(1)_*(X), E(1)_*(X))$.

The 'converse', so to speak, of the result mentioned above is also valid. By 'converse' I mean the following. In Thm. 9.1, it is shown that, for any $E(1)$-comodule M, and any object $\kappa \in \mathrm{Ext}^{2,1}_{E(1)\text{-}\mathsf{Comod}}(M, M)$, one can find an $E(1)$-local spectrum X such that $E(1)_*(X)$ is isomorphic, as a comodule, to M, and the invariant k_X is precisely κ. This essentially completes the algebraic characterization of the $E(1)$-local, and hence also the $\mathrm{KU}_{(p)}$-local stable homotopy category.

4.2 A fracture square

The goal of this section and the next one is to use the results of the previous chapters to shed new light on the constructions we come across in the last section. The primary goal is to give a geometric interpretation of Theorem 4.1.2, so that any further construction involving $\mathcal{B}(p)_*$ may be reconsidered from a geometric point of view.

We begin by outlining our plan. As Bousfield also observed, p-local complex K-theory is Bousfield equivalent to Morava E-theory $E(1)$ — something we proved in Lemma 3.3.11. Consequently, the $\mathrm{KU}_{(p)}$-local stable homotopy category is the same as the $E(1)$-local stable homotopy category, which means that we may as well look at $E(1)$ instead of $\mathrm{KU}_{(p)}$. In particular, we can turn our attention to the category of graded $E(1)$-comodules rather than the category of graded $\mathrm{KU}_{(p)}$-comodules.

In Section 2.5, we learned that comodules over a flat Hopf algebroid may equally well be viewed as quasi-coherent sheaves over the corresponding stacks, and by Lemma 3.3.7, the algebraic stack associated to $E(1)$ is $\mathcal{M}_{\mathrm{FG}}^{<2}$. Here, the comodules are ungraded, but the passage from ungraded to graded comodules will be dealt with later on. For now, let's focus on the category $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}^{<2})$. For the sake of keeping things general, let's focus on $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}^{<n+1})$, passing to the case $n = 1$ once that's needed.

We can essentially split up $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ into two parts: the closed substack $\mathcal{M}_{\mathrm{FG}}^{n}$, and the open complement $\mathcal{M}_{\mathrm{FG}}^{<n}$. Geometrically speaking, it is reasonable to believe that sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ should correspond to sheaves on $\mathcal{M}_{\mathrm{FG}}^{n}$ and $\mathcal{M}_{\mathrm{FG}}^{<n}$, along with an overlap condition. This can be made precise, which has lead to the fracture square in [4, Thm. 8.17]. Unfortunately, this result only manages to compare the *derived* categories of quasi-coherent sheaves, which is not enough for our purposes; moreover, the proof relies on a highly non-trivial result due to Greenlees and May. The main goal of this section is to prove a version of the theorem that does not pass to derived categories.

The starting point of this section is the following algebraic result, which is a special case of [11, Tag 05ER].

Lemma 4.2.1. Let R be a Noetherian commutative ring, and let $\widehat{R}$ be the completion of R at an element f. Then we have a Cartesian square of categories

$$\begin{array}{ccc}
\mathsf{QCoh}(\operatorname{Spec} R) & \longrightarrow & \mathsf{QCoh}(\operatorname{Spec} \widehat{R}) \\
\downarrow & & \downarrow \\
\mathsf{QCoh}(\operatorname{Spec} R_f) & \longrightarrow & \mathsf{QCoh}(\operatorname{Spec} \widehat{R}_f)
\end{array}$$

defined by the pullbacks of the obvious corresponding ring maps. More generally, let I be a finitely generated ideal of R, write U for the open complement of $V(I)$ in $\operatorname{Spec} R$ and $\widehat{U}$ for the open complement of $V(I)$ in $\operatorname{Spec} \widehat{R}$ (where the elements of I are now interpreted as lying in $\overline{R}$). Then we have a Cartesian square

$$\begin{array}{ccc}
\mathsf{QCoh}(\operatorname{Spec} R) & \longrightarrow & \mathsf{QCoh}(\operatorname{Spec} \widehat{R}) \\
\downarrow & & \downarrow \\
\mathsf{QCoh}(U) & \longrightarrow & \mathsf{QCoh}(\operatorname{Spec} \widehat{U})
\end{array}$$

defined in the same way. □

If the modules are coherent rather than quasi-coherent (or equivalently, since everything is Noetherian, finitely generated), then by [11, Tag 00MA], the horizontal maps are given by the completions $M \mapsto \widehat{M}$.

Example 4.2.2. Take R to be $\operatorname{Spec} \mathbb{Z}$, and let f be a prime p in $\mathbb{Z}$. Then $\widehat{R} = \mathbb{Z}_p$, while $\widehat{R}_f = \mathbb{Q}_p$. The above result then tells us that finitely generated $\mathbb{Z}$-modules correspond precisely to finite-dimensional $\mathbb{Q}$-vector spaces and finitely generated $\mathbb{Z}_p$-modules satisfying a compatibility relation.

The lemma is particularly valuable when it is interpreted geometrically. The vanishing set of f defines a closed subscheme of $\operatorname{Spec} R$. The scheme $\operatorname{Spec} R_f$ is precisely the open complement, while $\operatorname{Spec} \widehat{R}$, being, in a sense, a certain limit of the $\operatorname{Spec} R/f^n R$ for various n, is the closed subscheme along with infinitesimal information about its neighbourhood. The lemma essentially tells us that functions on $\operatorname{Spec} R$ are precisely functions on the closed subscheme and functions on the open complement, satisfying an overlap condition on the infinitesimal neighbourhood of the closed subscheme.

This point of view tells us that the result is local in nature, and should consequently generalize to schemes, and hopefully further to certain algebraic stacks. Proving this will be the main goal of this section. We expect something of the following form to be true. Let $\mathcal{M}$ be a stack satisfying a notion of being Noetherian. Write $\widehat{\mathcal{M}}$ for the completion defined at a principal ideal sheaf $\mathscr{I}$, or more generally, a finitely generated ideal sheaf, where 'finitely generated' means 'fpqc-locally finitely generated'. Then coherent sheaves on $\mathcal{M}$ are precisely coherent sheaves on $\widehat{\mathcal{M}}$ and coherent sheaves on the open complement $D(\mathscr{I})$, satisfying an overlap condition on the intersection.

At this point we remark that the envisaged generalization is not really the correct generalization of Lemma 4.2.1. As we also pointed out in Section 2.4, the scheme $\operatorname{Spec}\widehat{R}$ is not the formal completion of $\operatorname{Spec} R$ at I, but the spectrum of the completion of R at I. Rather, the formal completion is $\operatorname{Spf}\widehat{R}$. There are two ways to fix this issue. Either we find the correct stack-theoretic generalization of $\operatorname{Spec}\widehat{R}$, or we find a variant of Lemma 4.2.1 involving $\operatorname{Spf}\widehat{R}$. The first approach is unfeasible, however: as far as the author is aware, $\operatorname{Spec}\widehat{R}$ does not admit a simple expression from the functor-of-points perspective. We therefore choose to take the second approach. The next two lemmas will help us out.

Lemma 4.2.3. Let R be any commutative ring, and let $\operatorname{Spf}\widehat{R}$ be the formal completion of $\operatorname{Spec} R$ at the closed subscheme $\operatorname{Spec} R/I$. Then we have a limit diagram of categories

$$\begin{array}{ccccccc} \mathsf{QCoh}(\operatorname{Spf}\widehat{R}) & & & & & & \\ \downarrow & \searrow & & \searrow & & & \\ \mathsf{QCoh}(\operatorname{Spec} R/I) & \longrightarrow & \mathsf{QCoh}(\operatorname{Spec} R/I^2) & \longrightarrow & \mathsf{QCoh}(\operatorname{Spec} R/I^3) & \longrightarrow & \cdots \end{array}$$

in the sense that a quasi-coherent sheaf over $\operatorname{Spf}\widehat{R}$ corresponds precisely to a collection of R/I^n-modules M_n, for $n \geq 1$, such that $M_n \cong M_{n-1} \otimes_{R/I^{n-1}} R/I^n$, and a morphism of quasi-coherent sheaves over $\operatorname{Spf}\widehat{R}$ corresponds to compatible morphisms of R/I^n-modules. □

Proof: From our definition of quasi-coherent sheaves over an algebraic stack (of which quasi-coherent sheaves over $\operatorname{Spf}\widehat{R}$ are a special case), every quasi-coherent sheaf is determined by considering the pullback along all possible maps $\operatorname{Spec} A \to \operatorname{Spf}\widehat{R}$, along with potentially non-trivial isomorphisms, which we need not worry about as $\operatorname{Spf}\widehat{R}$ is fibred in sets. By our definition of formal completion, every such map is precisely a map $\operatorname{Spec} A \to \operatorname{Spec} R/I^n$ for some large enough n. Thus our sheaf on $\operatorname{Spf}\widehat{R}$ determines and is uniquely determined by the pullbacks along the maps $\operatorname{Spec} R/I^n \to \operatorname{Spf}\widehat{R}$. ■

Lemma 4.2.4. Let R be a Noetherian ring, and let I be a finitely generated non-trivial ideal in R. Then we have an equivalence of categories $\mathsf{Coh}(\operatorname{Spec}\widehat{R}) \cong \mathsf{Coh}(\operatorname{Spf}\widehat{R})$. □

Proof: The functor $\mathsf{Coh}(\operatorname{Spec}\widehat{R}) \to \mathsf{Coh}(\operatorname{Spf}\widehat{R})$ is described as follows. For any coherent sheaf M on $\operatorname{Spec}\widehat{R}$, we write M_n for the pullback to $\operatorname{Spec} R/I^n$ along the map of rings $\widehat{R} \to \widehat{R}/I^n\widehat{R} \cong R/I^nR$. By Lemma 4.2.3, this uniquely determines a quasi-coherent sheaf on $\operatorname{Spf}\widehat{R}$, which is easily seen to be coherent. This functor is essentially surjective: given a sheaf on $\operatorname{Spf}\widehat{R}$, one can obtain a sheaf on $\operatorname{Spec}\widehat{R}$ by taking the various R/I^n-modules, viewing them as R-modules, and taking their categorical limit. In the same way one shows that the functor is full.

We verify that the functor is faithful. Take a morphism $f\colon M \to N$ between two coherent $\widehat{R}$-modules such that the induced morphisms $M/I^n M \to N/I^n N$ are zero, Any element of M gets mapped into the intersection of the various $I^n N$. The Krull intersection theorem asserts that there exists an x in I such that $(1+x)\cdot \bigcap_n I^n N = 0$. The element $1+x$ is a unit in $\widehat{R}$ (although not necessarily in R — we really need the $\widehat{R}$-module structure at this point). Consequently, we may write $f(m) = (1+x)^{-1}(1+x)f(m) = 0$ and the result follows. ■

Example 4.2.5. The above lemma fails if we try to generalize to the non-finitely generated situation. Let the notation be as in Example 4.2.2. The functor $\mathsf{QCoh}(\operatorname{Spec}\mathbb{Z}_p) \to \mathsf{QCoh}(\operatorname{Spf}\mathbb{Z}_p)$, defined in the same way as the functor in the above proof, sends the non-coherent $\mathbb{Z}_p$-module $\mathbb{Q}_p$ to 0.

The above result now tells us that we can replace $\mathsf{Coh}(\operatorname{Spec}\widehat{R})$ in Lemma 4.2.1 with $\mathsf{Coh}(\operatorname{Spf}\widehat{R})$. We'd expect that we can do something similar with $\mathsf{Coh}(\operatorname{Spf}\widehat{R_f})$. But unfortunately, this causes problems: the open complement of the closed subscheme $V(f)$ of $\operatorname{Spf}\widehat{R}$ is empty, both set-theoretically and functorially.

It may seem that all hope is lost, but there's reason to believe that our goal can still be achieved. Lemma 4.2.1 is fundamentally about certain categories of sheaves, and only secondarily about the underlying schemes that define these sheaf categories. Just because our naive generalization of the underlying spaces to the level of stacks seems to fail, it doesn't necessarily mean that the sheaf categories cannot be realized in some other way. This is what we turn to now. We simply change the way we define our localizations and completions, and find that this new approach will give us the required categories of sheaves.

Let $\mathcal{M}$ be an algebraic stack. We will usually view $\mathcal{M}$ as a fibred category over Aff, and we also write $\mathsf{Aff}/\mathcal{M}$ to explicitly refer to this category. The fpqc topology on Aff lifts to a topology on $\mathsf{Aff}/\mathcal{M}$. We define a **structure sheaf** on $\mathcal{M}$ to be a sheaf $\mathscr{O}$ of rings defined on the site $\mathcal{M}$. The standard example is given by the 'identity functor', sending an object $\operatorname{Spec} A \to \mathcal{M}$ in the fibred category $\mathcal{M}$ to the ring A. In analogy with ringed spaces, let us call any pair $(\mathcal{M}, \mathscr{O})$ consisting of an algebraic stack along with a sheaf of rings a **ringed stack**.

We define a **quasi-coherent $\mathscr{O}$-module** $\mathscr{F}$ over a ringed stack $(\mathcal{M}, \mathscr{O})$ is an $\mathscr{O}$-module presheaf satisfying certain properties. More precisely, it consists of the following data.

- For every object ξ in the fibre $\mathcal{M}(\operatorname{Spec} R)$ (viewing $\mathcal{M}$ as a fibred category over Aff), a choice of $\mathscr{O}(\xi)$-module, denoted $\mathscr{F}(\xi)$;
- for any two objects $\xi \in \mathcal{M}(\operatorname{Spec} R)$, $\xi' \in \mathcal{M}(\operatorname{Spec} R')$, and any morphism $f\colon \xi' \to \xi$ in $\mathcal{M}$ lying over the morphism $p_{\mathcal{M}}(f)\colon \operatorname{Spec} R' \to \operatorname{Spec} R$, we want a pullback map $\operatorname{res}_f\colon \mathscr{F}(\xi) \to \mathscr{F}(\xi')$.

We ask for this data to satisfy the following conditions.

- The restriction maps behave well with respect to the module structures, in the sense that if $f\colon \xi' \to \xi$ is a morphism, we want the diagram

$$\begin{array}{ccc} \mathscr{F}(\xi) \otimes \mathscr{O}(\xi) & \longrightarrow & \mathscr{F}(\xi) \\ \downarrow{\scriptstyle \mathrm{res}_f \otimes \mathscr{O}(f)} & & \downarrow{\scriptstyle \mathrm{res}_f} \\ \mathscr{F}(\xi') \otimes \mathscr{O}(\xi') & \longrightarrow & \mathscr{F}(\xi') \end{array}$$

to always commute;
- the restriction map $\mathrm{res}_f\colon \mathscr{F}(\xi) \to \mathscr{F}(\xi')$ gives rise to a morphism $\phi_f\colon \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi') \to \mathscr{F}(\xi')$, and we ask for this map to be an isomorphism of $\mathscr{O}(\xi')$-modules;
- we ask for the isomorphisms ϕ_f to satisfy a cocycle condition: if $f\colon \xi' \to \xi$ and $g\colon \xi'' \to \xi'$ are two morphisms, the diagram

$$\begin{array}{ccc} \mathscr{F}(\xi'') & \xrightarrow{p_{\mathcal{M}}(g)} & \mathscr{F}(\xi') \otimes_{\mathscr{O}(\xi')} \mathscr{O}(\xi') \\ \downarrow{\scriptstyle p_{\mathcal{M}}(f\circ g)} & & \downarrow{\scriptstyle p_{\mathcal{M}}(f)\otimes \mathrm{Id}_{\mathscr{O}(\xi')}} \\ \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi'') & \xrightarrow{\sim} & \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi') \otimes_{\mathscr{O}(\xi')} \mathscr{O}(\xi') \end{array}$$

should commute.

We say our $\mathscr{O}$-module $\mathscr{F}$ is **coherent** if moreover the various $\mathscr{F}(\xi)$ are coherent as $\mathscr{O}(\xi)$-modules. A **morphism of (quasi-)coherent $\mathscr{O}$-modules** is just a natural transformation of functors. In this way, the quasi-coherent $\mathscr{O}$-modules over $(\mathcal{M}, \mathscr{O})$ form a category that we denote by $\mathsf{QCoh}(\mathcal{M}, \mathscr{O})$.

Lemma 4.2.6. Suppose the structure sheaf $\mathscr{O}$ preserves faithfully flat ring maps. That is, if $A \to B$ is faithfully flat, then the ring map $\mathscr{O}(A) \to \mathscr{O}(B)$ is also faithfully flat. Then quasi-coherent $\mathscr{O}$-modules over $(\mathcal{M}, \mathscr{O})$ are automatically sheaves over the fpqc topology of the fibred category defining $\mathcal{M}$. □

Proof: Let $\mathscr{F}$ be a quasi-coherent $\mathscr{O}$-module over $(\mathcal{M}, \mathscr{O})$. Take an fpqc covering $\{f_i\colon \xi_i \to \xi\}$ lying over an fpqc covering $\{p_{\mathcal{M}}(f_i)\colon \operatorname{Spec} R_i \to \operatorname{Spec} R\}$. We have to show that the diagram

$$\mathscr{F}(\xi) \longrightarrow \prod_{i\in I} \mathscr{F}(\xi_i) \rightrightarrows \prod_{i,j\in I} \mathscr{F}(\xi_i \times_\xi \xi_j)$$

is an equalizer diagram. As $\mathscr{F}$ is quasi-coherent, the diagram is isomorphic to the diagram

$$\mathscr{F}(\xi) \longrightarrow \prod_{i\in I} \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi_i) \rightrightarrows \prod_{i,j\in I} \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi_i) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi_j)$$

At this point we recall from Example 2.2.1 that we may impose the requirement that fpqc coverings be finite, which we do from this point on. As tensor products commute with finite products, the above diagram can be written as

$$\mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi) \longrightarrow \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \prod_{i\in I} \mathscr{O}(\xi_i) \rightrightarrows \mathscr{F}(\xi) \otimes_{\mathscr{O}(\xi)} \prod_{i,j\in I} \mathscr{O}(\xi_i) \otimes_{\mathscr{O}(\xi)} \mathscr{O}(\xi_j)$$

The map $\mathscr{O}(\xi) \to \prod_i \mathscr{O}(\xi_i)$ is faithfully flat so that this diagram is an equalizer diagram by virtue of [11, Tag 023M]. ■

There's only one important example that we'll need. Let $\mathscr{O}$ be the 'usual' structure sheaf on an algebraic stack. We want to define a completed structure sheaf $\widehat{\mathscr{O}}$. If all goes right, the category $\mathsf{QCoh}(\mathcal{M}, \widehat{\mathscr{O}})$ should coincide with $\mathsf{QCoh}(\widehat{\mathscr{O}})$. If this is true, then for our purposes, $(\mathcal{M}, \widehat{\mathscr{O}})$ and $\widehat{\mathcal{M}}$ are pretty much the same. But the difference is that $\widehat{\mathcal{M}}$ has 'lost' some objects, whereas the underlying stack of $(\mathcal{M}, \widehat{\mathscr{O}})$ hasn't. When it comes to making sure that intersections don't get empty, the latter seems like a safer option.

Perhaps the reader may wonder why this is the only example we'll need. What about localization? Given a structure sheaf $\mathscr{O}$ on an algebraic stack $\mathcal{M}$, and a finitely generated ideal sheaf $\mathscr{I}$ on $\mathcal{M}$, it turns out that we can also make sense of the localized sheaf $\mathscr{O}_{\mathscr{I}}$, and indeed it holds quite generally that $\mathsf{QCoh}(\mathcal{M}, \mathscr{O}_{\mathscr{I}})$ is equivalent to $\mathsf{QCoh}(\mathcal{U})$, where $\mathcal{U}$ is the open complement of the closed substack defined by $\mathscr{I}$. As it turns out, however, the completion suffices.

Example 4.2.7. Let $\mathcal{M}$ be an algebraic stack, and denote by $\mathscr{O}$ its usual structure sheaf. Let $\mathscr{I}$ be a finitely generated ideal sheaf on $\mathcal{M}$. We'd like to define a completion $\widehat{\mathscr{O}}$. We first define a presheaf $\mathscr{O}^{\mathrm{pre}}$ as follows. After localizing if needed, an object π in $\mathcal{M}(\operatorname{Spec} A)$ gets sent under $\mathscr{I}$ to a finitely generated ideal I of A. We define $\widehat{\mathscr{O}}^{\mathrm{pre}}(\pi)$ to be the completion $\widehat{A}$ of A at the ideal I, and we define a morphism $\pi \to \pi'$ above a ring map $A' \to A$ to be sent to the completed ring map $\widehat{A'} \to \widehat{A}$.

If all the rings involved were Noetherian, then $\widehat{\mathscr{O}}^{\mathrm{pre}}$ would preserves faithful flatness, and in particular it would be a sheaf. But in general this need not hold anymore, and rarely will it happen, perhaps barring the illusive empty scheme, that $\mathcal{M}$ won't admit at least some A-valued point for a non-Noetherian ring A. So we need to sheafify in order to ensure that $\widehat{\mathscr{O}}$ is a sheaf of rings on $\mathsf{Aff}/\mathcal{M}$.

Even after sheafifying, there's no reason to believe that $\widehat{\mathscr{O}}$ preserves faithfully flat ring maps. Consequently, the hypothesis of Lemma 4.2.6 may fail, and quasi-coherent sheaves over the ringed stack $(\mathcal{M}, \widehat{\mathscr{O}})$ behave quite badly. To remedy this, we need to restrict the fibred category on which we define our sheaves. This brings us to the following definition.

Let $\mathcal{M}$ be an algebraic stack, and write $\mathsf{Aff}/\mathcal{M}$ for the fibred category over Aff defining $\mathcal{M}$. The **small fppf site**, denoted $\mathsf{Fppf}/\mathcal{M}$, will be the full subcategory of $\mathsf{Aff}/\mathcal{M}$ consisting only of the faithfully flat maps $\operatorname{Spec} A \to \mathcal{M}$, endowed with the restricted Grothendieck topology coming from $\mathsf{Aff}/\mathcal{M}$. That means we have an fppf site with the finite fpqc topology. (By 'finite', we mean 'allowing only finite coverings' — a condition we needed to impose in the proof of Lemma 4.2.6.) Every sheaf on $\mathsf{Aff}/\mathcal{M}$ restricts to a sheaf on $\mathsf{Fppf}/\mathcal{M}$, although the discussion in the example above indicates that perhaps the converse may not hold. Let us write $\mathsf{QCoh}_{\text{fppf}}(\mathcal{M}, \mathscr{O})$ instead of $\mathsf{QCoh}(\mathcal{M}, \mathscr{O})$ whenever we're dealing with quasi-coherent sheaves on the small fppf site only.

Lemma 4.2.8. Let R be a *Noetherian* ring, and let I be a finitely generated ideal of R. Interpreting $\operatorname{Spec} R$ as a stack, we may define the small fppf site $\mathsf{Fppf}/\operatorname{Spec} R$ as above. Take the usual structure sheaf $\mathscr{O}$ on $\operatorname{Spec} R$ again, and define $\widehat{\mathscr{O}}$ to be the completion of $\mathscr{O}$ at I as in Example 4.2.7. Then $\mathsf{QCoh}_{\text{fppf}}(\operatorname{Spec} R, \widehat{\mathscr{O}})$ is equivalent to $\mathsf{QCoh}(\operatorname{Spec} \widehat{R})$. □

Proof: If $\operatorname{Spec} A \to \operatorname{Spec} R$ is a faithfully flat map, then R being Noetherian implies that A is Noetherian. Consequently, all affine schemes occurring in $\mathsf{Fppf}/\operatorname{Spec} R$ will be Noetherian. For any Noetherian ring A, the completion $\widehat{M}$ of an A-module M is just $M \otimes_A \widehat{A}$; moreover, faithfully flat maps between Noetherian rings are preserved under taking completions, so that Lemma 4.2.6 applies. This makes life easier. Given an $\widehat{R}$-module M, one now easily defines an object of $\mathsf{QCoh}_{\text{fppf}}(\operatorname{Spec} R, \widehat{\mathscr{O}})$ sending an A-valued point $\operatorname{Spec} A \to \operatorname{Spec} R$ to the $\mathscr{O}(A) = \widehat{A}$-module $M \otimes_{\widehat{R}} \widehat{A}$. Conversely, every quasi-coherent sheaf is, up to a potentially non-trivial isomorphism, defined in this way. ■

With the above results in place, we can now define completions of stacks without having to mess with the underlying functor of points. In particular, we can make sense of the completion of the localization without having to worry about the underlying functor of points becoming empty. We thus have the language needed to lift Lemma 4.2.1 to the generality of algebraic stacks.

At this point, the existence of an atlas is crucial, as this allows us to descend to the level of schemes whenever we are talking about sheaves over algebraic stacks. To make this precise, we work in a rather general setting — the generality is needed later. We consider the situation of an fpqc morphism $\mathcal{X} \to \mathcal{Y}$ of algebraic stacks. Our starting point is the descent property of quasi-coherent sheaves over schemes, wish we discussed in Example 2.4.1, and goal is to extend this statement to stacks. In order to do this, we need a stack-theoretic analogue of **descent data** of sheaves. Denote by p_1 and p_2 the projection maps from the 2-fibre product $\mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$ to the components. Define descent data at $\mathcal{Y}$ to be a choice of quasi-coherent sheaf $\mathscr{F}$ on $\mathcal{X}$, and an isomorphism $\varphi\colon p_1^*\mathscr{F} \to p_2^*\mathscr{F}$, such that on three-fold 2-fibre products, the isomorphism satisfies the usual cocycle conditions. A morphism of descent data is defined to be a morphism of the quasi-coherent sheaves on $\mathcal{X}$ that are compatible with the rest of the data.

Lemma 4.2.9. Let $\mathcal{X} \to \mathcal{Y}$ be an fpqc morphism of algebraic stacks. Then the category $\mathsf{QCoh}(\mathcal{Y})$ is equivalent to the category of descent data of sheaves at $\mathcal{Y}$ as defined above. □

Proof sketch: Some details will be omitted, like the verification of certain naturality conditions. Suppose we start out with descent data of sheaves at $\mathcal{Y}$, consisting of a sheaf $\mathscr{F}$ on $\mathcal{X}$, and an isomorphism φ between $p_1^*\mathscr{F}$ and $p_2^*\mathscr{F}$ satisfying the cocycle conditions. Our goal is to reconstruct the quasi-coherent sheaf $\mathscr{G}$ on $\mathcal{Y}$. We first look at the level of objects: given a morphism $f\colon T \to \mathcal{Y}$, where T is a scheme, we wish to find $f^*\mathscr{G}$.

Write p_1 and p_2 for the projections from $\mathcal{X} \times_{\mathcal{Y}} T$ to $\mathcal{X}$ and T, respectively. By relative representability, p_2 is an fpqc morphism of schemes. The descent data at $\mathcal{Y}$ contains a quasi-coherent sheaf $\mathscr{F}$ on $\mathcal{X}$, which pulls back to a quasi-coherent sheaf $p_1^*\mathscr{F}$ on T. Moreover, the compatibility relations in the descent data at $\mathcal{Y}$ pulls back to compatibility relations for $p_1^*\mathscr{F}$. To see this, we remark that there's a canonical isomorphism between $(\mathcal{X} \times_{\mathcal{Y}} T) \times_T (\mathcal{X} \times_{\mathcal{Y}} T)$ and $(\mathcal{X} \times_{\mathcal{Y}} \mathcal{X}) \times_{\mathcal{Y}} T$, so that the isomorphism φ on $\mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$ pulls back to the fibre product of $\mathcal{X} \times_{\mathcal{Y}} T$ over T. In the same way, the cocycle conditions on triple overlaps pull back as well. Thus we have found descent data of sheaves at T, which uniquely assemble into a sheaf $f^*\mathscr{G}$ on T.

Now suppose we have an isomorphism above the fibre $\mathcal{Y}(T)$, say between $f\colon T \to \mathcal{Y}$ and $g\colon T \to \mathcal{Y}$. Our goal is to establish a natural isomorphism between $f^*\mathscr{G}$ and $g^*\mathscr{G}$. We have two fibre products, $\mathcal{X} \times_{\mathcal{Y},f} T$ and $\mathcal{X} \times_{\mathcal{Y},g} T$, and they are canonically isomorphic. Both have projection maps to $\mathcal{X}$, say $p_{1,f}$ and $p_{1,g}$. They are not the same, but are isomorphic when composed with the map to $\mathcal{Y}$. This gives us a well-defined map from the 2-fibre product $\mathcal{X} \times_{\mathcal{Y}} T$ to $\mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$. The isomorphism φ on $\mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$ pulls back to an isomorphism between the two quasi-coherent sheaves $p_{1,f}^*\mathscr{F}$ and $p_{1,g}^*\mathscr{F}$ on $\mathcal{X} \times_{\mathcal{Y}} T$. This isomorphism is compatible with the compatibility isomorphisms that are part of the descent data, and so it assembles to a well-defined isomorphism between $f^*\mathscr{F}$ and $f^*\mathscr{G}$. ∎

A few remarks about this lemma. First, the most important special case of the lemma is when $\mathcal{X}$ is a scheme X. This allows us to reduce statements about quasi-coherent sheaves on algebraic stacks to statements about quasi-coherent sheaves on schemes. We've implicitly used this special case in Section 2.5 already. Second, we point out that the proof hardly uses the definition of quasi-coherent sheaves. Really, the lemma would work for any notion that satisfies fpqc descent. In particular, suppose $\mathscr{O}$ is a sheaf of rings on $\mathsf{Aff}/\mathcal{Y}$ that preserves faithfully flat ring maps. Then Lemma 4.2.6 tells us that quasi-coherent $\mathscr{O}$-modules are fpqc sheaves, so our lemma carries over to this situation. Let's write this down for concreteness.

Lemma 4.2.10. Let $f\colon \mathcal{X} \to \mathcal{Y}$ be an fpqc morphism of algebraic stacks, and let $\mathscr{O}$ be a structure sheaf on $\mathcal{Y}$, which we assume to preserve faithfully flat ring maps. Then the category $\mathsf{QCoh}(\mathcal{Y}, \mathscr{O})$ is equivalent to the category of descent data of quasi-coherent $\mathscr{O}$-modules at $\mathcal{Y}$. In particular, giving a quasi-coherent $\mathscr{O}$-module $\mathscr{F}$ on $\mathcal{Y}$ is equivalent to giving a quasi-coherent $f^*\mathscr{O}$-module $\mathscr{G}$ on $\mathcal{X}$ along with an isomorphism $\varphi\colon p_1^*\mathscr{G} \xrightarrow{\sim} p_2^*\mathscr{G}$ satisfying the cocycle conditions. ∎

So we need an atlas when generalizing Lemma 4.2.1, but even then, we cannot hope our generalization works for algebraic stacks. Indeed, when discussing completions of the structure sheaf, we saw that we needed quite stringent finiteness conditions for our constructions to make sense. Most notably, we required all our (affine) schemes to be Noetherian. When passing to algebraic stacks, we will need a similar such constraint. We use the following definition to make this precise. Let $\mathcal{M}$ be an algebraic stack. We shall call $\mathcal{M}$ **Noetherian** if it admits an fppf (not just fpqc) covering $\operatorname{Spec} A \to \mathcal{M}$ by a Noetherian affine scheme $\operatorname{Spec} A$.

Lemma 4.2.11. Let $\mathcal{M}$ be a Noetherian algebraic stack, and let $\widehat{\mathscr{O}}$ be the completion of the structure sheaf at a finitely generated ideal sheaf $\mathscr{I}$ on $\mathcal{M}$. Then $\mathsf{Coh}(\widehat{\mathcal{M}})$ is equivalent to $\mathsf{Coh}_{\mathrm{fppf}}(\mathcal{M}, \widehat{\mathscr{O}})$. □

Proof sketch: To simplify the discussion, I will only sketch the proof when $\mathcal{M}$ is in addition an Adams stack, so as to ensure that fibre products of affine schemes over $\mathcal{M}$ remain affine. In principle, this is probably not necessary at all — it's just that we've worked with affine schemes throughout the above discussion. Had we worked in the generality of arbitrary schemes, the lemma would likely easily generalize to arbitrary Noetherian algebraic stacks.

Take an fppf atlas of our stack $\mathcal{M}$ by $\operatorname{Spec} A$, where A is a Noetherian ring. As $\mathcal{M}$ is assumed to be an Adams stack, this atlas is moreover affine, so that the pullback $\operatorname{Spec} A \times_{\mathcal{M}} \operatorname{Spec} A$ will be a Noetherian affine scheme, say $\operatorname{Spec} \Gamma$. The pullback of $\mathscr{I}$ to $\operatorname{Spec} A$ is a finitely generated ideal sheaf, at which we may complete A. The formal completion $\widehat{\mathcal{M}}$ admits a map from $\operatorname{Spf} \widehat{A}$. This map defines a relatively representable fppf covering of $\widehat{\mathcal{M}}$, whose pullback along itself will be $\operatorname{Spf} \widehat{\Gamma}$. Now apply Lemma 4.2.9 to the case $\mathcal{X} = \operatorname{Spf} A$ and $\mathcal{Y} = \widehat{\mathcal{M}}$ to find that sheaves on $\widehat{\mathcal{M}}$ are sheaves on $\operatorname{Spf} A$ with a compatibility relation on $\operatorname{Spf} \Gamma$.

Now look at $(\mathcal{M}, \widehat{\mathscr{O}})$. The pullback of $\widehat{\mathscr{O}}$ along the maps $\operatorname{Spec} A \to \mathcal{M}$ and $\operatorname{Spec} \Gamma \to \mathcal{M}$ give completions of the structure sheaves of A and Γ, respectively. As we restricted our attention to the small fppf site of $\mathcal{M}$, and $\mathcal{M}$ is moreover Noetherian, $\widehat{\mathscr{O}}$ preserves faithful flatness so that Lemma 4.2.10 applies: a quasi-coherent $\widehat{\mathscr{O}}$-module on $\mathcal{M}$ is an $\widehat{\mathscr{O}}$-module on $\operatorname{Spec} A$ with a compatibility relation on $\operatorname{Spec} \Gamma$. Now apply Lemma 4.2.8 and Lemma 4.2.4, the latter invoking our assumption that our sheaves be coherent, to conclude that quasi-coherent $\widehat{\mathscr{O}}$-modules are modules on the small fppf site of $\operatorname{Spf} A$ with a compatibility relation on $\operatorname{Spf} \Gamma$.

We now see that the descriptions almost coincide, the only difference being that we work with the small fppf site when dealing with $(\mathcal{M}, \widehat{\mathscr{O}})$. After passing to the atlas, however, this distinction no longer makes a difference: any quasi-coherent sheaf defined on the small fppf site of a Noetherian (formal) scheme uniquely specifies a quasi-coherent sheaf on the big site. ■

Theorem 4.2.12. Let $\mathcal{M}$ be a Noetherian algebraic stack, and let $\mathscr{I}$ be a finitely generated ideal sheaf on $\mathcal{M}$. Write $\mathcal{U}$ for the open complement of the closed substack define by $\mathscr{I}$. Then we have a diagram

$$
\begin{array}{ccc}
\mathsf{QCoh}(\mathcal{M}) & \longrightarrow & \mathsf{QCoh}_{\mathrm{fppf}}(\mathcal{M}, \widehat{\mathscr{O}}) \\
\downarrow & & \downarrow \\
\mathsf{QCoh}(\mathcal{U}) & \longrightarrow & \mathsf{QCoh}_{\mathrm{fppf}}(\mathcal{U}, \widehat{\mathscr{O}}|_{\mathcal{U}})
\end{array}
$$

that defines a Cartesian square of categories. □

If, moreover, we restrict to coherent sheaves, we may replace $\mathsf{Coh}_{\mathrm{fppf}}(\mathcal{M}, \widehat{\mathscr{O}})$ by $\mathsf{Coh}(\widehat{\mathcal{M}})$ thanks to Lemma 4.2.11, so that under this restriction, we can truly say that sheaves on $\mathcal{M}$ are constructed from sheaves on $\mathcal{U}$ and on $\widehat{\mathcal{M}}$.

Proof sketch of Theorem 4.2.12: The stack $\mathcal{M}$ admits an fppf atlas by $\operatorname{Spec} R$, where R is an affine scheme. The pullback of the ideal sheaf $\mathscr{I}$ to $\operatorname{Spec} R$ corresponds to a finitely generated ideal I of R. We have a commutative cube

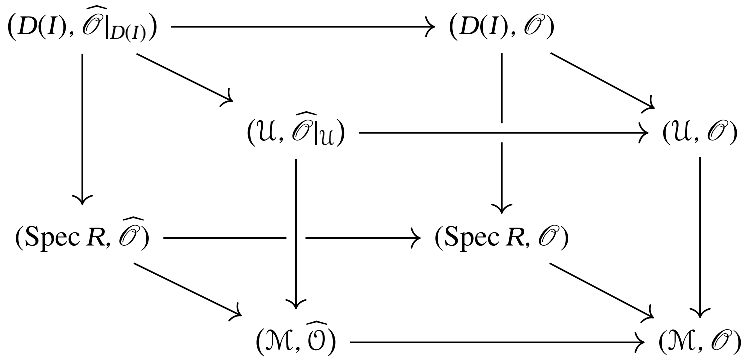

Thanks to Lemma 4.2.1, the back face yields a Cartesian square of categories of quasi-coherent sheaves. Now use descent (Lemma 4.2.9 and 4.2.10) to describe the quasi-coherent sheaves of the stacks in the front face in terms of quasi-coherent sheaves of the stacks in the back face. ■

There's a catch. While the stack $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ admits an fpqc cover by a Noetherian scheme thanks to Lemma 3.2.3, the fibre product of this cover with itself over $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ is not Noetherian anymore. Moreover, this tells us that the cover is not fppf. Consequently, $\mathsf{Fppf}/\mathcal{M}_{\mathrm{FG}}^{<n+1}$ does not contain the fpqc cover, and in fact, it may well be that it doesn't contain any object at all.

The solution to our issue lies in Lemma 3.1.2. The stack $\mathcal{M}_{\mathrm{FG}}$ may not be Noetherian, but it is the homotopy limit of Noetherian stacks, which are defined by k-buds of formal group laws, for various k. These stacks have reasonably defined fppf sites, and what's more, we have the following result. The proof is reasonably simple and can be found in [4, Thm. 3.25].

Lemma 4.2.13. The truncation maps $\mathcal{M}_{\mathrm{FG}} \to \mathcal{M}_{\mathrm{FG}}\langle k\rangle$, as k ranges over the positive integers, give rise to pullback functors $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}\langle k\rangle) \to \mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}})$, which in turn assemble into a single functor $\varinjlim_k \mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}\langle k\rangle) \to \mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}})$. This functor is faithful and induces an equivalence on the full subcategories of coherent sheaves. In particular, every coherent sheaf on $\mathcal{M}_{\mathrm{FG}}$ is the pullback of a coherent sheaf on $\mathcal{M}_{\mathrm{FG}}\langle k\rangle$ for some sufficiently large k. □

The stacks $\mathcal{M}_{\mathrm{FG}}\langle k\rangle$ fit within the framework of Theorem 4.2.12. This raises the question whether $\mathcal{M}_{\mathrm{FG}}\langle k\rangle$ admits a similar filtration into a heights. When we say that a formal group law $f(x, y)$ is of height $\geq n$ or $< n+1$, we put requirements on the coefficients $v_0(f), \ldots, v_{n-1}(f)$ of $f(x, y)$. For every n, these coefficients $v_0(f), \ldots, v_{n-1}(f)$ are always determined by finitely many coefficients of the formal group laws $f(x, y)$. Consequently, so long as k is large enough that none of the relevant coefficients get cut off, being of height $\geq n$ or $< n+1$ makes sense for k-buds. We thus find the following application of Lemma 4.2.11 and Theorem 4.2.12:

Lemma 4.2.14. Let n be a positive integer, and let k be sufficiently large so that the substacks $\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle$ and $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle$ are well-defined. Then we have a fraction square

$$\begin{array}{ccc} \mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle) & \longrightarrow & \mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle, \widehat{\mathscr{O}}) \\ \downarrow & & \downarrow \\ \mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle) & \longrightarrow & \mathsf{Coh}\left(\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle, \widehat{\mathscr{O}}|_{\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle}\right) \end{array}$$

Moreover, $\mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle, \widehat{\mathscr{O}})$ is equivalent to $\mathsf{Coh}(\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle)$. Thus, coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle$ are equivalent to coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle$ and on $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle$, along with a suitable overlap condition. ■

Combining this with Lemma 4.2.13, we arrive at the main theorem of this section.

Theorem 4.2.15. We have an equivalence of categories

$$\mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n+1}) \cong \varinjlim_{k} \left(\mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle) \times_{\mathsf{Coh}\left(\mathcal{M}_{\mathrm{FG}}^{<n+1}\langle k\rangle, \widehat{\mathscr{O}}|_{\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle}\right)} \mathsf{Coh}(\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle)\right).$$

In human terms, for every coherent sheaf $\mathscr{F}$ on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$, there is some sufficiently large k such that $\mathscr{F}$ is uniquely determined by a coherent sheaf on $\mathsf{Coh}(\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle)$ and on $\mathsf{Coh}(\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle)$ satisfying a compatibility relation. ■

This is the fracture square we've been dreaming of, albeit at the cost of having to pass to k-buds, and restricting attention of coherent sheaves. Among other things, the passage to the category of quasi-coherent sheaves will be dealt with in the next section.

4.3 The category of $\mathrm{KU}_{(p)}$-comodules revisited

Recall that our aim was to find a geometric interpretation of Theorem 4.1.2. Our outline is as follows. First, in Theorem 4.2.15 we have obtained a decomposition theorem for coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$. Our goal, however, is to describe quasi-coherent sheaves. In the first part of this section, we explain how to pass from coherent sheaves to quasi-coherent sheaves. In hindsight, we'll find that this passage corresponds precisely to the passage from $\mathcal{B}(p)_f$ to $\mathcal{B}(p)$ in Bousfield's construction.

Having done that, we will have concluded that the category of quasi-coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ can be stated entirely in terms of the category of coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n}\langle k\rangle$ and $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle$, still for k sufficiently large. This tells us that, if we understand the category of coherent sheaves on $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle$, and on $\mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle$, then, at least in principle, by induction we can understand the category of coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ for all n as well. This leads us to a study of the coherent sheaves on $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle$ and $\mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle$. To study the former we use Lubin–Tate deformation theory as developed in Section 3.4, while the latter can be tackled directly.

At the end of the journey, we should have a purely algebraic description of the category of quasi-coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$, and therefore by Section 2.5 of the category of ungraded comodules over $E(n)$. At least in principle, that is. The final step is to explain how to pass from the category of ungraded comodules to graded comodules. In Bousfield's language, this will mark the passage from $\mathcal{B}(p)$ to $\mathcal{B}(p)_*$.

Let's start with the first step, which is to express the category of quasi-coherent sheaves of $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ in terms of the category of coherent sheaves. We begin with the following result, which already suggests the road we will be taking.

Lemma 4.3.1. Let R be a Noetherian ring, and let M be a module over R. Then M is a colimit over its filtered system of coherent submodules. □

Proof: For Noetherian rings, being finitely generated is equivalent to being coherent. Therefore if $\{m_1, \ldots, m_k\}$ is a collection of finitely many elements of M, the submodule $(m_1, \ldots, m_k)$ generated by these elements are Noetherian. We can represent all these submodules as a filtered system of R-modules under inclusion. The colimit over such a system is given by the union of the R-modules involved, which must be M. ■

Notice that we need not take all finitely generated submodules of M. Any collection of submodules whose union is all of M will do just fine. Our goal is now to generalize the above result to algebraic stacks. We will use Lemma 4.2.9, applied to the case where $\mathcal{X}$ is a scheme, to reduce the proof of the generalization to the above lemma.

Theorem 4.3.2. Let $\mathcal{X}$ be an algebraic stack, and assume that it admits an fpqc covering by a Noetherian affine scheme. Then quasi-coherent sheaves on $\mathcal{X}$ are precisely filtered colimits of coherent sheaves, and in fact, every quasi-coherent sheaf is a filtered colimit over its coherent subsheaves. □

Proof: It may be useful to first consider what colimits are. As colimits commute with pullbacks, colimits of quasi-coherent sheaves (and hence of coherent sheaves) over an algebraic stack have a straightforward definition: if $\mathscr{F}_i$ are quasi-coherent sheaves over $\mathcal{X}$, then for any morphism $T \to \mathcal{X}$ from a scheme T, the pullback of $\varinjlim \mathscr{F}_i$ is just the colimit of the pullbacks of the $\mathscr{F}_i$.

The stack $\mathcal{X}$ admits an fpqc atlas from an affine scheme $\operatorname{Spec} R$, where R is a Noetherian ring. Lemma 4.2.9 now tells us that a quasi-coherent sheaf on $\mathcal{X}$ is just a quasi-coherent sheaf on $\operatorname{Spec} R$ along with an isomorphism φ of pulled back sheaves on $\operatorname{Spec} R \times_{\mathcal{X}} \operatorname{Spec} R$. So let's start with a quasi-coherent sheaf $\mathscr{F}$, and pull it back to an R-module M. By Lemma 4.3.1, M may be described as the colimit over a filtered diagram $\{M_i\}$ of finitely generated submodules. Take one such submodule M_i. The module M_i along with the restriction $\varphi|_{M_i}$ form descent data on their own, thus describing a quasi-coherent sheaf $\mathscr{F}_i$ on $\mathcal{X}$. This quasi-coherent sheaf is in fact coherent thanks to descent of finiteness properties (see [11, Tag 05AY]).

Consider the various $\mathscr{F}_i$ constructed in this way, along with their inclusion relations. We claim that their colimit $\varinjlim \mathscr{F}_i$ is $\mathscr{F}$. We know that, whatever the colimit is, it should be preserved when pulling back along the atlas, and on the atlas, the colimit is M by construction; moreover, the (trivial) isomorphism of the two pullbacks of $\varinjlim \mathscr{F}_i$ to $\operatorname{Spec} R \times_{\mathcal{X}} \operatorname{Spec} R$ is the same as that of M, so that the colimit must in fact be $\mathscr{F}$. ∎

At this point, we are able to describe quasi-coherent sheaves on $\mathcal{M}_{\mathrm{FG}}^{<n+1}$ entirely in terms of coherent sheaves on $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}\langle k\rangle$ and on $\mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle$. The next step is to describe what coherent sheaves on these two stacks should look like. As the restriction to coherent sheaves won't be relevant for this step, we will actually aim describe the quasi-coherent sheaves on both of these stacks, and we will also discuss how things work for the non-truncated stacks $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}$ and on $\mathcal{M}_{\mathrm{FG}}^{<1}$ — as it turns out, the truncation will make little difference.

We begin with $\mathcal{M}_{\mathrm{FG}}^{<1}$, and look at the truncated version in a moment. We proved in Lemma 3.2.4 that $\mathcal{M}_{\mathrm{FG}}^{<1}$ is isomorphic, as an algebraic stack, to $B\mathbb{G}_m \times \operatorname{Spec} \mathbb{Q}$. To understand the quasi-coherent sheaves over this stack, we take a more general approach. Let G be a group scheme acting on a scheme X, and denote by $[X/G]$ the resulting quotient stack introduced in Example 2.3.2. What are the quasi-coherent sheaves on X/G?

To answer this question, we introduce some general terminology that will also be useful when analyzing $\widehat{\mathcal{M}}_{\mathrm{FG}}^{n}$. Let $\mathcal{F} \to \mathcal{C}$ be a fibred category, and let G be a group object in $\mathcal{C}$, acting on an object X in $\mathcal{C}$. Take an object ξ of $\mathcal{F}(X)$. We call it a G-**equivariant object** if, for all $\eta \in \mathcal{F}(U)$, $\operatorname{Hom}_{\mathcal{F}}(\eta, \xi)$ admits an action of $\operatorname{Hom}_{\mathcal{C}}(U, G)$, and the following two conditions are satisfied.

- For any morphism $\eta' \to \eta$ in $\mathcal{F}$ above a morphism $U' \to U$ in $\mathcal{C}$, the induced map $\operatorname{Hom}_{\mathcal{F}}(\eta, \xi) \to \operatorname{Hom}_{\mathcal{F}}(\eta', \xi)$ is equivariant with respect to the homomorphism $\operatorname{Hom}_{\mathcal{C}}(U, G) \to \operatorname{Hom}_{\mathcal{C}}(U', G)$;
- the function $\operatorname{Hom}_{\mathcal{F}}(\eta, \xi) \to \operatorname{Hom}_{\mathcal{C}}(U, X)$ induced by the projection map $p_{\mathcal{F}}$ is $\operatorname{Hom}_{\mathcal{C}}(U, G)$-equivariant.

We write $\mathcal{F}^G(X)$ for the category of G-equivariant objects in $\mathcal{F}(X)$ along with their equivariant morphisms. The following result can be found as [14, Prop. 3.49].

Lemma 4.3.3. Let the notation be as above. Write m for the multiplication map $G \times G \to G$, ρ for the action map $G \times X \to X$, and write p_2 and p_{23} for the quotient maps $G \times X \to X$ and $G \times G \times X \to G \times X$, respectively. Let ξ be an object of $\mathcal{F}(X)$. Then G-equivariant structures on ξ are in one-to-one correspondence with isomorphisms $\varphi\colon p_2^*\xi \xrightarrow{\sim} \rho^*\xi$ in $\mathcal{F}(G \times X)$ such that the diagram

$$\begin{array}{ccc} p_{23}^* p_2^* \xi & \xrightarrow{\quad p_{23}^*\varphi \quad} & p_{23}^*\rho^*\xi \\ & {\scriptstyle (m \times \mathrm{Id}_X)^*\varphi} \searrow \quad \swarrow {\scriptstyle (\mathrm{Id}_G \times \rho)^*\varphi} & \\ & (\mathrm{Id}_G \times \rho)^*\rho^*\xi & \end{array}$$

commutes. □

Lemma 4.3.4. Let G be a group scheme acting on a scheme X, and denote by $[X/G]$ the corresponding quotient stack. Then $\mathsf{QCoh}([X/G])$ is equivalent to the category $\mathsf{QCoh}^G(X)$ of G-equivariant quasi-coherent sheaves of X. □

Proof: There's an obvious morphism $q\colon X \to [X/G]$, which we claim is an fpqc atlas. To prove this, we differentiate between the prestack $[X/G]^{\mathrm{pre}}$ introduced in Example 2.3.2, and the actual stack $[X/G]$. The map $X \to [X/G]$ factors through $[X/G]^{\mathrm{pre}}$. Take a morphism $\operatorname{Spec} R \to [X/G]^{\mathrm{pre}}$ corresponding to a map $\operatorname{Spec} R \to X$. Looking at the definitions, the 2-fibre product $X \times_{[X/G]^{\mathrm{pre}}} \operatorname{Spec} R$ is given by $G \times \operatorname{Spec} R$, and the structure map to $\operatorname{Spec} R$ is just the quotient map. This map is clearly fpqc. Now replace $[X/G]^{\mathrm{pre}}$ with $[X/G]$. The 2-fibre product should be stackified, but as $G \times \operatorname{Spec} R$ is already a stack, nothing changes. Thus we have almost shown $X \to [X/G]$ to be relatively representable by fpqc morphisms. Almost, we say, because upon replacing $[X/G]^{\mathrm{pre}}$ with $[X/G]$, new R-valued points are created, whose 2-fibre product should a priori also be investigated. As being fpqc is an fpqc-local property, however, this is unnecessary, so we're done.

The pullback of $X \to [X/G]$ over itself is $G \times X$, the structural maps being the projection and action map. Lemma 4.2.9 now tells us that a quasi-coherent sheaf $\mathscr{F}$ over $[X/G]$ is exactly a quasi-coherent sheaf $q^*\mathscr{F}$ on X along with an isomorphism $p_2^* q^* \mathscr{F} \xrightarrow{\sim} \rho^* q^* \mathscr{F}$ satisfying the cocycle conditions, which in turn is precisely the data needed to describe a G-equivariant sheaf on X thanks to Lemma 4.3.3. ■

We are now ready to get to the goal that we started out with, namely describing the sheaves of $\mathcal{M}_{\mathrm{FG}}^{<1}$. What about $\mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle$, for various k? Investigating the proof of Lemma 3.2.4, one easily finds that the truncation makes no difference, as formal group laws of height < 1 are entirely determined by their lowest coefficients. It shows that the categories of (quasi-)coherent sheaves over $\mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle$ coincides with that of $\mathcal{M}_{\mathrm{FG}}^{<1}$, so we might as well treat the truncated and non-truncated cases in one go.

Theorem 4.3.5. The category $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}^{<1})$ is equivalent to the category of $\mathbb{Z}$-graded $\mathbb{Q}$-vector spaces. The truncations $\mathcal{M}_{\mathrm{FG}}^{<1} \to \mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle$ are equivalences of stacks for all positive integers k, so the categories $\mathsf{QCoh}(\mathcal{M}_{\mathrm{FG}}^{<1}\langle k\rangle)$ are the same. □

Proof: By Lemma 3.2.4, $\mathcal{M}_{\mathrm{FG}}^{<1}$ is equivalent to $B\mathbb{G}_m \times \operatorname{Spec}\mathbb{Q}$, which in turn, as one can check, is equivalent to $[\operatorname{Spec}\mathbb{Q}/\mathbb{G}_m]$. Now Lemma 4.3.4 tells us that the category $\mathsf{QCoh}([\operatorname{Spec}\mathbb{Q}/\mathbb{G}_m])$ is equivalent to the category of $\mathbb{G}_m$-equivariant sheaves over $\operatorname{Spec}\mathbb{Q}$. Unravelling the definition, a $\mathbb{G}_m$-equivariant sheaf over $\operatorname{Spec}\mathbb{Q}$ is just a $\mathbb{Q}$-module M along with a $\mathbb{Q}[u^{\pm 1}]$-module isomorphism $M[u^{\pm 1}] \to M[u^{\pm 1}]$. Such an isomorphism will be of the form $m \otimes 1 \mapsto m \otimes u^d$, and we may set d to be the degree of m. This uniquely determines a grading of M, and conversely, from any grading we may reconstruct the isomorphism. See also [11, Tag 03LE] for more details. ■

Next, we consider the sheaves on $\widehat{\mathcal{M}}_{\mathrm{FG}}^n$, for various finite n, and then we discuss what changes when passing to the truncations $\widehat{\mathcal{M}}_{\mathrm{FG}}^n\langle k\rangle$. Theorem 3.4.5 and Lemma 4.2.9 together tell us that quasi-coherent sheaves on $\widehat{\mathcal{M}}_{\mathrm{FG}}^n$ correspond to quasi-coherent sheaves on $\operatorname{Spf} R(k,f)$ satisfying suitable overlap conditions on the fibre product $\operatorname{Spf} R(k,f) \times_{\widehat{\mathcal{M}}_{\mathrm{FG}}^n} \operatorname{Spf} R(k,f)$. It makes sense to further investigate this fibre product.

Lemma 4.3.6. Let the notation be as above. The automorphism group $G = \operatorname{Aut}(f(x,y))$, the group structured being composition, acts on $\operatorname{Aut}(\operatorname{Spf} R(k,f)/\widehat{\mathcal{M}}_{\mathrm{FG}}^n)$ □

Before we move to the proof, we ask ourselves what automorphisms of $\operatorname{Spf} R(k,f)$ should be in the first place. We defined formal completions in a purely functorial manner, so that a morphism $\operatorname{Spf} R(k,f) \to \operatorname{Spf} R(k,f)$ is, by definition, a natural transformation of the functor of points. The following lemma helps us out. We omit the proof, but one can check that it follows as a special case of [11, Tag 0AN0].

Lemma 4.3.7. Every morphism $\operatorname{Spf} R(k,f) \to \operatorname{Spf} R(k,f)$ uniquely corresponds to a continuous ring map $R(k,f) \to R(k,f)$, where $R(k,f)$ is endowed with its $\mathfrak{m}$-adic topology. □

Unravelling the definitions, continuous ring maps $R(k,f) \to R(k,f)$ are precisely those ring maps for which the pre-image of $\mathfrak{m}^n$ is $\mathfrak{m}^{F(n)}$ for some integer $F(n)$. Equivalently, a morphism $\operatorname{Spf} R(k,f) \to \operatorname{Spf} R(k,f)$ is a collection of maps $\operatorname{Spec} R(k,f)/\mathfrak{m}^n \to \operatorname{Spf} R(k,f)$ for all n, that are compatible with the immersions $\operatorname{Spec} R/\mathfrak{m}^n \to \operatorname{Spec} R/\mathfrak{m}^{n+1}$.

Proof sketch of Lemma 4.3.6: We begin by describing the G-action. Let $h\colon f(x,y) \to f(x,y)$ be an automorphism of the formal group law $f(x,y)$ over k. Let A be a ring in $\mathsf{Art}(k)$. We begin by defining a map $F_A\colon \operatorname{Def}(A) \to \operatorname{Def}(A)$. Take a deformation class $[f_A]$ in $\operatorname{Def}(A)$, and represent it by a formal group law $f_A(x,y)$. It'll be a lift of $f(x,y)$. The map $W \to k$ representing the isomorphism h may be lifted to A, and looking at Lemma A.2.3 we can choose the lift in such a way, so tha the source $L \to W \to k$ represents $f_A(x,y)$. Write $f'_A(x,y)$ for the target of the lifted isomorphism, and take its deformation class $[f'_A]$. We now define F_A by sending the class $[f_A]$ to $[f'_A]$. It is easily seen that this map is well-defined, and that the various F_A satisfy the desired naturality conditions for it to become a natural isomorphism of functors $\mathrm{Def}(\cdot) \to \mathrm{Def}(\cdot)$. As we know, the functor is representable by $\mathrm{Spf}\, W(k)[[t_1, \dots, t_{n-1}]]$, so that thanks to the Yoneda Lemma, it gives rise to an automorphism $\mathrm{Spf}\, R(k,f) \to \mathrm{Spf}\, R(k,f)$.

Up to isomorphism, the automorphism commutes with the map $\pi\colon \mathrm{Spf}\, R(k,f) \to \widehat{\mathcal{M}}^n_{\mathrm{FG}}$. By Lemma 4.3.7, the automorphism is described entirely by the various maps $\mathrm{Spec}\, R(k,f)/\mathfrak{m}^n \to \mathrm{Spf}\, R(k,f)$. These maps should correspond to a deformation of $f(x,y)$ over $R(k,f)/\mathfrak{m}^n$. Which deformation? By construction, it's described as follows. Writing p for the quotient map $R(k,f) \to R(k,f)/\mathfrak{m}^n$, it should be the image of the deformation class of $[p^* f_{\mathrm{univ}}(x,y)]$ under the map $F_{R(k,f)/\mathfrak{m}^n}\colon \mathrm{Def}\,(R(k,f)/\mathfrak{m}^n) \to \mathrm{Def}\,(R(k,f)/\mathfrak{m}^n)$. Passing to the limit, we find that the automorphism of $\mathrm{Spf}\, R(k,f)$ classifies a possibly non-trivial deformation of $f_{\mathrm{univ}}(x,y)$, which is nonetheless isomorphic, as a formal group law, to $f_{\mathrm{univ}}(x,y)$. This results in a 2-commutative diagram

$$\begin{array}{ccc} \mathrm{Spf}\, R(k,f) & \longrightarrow & \mathrm{Spf}\, R(k,f) \\ & {\scriptstyle F_{R(k,f)}(f_{\mathrm{univ}})} \searrow \quad \swarrow {\scriptstyle f_{\mathrm{univ}}} & \\ & \widehat{\mathcal{M}}^n_{\mathrm{FG}} & \end{array}$$

Now, for any two automorphisms of $f(x,y)$ that give rise to two automorphisms of $\mathrm{Spf}\, R(k,f)$, one can verify from the definitions that the composition of the automorphisms should give to the composed automorphism of $\mathrm{Spf}\, R(k,f)$, from which it follows that we have the desired G-action. ∎

For simplicity, let us restrict attention to the field $k = \mathbb{F}_p$, and look at the natural map of stacks

$$G \times \mathrm{Spf}\, R(\mathbb{F}_p, f) \xrightarrow{\text{(action,projection)}} \mathrm{Spf}\, R(\mathbb{F}_p, f) \times_{\widehat{\mathcal{M}}^n_{\mathrm{FG}}} \mathrm{Spf}\, R(\mathbb{F}_p, f)$$

As it turns out, when we restrict our attention to substacks defined over the subcategory $\mathsf{Art}(\mathbb{F}_p)^{\mathrm{op}}$ of spectra of local Artin rings with residue field $\mathbb{F}_p$, the above map defines an equivalence of stacks, thus turning $\mathrm{Spf}\, R(\mathbb{F}_p, f)$ into a torsor over $\widehat{\mathcal{M}}^n_{\mathrm{FG}}$, as was proved in [4, Thm. 7.17]. If k is a more general field, $\mathrm{Spf}\, R(k,f)$ remains a torsor over $\widehat{\mathcal{M}}^n_{\mathrm{FG}}$, but the group G involved will need to incorporate potential non-trivial automorphisms of k. Also, it should be noted that, if we stick to keeping such automorphisms as part of the data of the objects in $\mathsf{Art}(k)$, then $\mathsf{Art}(k)^{\mathrm{op}}$ will no longer be a subcategory of Aff.

Our aim is to apply Galois descent to the G-torsor structure of $\mathrm{Spf}\, R(\mathbb{F}_p, f)$ over $\widehat{\mathcal{M}}^n_{\mathrm{FG}}$ so as to find an algebraic desciption of the category of quasi-coherent sheaves over $\widehat{\mathcal{M}}^n_{\mathrm{FG}}$. But we only know that $\mathrm{Spf}\, R(\mathbb{F}_p, f)$ is a torsor when viewed as a category over $\mathsf{Art}(\mathbb{F}_p)^{\mathrm{op}}$. There are two ways we might proceed.

- Argue that quasi-coherent sheaves over the relevant stacks are entirely determined by their values on local Artin rings; or

- prove that $\operatorname{Spf} R(\mathbb{F}_p, f)$ is in fact a G-torsor more generally as a stack over Aff.

Let's look at the first approach, and for simplicity moreover set $n = 1$ and $f(x, y) = x + y + xy$. Suppose we wish to describe a quasi-coherent sheaf over $\widehat{\mathcal{M}}^1_{\mathrm{FG}}$, but we only know what this quasi-coherent sheaf does along pullbacks $\operatorname{Spec} A \to \widehat{\mathcal{M}}^1_{\mathrm{FG}}$ for local Artin rings A in $\mathsf{Art}(\mathbb{F}_p)$. Does this entirely determine the sheaf? There's good reason to believe that the answer is no. Take the ring $R = \mathbb{F}_p[u]$, and consider the formal group law $f(x, y)$ corresponding to the map $L_{(p)} \cong \mathbb{Z}_{(p)}[t_1, t_2, \ldots] \to R$ sending t_{p-1} to 1, sending t_{p^n-1} to 0 for all remaining $n > 1$, but sending at least one other t_i to u. By Corollary A.3.10 this formal group law is of height 1. For any local Artin ring A, any ring map $A \to \mathbb{F}_p[u]$ will not have u within its image. In no way can $f(x, y)$ therefore be obtained as the pullback along such a ring map.

Perhaps we can make the first approach work by using the fpqc map from $\operatorname{Spf} R(\mathbb{F}_p, f)$. By Lemma 4.2.9 we know that quasi-coherent sheaves over $\widehat{\mathcal{M}}^1_{\mathrm{FG}}$ can be described as certain quasi-coherent sheaves over $\operatorname{Spf} R(\mathbb{F}_p, f)$, along with an isomorphism on the 2-fibre product. Thanks to Lemma 4.2.3, we know the quasi-coherent sheaves over $\operatorname{Spf} R(\mathbb{F}_p, f)$ are determined by their behaviour on local Artin rings, but what we *don't* know right away is whether the isomorphism on the 2-fibre product is determined on local Artin rings too.

The last question will admit a positive solution if the following holds. Given an R-valued point $\operatorname{Spec} R \to \operatorname{Spf} R(\mathbb{F}_p, f) \times_{\widehat{\mathcal{M}}^1_{\mathrm{FG}}} \operatorname{Spf} R(\mathbb{F}_p, f)$, does there always exist a local Artin ring A with residue field $\mathbb{F}_p$ such that the above map factors through some A-valued point of $\operatorname{Spf} R(\mathbb{F}_p, f) \times_{\widehat{\mathcal{M}}^1_{\mathrm{FG}}} \operatorname{Spf} R(\mathbb{F}_p, f)$? The answer seems to be negative. Take A to be the ring $\mathbb{F}_2[u]/(u^2 - u)$, and let $f(x, y) = x + y + xy$. We claim that $h(t) = t + ut^2 + ut^3$ describes an endomorphism of $f(x, y)$, which must be invertible by Lemma A.1.4. Simply expand

$$\begin{aligned} h(f(x, y)) - f(h(x), h(y)) = {} & -u^2x^3y^3 - u^2x^3y^2 - u^2x^2y^3 - u^2x^2y^2 + ux^3y^3 \\ & + 3ux^3y^2 + 4ux^3y + 2ux^3 \\ & + 3ux^2y^3 + 7ux^2y^2 + 6ux^2y \\ & + 2ux^2 + 4uxy^3 + 6uxy^2 \\ & + 2uxy + 2uy^3 + 2uy^2 \end{aligned}$$

Replace the u^2 with u, simplify, and take the equation modulo 2, to find that it vanishes.

Now focus on the second approach. In order to prove that $\operatorname{Spf} R(\mathbb{F}_p, f)$ is a G-torsor over Aff, we need to verify that the morphism

$$G \times \operatorname{Spf} R(\mathbb{F}_p, f) \xrightarrow{\text{(action,projection)}} \operatorname{Spf} R(\mathbb{F}_p, f) \times_{\widehat{\mathcal{M}}^1_{\mathrm{FG}}} \operatorname{Spf} R(\mathbb{F}_p, f)$$

is an equivalence of stacks over Aff. Given an R-valued point on the right-hand side, does it correspond to some R-valued point on the left-hand side? It would seem that the same example provided above gives us a no, as the automorphism $h(t) = t + ut^2 + ut^3$ cannot be a pullback along any map $\mathbb{Z}/p^n\mathbb{Z} \to R$.

Nonetheless we notice something hopeful in our example. The scheme $\operatorname{Spec}(\mathbb{F}_2[u]/(u^2-u))$ is a disjoint union of two copies of $\operatorname{Spec}\mathbb{F}_2$; taking the pullback of $h(t)$ gives two $\mathbb{F}_2$-valued points of the 2-fibre product, and these points *do* factor through a local Artinian ring — in fact, they are objects of $\mathsf{Art}(\mathbb{F}_2)$ themselves.

A quasi-coherent sheaf on an algebraic stack is fully described once we know what it does on pullbacks to *connected* schemes, which raises the question whether all counterexamples are like the one presented above, and this indeed seems to be the case. Let R a ring. We may assume p is nilpotent, say $p^n = 0$, since we ask for it to admit a map to $\widehat{\mathcal{M}}^1_{\mathrm{FG}}$. Let $h(t)$ define an endomorphism of the multiplicative formal group law $f(x,y) = x+y+xy$. Then it necessarily satisfies the functional equation $h(x+y+xy) = h(x)+h(y)+h(x)h(y)$. Write $h(t) = \sum_{n\geq 1} h_n t^n$, expand both sides of the functional equation, and compare the coefficients in front of $x^{p^k}y^{p^k}$, for all $k \geq 1$. Using some elementary combinatorics, one finds the relation

$$h_{p^k} + \frac{(p^k+1)!}{1!1!(p^k-1)!}h_{p^k+1} + \frac{(p^k+2)!}{2!2!(p^k-1)!}h_{p^k+2} + \cdots$$
$$\cdots + \frac{(p^k+(p^k-1))!}{(p^k-1)!(p^k-1)!1!}h_{p^k+(p^k-1)} + \frac{(p^k+p^k)!}{(p^k)!(p^k)!}h_{p^k+p^k} = h_{p^k}^2$$

With help of Legendre's formula if needed, one finds that all the non-trivial coefficients that occur in the above equation are divisible by p, and hence nilpotent. In particular, it follows that $h_{p^k}^2 - h_{p^k}$ is a multiple of p, and hence nilpotent. (Moreover, if k is large enough, $h_{p^k}^2 - h_{p^k}$ must in fact be zero.) In the second part of the proof of Theorem A.4.2, we saw that, at least over $\mathbb{F}_p$, every automorphism $h(t)$ was entirely determined by its values on h_1, h_p, h_{p^2}, and so on. This part carries over just fine to our situation, so what we may conclude is that $h(t)$ is entirely determined by elements whose difference with their squares are nilpotent.

Ignoring the $h_{p^k+1}, \dots, h_{p^k+(p^k-1)}$ for a moment, the equation $x^2 - x - cp \equiv 0 \pmod{p^n}$ tends to have two distinct solutions; in fact, it probably does so for all n and all constants c, but I'm not the right person for that. The splitting of this polynomial gives rise to an idempotent of the form $x-d$ for some $d \in R$, provided that the characteristic p is not 2. If this all works out, we are able to split up the ring R into components $R/(x-d)$ and $R/(x-d-1)$, and we may repeat this discussion for all the coefficients of $h(t)$ for which this needed. Of course, there tend to be infinitely many coefficients, but this problem goes away when considering k-buds. All things considered, it seems reasonably safe to state the result we're expecting in the form of a question, and continue with it for the rest of this section.

Question 4.3.8. Let $f(x,y)$ be a formal group law of height 1 over the field $k = \mathbb{F}_p$. With respect to the action of $G = \operatorname{Aut}(f(x,y))$ on $\operatorname{Spf} R(k,f)$ over $\widehat{\mathcal{M}}^1_{\mathrm{FG}}$ as constructed in Lemma 4.3.6, we have an equivalence of categories between $\mathsf{QCoh}(\widehat{\mathcal{M}}^1_{\mathrm{FG}})$ and the category $\mathsf{QCoh}^G(\operatorname{Spf} R(\mathbb{F}_p,f))$ of G-equivariant quasi-coherent sheaves over $\operatorname{Spf} R(\mathbb{F}_p,f)$. □

We have to be a little bit careful when we say 'G-equivariant', as our original definition was an equivariant object in the fibred category $\mathsf{QCoh} \to \mathsf{Sch}$, which does not apply in this context as neither G (being abstract and infinite) nor $\operatorname{Spf} R(k,f)$ (being formal) are schemes. However, the equivalent definition in Lemma 4.3.3 generalizes easily to this context.

Question 4.3.9. Let $f(x,y)$ be a k-bud of height 1 over $\mathbb{F}_p$. The group G of k-bud automorphisms of $f(x,y)$ can be seen to act on $\operatorname{Spf} R(k,f)$ over $\widehat{\mathcal{M}}^1_{\mathrm{FG}}\langle k\rangle$ in a manner analogous to the construction in Lemma 4.3.6. With respect to this action, we have an equivalence of categories between $\mathsf{QCoh}(\widehat{\mathcal{M}}^1_{\mathrm{FG}}\langle k\rangle)$ and the category $\mathsf{QCoh}^G(\operatorname{Spf} R(\mathbb{F}_p,f))$ of G-equivariant quasi-coherent sheaves over $\operatorname{Spf} R(\mathbb{F}_p,f)$. □

There are obvious generalizations of the above questions to higher heights, but we'll stick with the current level of generality. There are now three questions we should ask ourselves.

- What is G? This one has already been answered in Theorem A.4.2 and Theorem A.4.3 in the non-truncated and truncated cases, respectively.
- Can we describe the action of G on $\operatorname{Spf} R(k,f)$ algebraically?
- What does it mean, algebraically, to be a G-equivariant sheaf over an affine scheme (or a formal affine scheme) when G is an abstract group?

Let us turn to the second question. We know from Lemma 4.3.7 that an automorphism of $\operatorname{Spf} R(k,f)$ should correspond to a continuous ring automorphism of $R(k,f)$. In our case, we're looking for continuous ring automorphisms of $\mathbb{Z}_p$. Apart from the identity, there are none. Thus, every element in $\mathbb{Z}_p^*$ acts on $\mathbb{Z}_p$ via the identity map.

This result is confusing at first, and in fact it seems in blatant contradiction with [4, Thm. 7.19], where it is shown that the automorphism group of $\operatorname{Spf} \mathbb{Z}_p$ should in fact be in bijection with $\mathbb{Z}_p^*$. The results are perfectly compatible however, and this is because of the 2-commutativity, rather than the strict commutativity, with the structure maps to $\widehat{\mathcal{M}}^1_{\mathrm{FG}}$; the 2-functor expressing this 2-commutativity should be taken as part of the data of an automorphism.

The above paragraph may or may not be disappointing to the reader. If it is, allow me to make up for that by saying that, for higher n, the action of G on $\operatorname{Spf} R(k,f)$ is highly non-trivial, or so I've been told.

We're now ready for the third question, which is to find an algebraic characterization of our G-equivariant sheaves. We will impose the additional constraint that G is finite — something which is reasonable so long as one works with k-buds. It is unclear to me how the generalization to infinite abstract groups should go.

We begin with a finite abstract abelian group G acting on an affine scheme $X = \operatorname{Spec} A$, and let's assume, just because we can, that we are working over an affine base scheme $\operatorname{Spec} R$. We first remark that we said G is an abstract group, rather than a group scheme. This makes little difference, however. Let us write $R[G] = \bigoplus_{g\in G} R$ for the direct sum of $|G|$ copies of R. An element of $R[G]$ is written $\sum_{g\in G} r_g$. For convenience, we view elements of G as being within $R[G]$, so that when we write $h \in R[G]$, what we really mean is $\sum_{g\in G} r_g$, where $r_g = 1$ when $g = h$ and $r_g = 0$ otherwise.

The spectrum $\operatorname{Spec} R[G]$ is a group scheme over R, or equivalently, $R[G]$ is a Hopf algebra. The cogroup structure $m^\sharp \colon R[G] \to R[G] \otimes_R R[G]$ is defined by

$$m^\sharp \colon g \mapsto \sum_{h\in G} g \cdot (h^{-1} \otimes h)\,,$$

and extending R-linearly. The counit of this cogroup structure is the map sending $\sum_{g\in G} r_g$ to r_e, where e denotes the unit of G. We have a coaction $\rho^\sharp$ of $R[G]$ on A defined by sending $a \in A$ to $\sum_{g\in G} g \otimes g(a) \in R[G] \otimes_R A$, which gives rise to an action ρ of $\operatorname{Spec} R[G]$ on $\operatorname{Spec} A$. Conversely, starting out with the action ρ of $\operatorname{Spec} R[G]$ on $\operatorname{Spec} A$, we can easily recover the original G-action.

Take a quasi-coherent sheaf $\mathscr{F}$ on X, corresponding to an A-module $\mathscr{F}(X) = M$. Recall from Lemma 4.3.3 that a G-equivariant structure on $\mathscr{F}$ is a choice of isomorphism $\phi \colon p_2^* \mathscr{F} \xrightarrow{\sim} \rho^* \mathscr{F}$ satisfying certain cocycle conditions. Let's take global sections to see what happens on the level of modules. The map $\phi(X)$ is a map of $R[G] \otimes_R A$-modules $M \otimes_{A,p_2} (R[G] \otimes_R A) \xrightarrow{\sim} M \otimes_{A,\rho} (R[G] \otimes_R A)$. Being a morphism of $R[G] \otimes_R A$-modules, the image of $m \otimes 1_{R[G]} \otimes 1_A$ under the map $\phi(X)$ can always be uniquely written into the form $\sum_{g\in G} m_g \otimes g \otimes 1$, and, these images determine $\phi(X)$. This defines a map $G \times M \to M$, sending (g, m) to the element m_g that we see written above. Thus we find that the coaction is encoded within some kind of a map $G \times M \to M$, and conversely, certain such maps will always give always give rise to a morphism of $R[G] \otimes_R A$-modules.

Which 'certain such maps'? Being an $R[G] \otimes_R A$-module map imposes the following algebraic condition. Take $a \in A$ and $m \in M$, and consider the image of $am \otimes 1_{R[G]} \otimes 1_A$ under $\phi(X)$. It must be $a \cdot \sum_{g\in G} m_g \otimes g \otimes 1_A$. It'd be tempting to equate this to $\sum_{g\in G} m_g \otimes g \otimes a$, but this is false, since the $R[G] \otimes_R A$-module structure of the codomain is described by the coaction ρ. The real image is

$$\begin{aligned}\phi(X) \colon am &\mapsto \sum_{g\in G}\sum_{h\in G} m_g \otimes gh \otimes h(a)\\ &= \sum_{g\in G} m_g \otimes g \otimes g(a)\\ &= \sum_{g\in G} g(a) m_g \otimes g \otimes 1_A\end{aligned}$$

Thus we find the algebraic condition $g(am) = g(a)g(m)$, and by reversing the procedure, we find that any map $G \times M \to M$ satisfying this condition will give to a morphism like $\phi(X)$.

What about the cocycle conditions? These turn out to precisely encode the additional compatibility relation that $(gh)m = g(hm)$. In other words, G acts on M. We thus end up with a remarkably simple conclusion: G-equivariant quasi-coherent sheaves over $\operatorname{Spec} A$ are precisely A-modules admitting a G-action that is twisted over the module structure in the sense that $g(am) = (ga)(gm)$.

We'd like to apply this to the case where G is $(\mathbb{Z}/p^k\mathbb{Z})^*$ and A is $\operatorname{Spf} \mathbb{Z}_p$. Of course, we should be careful, as we're dealing with a formal scheme rather than an ordinary scheme. But in our situation, this poses no issues, because for every $g \in G$, the automorphism on $\operatorname{Spf} \mathbb{Z}_p$ trivially descends to the various $\operatorname{Spec} \mathbb{Z}/p^n\mathbb{Z}$; indeed, all the automorphisms were just identity maps. Being identity maps, this also tells us that $ga = a$ for all $g \in G$ and $a \in \mathbb{Z}_p$. The twisting $g(am) = (ga)(gm)$ therefore simplifies to $g(am) = ag(m)$. This brings us to the following theorem.

Theorem 4.3.10. Let k be an integer greater than p, and write $r = \lfloor \log_p(k) \rfloor$. Then $\mathsf{QCoh}(\widehat{\mathcal{M}}^1_{\mathrm{FG}}\langle k \rangle)$ is equivalent to the category whose objects consists of all towers

$$\cdots \longrightarrow M_3 \longrightarrow M_2 \longrightarrow M_1$$

of $\mathbb{Z}_p$-modules satisfying the following properties.

- The $\mathbb{Z}_p$-module structure on M_n descends to a $\mathbb{Z}/p^n\mathbb{Z}$-module structure;
- the induced morphisms $\mathbb{Z}/p^n\mathbb{Z} \otimes_{\mathbb{Z}/p^{n+1}\mathbb{Z}} M_{n+1} \to M_n$ are isomorphisms of modules;
- each M_n admits an action from the group $G = (\mathbb{Z}/p^r\mathbb{Z})^*$, and the action is compatible with both the tower maps and the $\mathbb{Z}/p^n\mathbb{Z}$-module structure, the latter meaning that, for all $m \in M_k$, $g \in G$ and $a \in \mathbb{Z}/p^k\mathbb{Z}$, we have $g(am) = a(gm)$.

If we restrict attention to $\mathsf{Coh}(\widehat{\mathcal{M}}^1_{\mathrm{FG}}\langle k \rangle)$, then this simplifies to the category of modules over the group ring $\mathbb{Z}_p[(\mathbb{Z}/p^r\mathbb{Z})^*]$ such that the underlying $\mathbb{Z}_p$-module is finitely generated. ■

Combining Theorem 4.2.15, Theorem 4.3.2 and Theorem 4.3.5, we should be able to tell at this point what the coherent sheaves over $\mathcal{M}^{<2}_{\mathrm{FG}}$ should be. They should correspond to a $\mathbb{Z}_p[(\mathbb{Z}/p^r\mathbb{Z})^*]$-module M and a $\mathbb{Z}$-graded finite-dimensional $\mathbb{Q}$-vector space V satisfying at least the following two conditions.

- The underlying $\mathbb{Z}_p$-module of M is finitely generated;
- the tensor product $M \otimes_{\mathbb{Q}_p} \mathbb{Z}_p$ of the underlying $\mathbb{Z}_p$-module of M with $\mathbb{Q}_p$ admits a grading making it naturally isomorphic to the $\mathbb{Z}$-graded $\mathbb{Q}_p$-vector space $V \otimes_{\mathbb{Q}} \mathbb{Q}_p$.

One additional condition is missing, however, namely the compatibility requirement that is needed for the $(\mathbb{Z}/p^r\mathbb{Z})^*$-action. In principle, this should be 'trivial' in the sense that it should follow from writing out the definitions, but with sufficiently many definitions involved, such trivial things aren't all that trivial anymore. Let us put a number in front of my confusion.

Question 4.3.11. What does the compatibility relation of the coherent sheaves on $\mathcal{M}^{<1}_{\mathrm{FG}}$ and $\widehat{\mathcal{M}}^1_{\mathrm{FG}}$ translate to algebraically? □

Whatever the condition may be, the resulting category already looks similar to the description of the category $\mathcal{B}(p)_f$ of Bousfield's. After having specified the compatibility relation, the passage from coherent to quasi-coherent sheaves is described by Theorem 4.3.2 applied to the stack $\mathcal{M}_{\mathrm{FG}}^{<2}$, much in the same way that Bousfield went from $\mathcal{B}(p)_f$ to $\mathcal{B}(p)$.

Finally, the passage from ungraded to graded quasi-coherent sheaves corresponds to going from $\mathcal{B}(p)$ to $\mathcal{B}(p)_*$. The stack $\mathcal{M}_{\mathrm{FG}}^{<2}$ admits an fpqc atlas by $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1^{\pm 1}]$, and quasi-coherent sheaves over $\mathcal{M}_{\mathrm{FG}}^{<2}$ should correspond to modules over $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$ with a certain isomorphism on the fibre product. Asking for graded quasi-coherent sheaves over $\mathcal{M}_{\mathrm{FG}}^{<2}$ is simply equivalent to asking for graded modules over $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$. The element v_1 has degree $2(p-1)$, from which we find that graded modules over $\mathbb{Z}_{(p)}[v_1^{\pm 1}]$ are in a natural one-to-one correspondence with $2p-2$ copies of ungraded such modules. Looking back at the definition of Bousfield's $\mathcal{B}(p)_*$, this is actually no surprise: the condition that $T^{p-1}M_n \cong M_{n+2p-2}$ for all n tells us that objects in $\mathcal{B}(p)_*$ are entirely determined by $2p-2$ objects $M_1, \ldots, M_{2p-2}$ in $\mathcal{B}(p)$.

4.4 Conclusion and scope for future work

Regardless of the final details, there should be no doubt remaining that the analysis in the last two sections provides a promising new look on Bousfield's initial insights. Even though the details aren't entirely filled in, we can take a look back at Section 4.1 and ask how the various details may be interpreted from our new perspective.

For example, the fpqc atlas $\operatorname{Spec} \mathbb{Z}_{(p)}[v_1^{\pm 1}]$ of $\mathcal{M}_{\mathrm{FG}}^{<2}$ gives rise to a pullback functor of quasi-coherent sheaves, which, by Section 2.5, corresponds to the forgetful functor sending $E(1)$-comodules to their underlying modules. A right adjoint, like the functor $\mathcal{U}$, is now easily found: just take the pushforward functor.

Here's another example. For a Noetherian ring R, it can be verified that the category of R-modules is equivalent to the category of coherent R-modules. It is easily envisaged that this generalizes to Noetherian schemes, and further to Noetherian stacks. $\mathcal{M}_{\mathrm{FG}}^{<2}$ being an example of such a stack, we can compute the homological dimension of its category of quasi-coherent sheaves by looking at its category of coherent sheaves instead. This simplifies matters.

We might wonder what distinguishes the case $p = 2$, geometrically speaking. As we know by now, we required p to be an odd prime in Bousfield's paper, but at no point in the last two sections was a hypothesis on the prime p needed. The answer lies in the structure of the group G that occurs when describing the formal Lubin–Tate spectrum as a G-torsor over $\widehat{\mathcal{M}}_{\mathrm{FG}}^{1}\langle k\rangle$. Indeed, we determined that G is $(\mathbb{Z}/p^r\mathbb{Z})^*$, and if p is odd then this group is cyclic, while if $p = 2$, it isn't. This is of course analogous to the fact that the group Γ^n that Bousfield used in the definition of $\mathcal{B}(p)_*$ fails to be cyclic when $p = 2$.

So what *does* happen in the case that $p = 2$? As mentioned earlier, it is known that the comodule category is more complicated algebraically, and in particular, we cannot ensure the degeneration of Adams' spectral sequence. In 1990 a 'follow-up' paper [3] was published in which the restriction of arithmetic localization is eliminated altogether. Spending a few words on it right now is undeniably worth the effort.

To resolve the issue of the infinite homological dimension, Bousfield replaces complex K-theory by a functor K^{CRT} that he calls 'united K-theory'. Given a space X, the united K-homology of X is given by $\{\mathrm{KU}_*(X), \mathrm{KO}_*(X), \mathrm{KT}_*(X)\}$. Here KO is real K-theory and KT is the 4-periodic spectrum of self-conjugate K-theory — God knows what the T stands for. We view this three-element set as an object in a certain category ACRT which takes into account all of the many KO-module operations that exist between the three spectra KU, KO, and KT, along with the stable Adams operations.

In a sense, this category ACRT subsumes our category of $\mathrm{KU}_{(p)}$-comodules, for all p. Miraculously, although the $\mathrm{KU}_{(2)}$-comodule category has infinite homological dimension, ACRT doesn't: its homological dimension is at most 2. In similar fashion to our outline in Section 4.1, this allows Bousfield to embark on an algebraic investigation of the category ACRT, and as a consequence, of the structure of the KU-local homotopy category.

In view of the previous hundred-or-so pages, it would be natural to ask if the constructions in this paper, too, have some sort of geometric interpretation. Bousfield proves that the category ACRT is abelian and has enough injectives — properties which are typical of categories of sheaves over spaces, whatever 'space' may still mean at this point.

Another promising point I'd like to mention is that the methods we have described are amenable to generalizations to higher heights, modulo severe complications. It's better than nothing, however: for instance, as far as I'm aware, no conjectural algebraic description of the category of $E(2)$-comodules seems to exist as of today, but I may well be wrong about this.

Taking into account the points stated above, as well as others that will surely exist (for one, my supervisor explained to me a fun application, which I have forgotten by now), it is my hope that what I've written will at some point spark the enthusiasm of a bright mind, eager to finish up the chaotic remnants that I hereby leave behind.

Appendix A

Formal group laws

Roughly speaking, a formal group law is a formal power series in two variables over a commutative ring that satisfies certain associativity, unitality and commutativity axioms. Formal group laws enter the stage of algebraic topology via the theory of complex-oriented cohomology theories, where they express how Chern classes of line bundles transform under tensor products.

As the language of formal group laws is prominently featured throughout this thesis, we dedicate an appendix to this topic, serving as a reference for the rest of this thesis. We will not go far into the theory, introducing only that which we will need later on, most notably heights and endomorphism rings.

A.1 First definitions

In this section we define formal group laws and formal groups. There are several different but equivalent definitions in the literature. For completeness, we will give two of them here. Start with a commutative ring R. A **formal group law** over R shall be defined to be a power series $f(x, y) \in R[[x, y]]$ satisfying the identities

$$f(x, 0) = f(0, x) = x\,, \qquad f(x, y) = f(y, x) \qquad \text{and} \qquad f(x, f(y, z)) = f(f(x, y), z))\,.$$

The rules should be interpreted as unitality, commutativity, and associativity of the 'operation' defined by $f(x, y)$.

Example A.1.1. There are two simple examples of formal group laws, defined for all rings R. One is the **additive formal group law**, defined by $f(x, y) = x + y$; the other is the **multiplicative formal group law** $f(x, y) = x + y + xy$. There are many other more complicated formal group laws — can the reader find some?

We now recall the definition of the **formal affine line** $\widehat{\mathbb{A}}^1_R$. It is defined to be the functor $R\text{–}\mathsf{Alg} \to \mathsf{Set}$ sending an R-algebra A to the set of nilpotent elements of A. This definition turns $\widehat{\mathbb{A}}^1_R$

into a Zariski sheaf, and in fact, it is a formal scheme, as it is obtained as the filtered colimit over $\operatorname{Spec} R[x]/(x^N)$, as N ranges over the positive integers.

Lemma A.1.2. Formal group laws over R are in one-to-one correspondence with lifts of the functor $\widehat{\mathbb{A}}^1_R$ to Ab. □

Proof: Given a formal group law f over R, we may define a lift F of $\widehat{\mathbb{A}}^1_R$ to Ab as follows. For every R-algebra A, we attempt to endow $\widetilde{\mathbb{A}}^1_R(\operatorname{Spec} A)$ with an abelian group structure by sending two nilpotent elements x and y to $f(x, y)$. As x and y are nilpotent, the expression $f(x, y)$ is finite, so this is a well-defined operation. Let us verify that it indeed defines a group structure on $\widehat{\mathbb{A}}^1_R(\operatorname{Spec} A)$.

As x and y are nilpotent, the resulting expression is finite, so this is well-defined; moreover, the identities defining a formal group law ensure that this structure is indeed that of a commutative monoid. It remains to be shown that inverses of non-zero element exist. We are done if we find a power series $\iota(x) \in R[\![x]\!]$ such that $f(x, \iota(x)) = f(\iota(x), x) = 0$. This power series $\iota(x)$ can be constructed term by term. That is, write $f(x, y) = \sum_{i,j} c_{ij} x^i y^j$ and start with the expression

$$\sum_{i,j} c_{ij} x^i \Big(\sum_k d_k x^k \Big)^j = 0 \,.$$

Write down the constant term on the left-hand side, and equate it to 0. We find that this yields nothing new. Next, write down the coefficient of the x^1-term on the left-hand side, and set it to zero, giving us

$$c_{10} + c_{01} d_1 + c_{11} d_0 = 0 \,.$$

Put $d_0 = 0$, and $d_1 = -1$. The coefficient of the x^2-term gives us

$$c_{20} + c_{11} d_1 + c_{02} \, (2 d_0 d_2 + d_1^2) = 0 \,,$$

so that we can figure out what d_2 is. Going on, we find that the coefficient of the x^3-term is an expression in the c_{ij}, in d_0, d_1, d_2, and d_3. As we already know d_0, d_1, d_2, we can solve this equation as well. As one can verify, this process continues indefinitely, thus giving us a way to inductively construct every coefficient of the formal power series of $\iota(x)$.

Associating, to every set $\widehat{\mathbb{A}}^1_R(\operatorname{Spec} A)$, the group structure as defined above gives rise to the desired functor $F \colon R\text{–}\mathsf{Alg} \to \mathsf{Ab}$. Indeed, our association is natural in the sense that, if $\pi^\sharp \colon A \to A'$ is a map of R-algebras, the associated map $F(\pi) \colon F(\operatorname{Spec} A) \to F(\operatorname{Spec} A')$ sending a nilpotent element x to $\pi^\sharp(x)$ is a group morphism: $\pi^\sharp$ preserves 0 by construction, and the identity $\pi^\sharp(f(x, y)) = f(\pi^\sharp(x), \pi^\sharp(y))$ easily follows from the definition.

Conversely, given a lift F of $\widehat{\mathbb{A}}^1_R$ to Ab, one can reconstruct the formal group law $f(x, y)$ from which it is defined by considering the group structure on $\widehat{\mathbb{A}}^1_R(\operatorname{Spec} A_N)$ for increasing N, where A_N is the R-algebra $R[x, y]/(x^N, y^N)$. ■

We now know that a formal group law over a ring R can be interpreted as either a lifted functor $F\colon R\text{–}\mathsf{Alg} \to \mathsf{Ab}$ and as a formal power series $f(x, y)$ in R. What do the natural transformations between two lifted functors represent? This brings us to the following definition. A **morphism of formal group laws** $h\colon f(x, y) \to g(x, y)$ over R is a formal power series $h(t) \in R[[t]]$, with constant term 0, such that $f(h(x), h(y)) = h(g(x, y))$.

The requirement that the constant term be 0 makes life easier rather than harder. This is because compositions of formal power series need not be well-defined without this assumption. For instance, if $f(t) = 1 + t + t^2 + \cdots$, and $g(t) = 1$, then $f(g(t))$ is not a well-defined power series. Another motivation for adding this assumption is the following lemma.

Lemma A.1.3. Let $f(x, y)$ and $g(x, y)$ be two formal group laws over a ring R. Then there is a one-to-one correspondence between morphisms of formal group laws $f(x, y) \to g(x, y)$, and natural transformations from the functor F corresponding to the formal group law $f(x, y)$, to the functor G corresponding to the formal group law $g(x, y)$. □

Proof: Given a morphism of formal group laws, define the natural transformation $F \to G$ as follows. For every affine R-scheme $\operatorname{Spec} A$, define a map $F(\operatorname{Spec} A) \to G(\operatorname{Spec} A)$ by sending a nilpotent element x to the nilpotent element $h(x)$. Essentially by assumption, this respects the group structures. Conversely, given a natural transformation $F \to G$, we can reconstruct the formal power series defining the morphism of formal group laws by considering the maps $F(\operatorname{Spec} R[x]/(x^N)) \to G(\operatorname{Spec} R[x]/(x^N))$, as N ranges over the positive integers. ■

If $h_1 : f(x, y) \to f'(x, y)$ and $h_2 : f'(x, y) \to f''(x, y)$ are two morphisms of formal group laws over a ring R, then there is an obvious way to compose these two morphisms: the identity

$$f(h_1 h_2(x), h_1 h_2(y)) = h_1(f'(h_2(x), h_2(y))) = h_1 h_2(f''(x, y))$$

yields a composed morphism $h_1 h_2 : f(x, y) \to f''(x, y)$, and this composition is associative in an appropriate sense. In particular, this allows us to define what it means for a morphism $h\colon f(x, y) \to g(x, y)$ to have an inverse, which we denote by h^{-1}. We emphasize that this inverse is not the *multiplicative* inverse of $h(t)$, but the inverse under *composition*.

Lemma A.1.4. If $h(t)$ is a formal power series in a single variable such that $h(0) = 0$. Then a formal power series $h^{-1}(t)$ such that $h^{-1} \circ h(t) = h \circ h^{-1}(t) = t$ exists if and only if $h'(0)$ is a unit. □

Proof: The coefficients of the power series $h^{-1}(t)$ can be found inductively. Write $h(t) = a_1 t + a_2 t^2 + \cdots$, and $h^{-1}(t) = b_1 t + b_2 t^2 + \cdots$. Take the equation

$$a_1(b_1 t + b_2 t^2 + \cdots) + a_2(b_1 t + b_2 t^2 + \cdots) + \cdots = t$$

and compare the coefficients on both sides; this yields the equations

$$\begin{aligned} a_1 b_1 &= 1\,, \\ a_1 b_2 + a_2 b_1^2 &= 0\,, \\ a_1 b_3 + 2a_2 b_1 b_2 + a_3 b_1^3 &= 0\,, \end{aligned}$$

and so on. Clearly b_1 can only exist if a_1 is a unit, and once we know that we see that b_{i+1} is determined uniquely by the a_j and $b_1, \ldots, b_i$. ■

If $h\colon f(x,y) \to g(x,y)$ is an invertible morphism of formal group laws, we call it an **isomorphism**. If, in fact, $h'(0)$ is not just a unit but it equals 1, we say that we have a **strict isomorphism of formal group laws**.

Let $\pi\colon\ \operatorname{Spec} R' \to \operatorname{Spec} R$ be a map of affine schemes, corresponding to a ring map $\pi^\sharp\colon R \to R'$. Let $f(x,y)$ be a formal group law over R represented by a functor $F\colon R\text{–}\mathsf{Alg} \to \mathsf{Ab}$. We construct a functor $F'\colon R'\text{–}\mathsf{Alg} \to \mathsf{Ab}$ as follows. Whenever we have an affine R'-scheme $\operatorname{Spec} A \to \operatorname{Spec} R'$, compose the structure map of $\operatorname{Spec} A$ with π so as to view $\operatorname{Spec} A$ as an R-scheme. Now let $F'(\operatorname{Spec} A)$ be $F(\operatorname{Spec} A)$.

Lemma A.1.5. With the notation as above, F' corresponds to a formal group law f' over R'. Explicitly, the formal power series $f'(x,y)$ is given by $\sum_{i,j} \pi^\sharp(c_{ij})\, x^i y^j$. □

Depending on whether you prefer rings or affine schemes, the formal power series f' obtained in the lemma above can be either called the **pushforward** or the **pullback** of f. We stick to the latter convention, and also denote f' by $\pi^* f$.

Proof of Lemma A.1.5: This is true essentially by construction. For every affine R'-scheme $\operatorname{Spec} A$, the group structure on $F'(\operatorname{Spec} A)$ is defined to be inherited from the group structure on $F(\operatorname{Spec} A)$. It must therefore be the case that the corresponding formal power series is obtained by substitution. ■

There is one more definition we wish to discuss. In Chapter 3 we will be interested in the so-called moduli stack of formal groups. Roughly speaking, this moduli stack is constructed as follows. We have a groupoid-valued presheaf that sends a ring R to the groupoid of formal group laws over R, where the morphisms are given by isomorphisms defined in the sense above. This presheaf fails a descent condition, so we sheafify our presheaf; the resulting sheaf defines this moduli stack.

To accurately describe this sheaf, it can be useful to define a ‘global’ version of a formal group law, i.e. an object that is essentially just a formal group law Zariski-locally. Confusingly, there are several ways of going about this, and the terminology used to describe such an object is not fixed, but they may rightfully be called **formal groups**. Luckily for us, we will never need such models in any critical way, and we will instead refer the reader to [4, Ch. 2] or [5, Lecture 11] for two ways of making formal groups precise.

A.2 Lazard's theorem

For every ring R, we can consider the collection of all formal group laws over R, denoted $\mathrm{FGL}(R)$. Given a morphism $\pi^\sharp\colon R \to R'$ of commutative rings, we just learned that we have a pullback map $\mathrm{FGL}(R) \to \mathrm{FGL}(R')$. Thus FGL defines a functor.

Lemma A.2.1. The functor FGL is representable. That is, there's is a ring L such that formal group laws over R are in one-to-one correspondence with ring morphisms from L to R. □

The ring L representing FGL is called the **Lazard ring**, and is traditionally denoted by L.

Proof of Lemma A.2.1: Every power series $f(x,y) \in R[[x,y]]$ can be written as a formal sum $f(x,y) = \sum_{i,j} c_{ij}x^iy^j$, where c_{ij} are unspecified elements of R. As said before, the identities defining a formal group law put constraints on the c_{ij}, namely $c_{10} = c_{01} = 1$, $c_{i,0} = c_{0,i} = 0$ for all $i \neq 1$, and $c_{ij} = c_{ji}$ for all i and j. Finally, the associativity identity $f(x, f(y,z)) = f(f(x,y),z)$ imposes additional polynomial constraints that nobody ever writes down as the precise form is irrelevant. We do it anyway, and advice the reader not to read this more than once. We start out with the identity

$$\sum_{i,j} c_{ij}x^i\Big(\sum_{k,\ell} c_{k\ell}y^kz^\ell\Big)^j = \sum_{i,j} c_{ij}\Big(\sum_{k,\ell} c_{k\ell}x^ky^\ell\Big)^i z^j\,.$$

Now fix $i = \alpha$, $j = \beta$, and $z = \gamma$, and equate the coefficients of $x^\alpha y^\beta z^\gamma$ that both sides of the above equation give us. This yields

$$\sum_{j=0}^{\infty} \sum_{\substack{k_1,\dots,k_j \\ k_1+\dots+k_j=\beta}} \sum_{\substack{\ell_1,\dots,\ell_j \\ \ell_1+\dots+\ell_j=\gamma}} c_{\alpha j}\, c_{k_1\ell_1}\cdots c_{k_j\ell_j} = \sum_{i=0}^{\infty} \sum_{\substack{k_1,\dots,k_i \\ k_1+\dots+k_i=\alpha}} \sum_{\substack{\ell_1,\dots,\ell_i \\ \ell_1+\dots+\ell_i=\beta}} c_{i\gamma}\, c_{k_1\ell_1}\cdots c_{k_i\ell_i}\,.$$

Notice that this is a polynomial equation despite the infinite summation. Define the ring L to be the ring of formal power series $\mathbb{Z}[c_{ij}]$ modulo the infinitely many polynomial constraints imposed by the identities of the formal group laws. Essentially by construction then, every map $L \to R$ corresponds to a formal group law. ■

Despite its messy construction, the structure of the Lazard ring is surprisingly simple, as the next result shows. What follows is based on [5, Lecture 2].

Theorem A.2.2 (Lazard's Theorem). The Lazard ring is non-canonically isomorphic to the free polynomial algebra $\mathbb{Z}[t_1, t_2, \dots]$. □

We endow the Lazard ring with a natural grading. The variables c_{ij} carry degree $2(i+j-1)$, while the variables t_i carry degree $2i$. With respect to this grading, the isomorphism becomes an isomorphism of graded rings. The choice of grading, which may seem somewhat unnatural, will explain itself in Theorem 1.3.9.

Proof sketch of Theorem A.2.2: If $f(x, y)$ is a formal group law over a ring R, and we make the substitution $g(t) = t + b_1 t^2 + b_2 t^3 + \cdots$, then the formal power series $g(f(g^{-1}(x), g^{-1}(y)))$ is also a formal group law over R. In particular, if $f(x, y) = x + y$, then $g(g^{-1}(x) + g^{-1}(y))$ can be viewed as a formal group law over the polynomial ring $\mathbb{Z}[b_1, b_2, \ldots]$. This yields a map $\varphi\colon L \to \mathbb{Z}[b_1, b_2, \ldots]$. If we endow the b_n with degree $2n$, one can see that φ becomes a map of graded rings.

The map φ need *not* be an isomorphism. However, we have the resulting result, whose proof can be found in [5, Lecture 3]. Let I and J be the ideals of L and $\mathbb{Z}[b_1, b_2, \ldots]$, respectively, generated by the elements of positive degree. Then the map φ induces an injection $(I/I^2)_{2n} \to (J/J^2)_{2n} \cong \mathbb{Z}\, b_n$; the image of this map is $p\mathbb{Z}$ if n is of the form $p^k - 1$ for some k, and is $\mathbb{Z}$ otherwise.

As a consequence, $(I/I^2)_{2n}$ is canonically isomorphic to $\mathbb{Z}$. For all n, choose a lift $t_n \in L_{2n}$ of a generator of $(I/I^2)_{2n}$. This defines a map of rings $\partial\colon \mathbb{Z}[t_1, t_2, \ldots] \to L$. The map ∂ is the desired isomorphism, which we prove now. It is injective, as can be seen by noting that the composition $\varphi \circ \partial$, which sends t_n to $b_n \left(\bmod J^2\right)$ must be injective too. It is surjective, as is proved by induction on the degree. First note that, in degree 2, surjectivity of ∂ can be verified explicitly as the relations in L are simple in low degree. Now say we know that ∂ in degree $2n$. Then it must contain the degree-$(2n{+}2)$ part of I^2. But by assumption, the image of ∂ also has a generator for $(I/I^2)_{2n+2} \cong \mathbb{Z}$. Hence ∂ reaches all if L_{2n+2}. ■

Given a ring R, we can not only consider the set of formal group laws over R, but also set of isomorphisms between any two formal group laws over R (where we do *not* identify between an isomorphism and its inverse), temporarily denoted IsoFGL(R). Given a ring map $\pi^\sharp\colon R \to R'$, any isomorphism $h\colon f(x, y) \to g(x, y)$ of formal group laws over R can be pulled back to an isomorphism $\pi^* h\colon \pi^* f(x, y) \to \pi^* g(x, y)$ of formal group laws over R' by substitution of the coefficients of the power series defining h. This turns IsoFGL into a functor $\mathsf{Ring} \to \mathsf{Set}$.

Lemma A.2.3. The functor IsoFGL is representable by a ring, commonly denoted W, and it equals $L[a_1^{\pm 1}, a_2, a_3, \ldots]$. □

Proof: An isomorphism $h\colon f(x, y) \to g(x, y)$ of formal group laws over a ring R is entirely determined by the formal group law $f(x, y)$ together with the formal power series $h(x)$. This is because we can recover g by the formula $g(x, y) = h^{-1}(f(h(x), h(y)))$. Thus every isomorphism of formal group laws over R is uniquely determined by a single formal group law and a single invertible power series. This data corresponds precisely to a morphism $L[a_0^{\pm 1}, a_1, a_2, \ldots] \to R$. ■

Given a ring R, the formal group laws over R and the isomorphisms between them give rise to a groupoid of formal group laws. By Lemma A.2.1 and Lemma A.2.3, it follows that the pair

$(\mathrm{Hom}(L,R), \mathrm{Hom}(W,R))$ comprise the objects and morphisms of a groupoid. In more technical terms, this tells us that the pair (L, W) is a **Hopf algebroid**. We return to this in Section 2.1.

Finally, we remark that we could repeat much of the story above when replacing isomorphisms with strict isomorphisms. The proof of Lemma A.2.3 carries over to show that strict isomorphisms are representable by some ring W^s, namely the ring $L[a_2, a_3, \dots]$, and this turns (L, W^s) into another example of a Hopf algebroid.

A.3 Heights

A reasonable question to ask is how we can prove that two given formal group laws are *not* isomorphic. In this section we introduce an invariant that can tell certain group laws apart, but which is also of interest from a more theoretical point of view. I also refer the reader to [5, Lectures 12–14 and 19] and [7, Section 4.2] for other sources covering heights of formal group laws.

Fix a formal group law $f(x,y)$ over a ring R. For every positive integer n, the n-**series** $[n](t)$ of f is defined recursively via $[0](t) = 0$ and $[n+1](t) = f([n](t), t)$. The intuition here is that the n-series represents a multiplication by n. Now fix a prime number p, and write $v_n(f)$, or v_n for short if the context is clear, for the coefficient of t^{p^n} in the p-series $[p](t)$ of $f(x,y)$. We say $f(x,y)$ has **height** $\geq n$ if $v_k = 0$ for all $k < n$, and $f(x,y)$ has **height** n if moreover v_n is an invertible element of R. We remark that, for any formal group law $f(x,y)$ over R, v_0 equals p, so saying that a formal group law is of height $\geq n$ for some non-trivial n implicitly tells you something about the ring R as well: it is of characteristic p.

Lemma A.3.1. Let R be a ring of characteristic p, and let $f(x,y)$ be a formal group law over R. Then either $[p](t) = 0$, or there exists a unique $n \geq 0$, namely the height, such that $[p](t) = g(t^{p^n})$ for some formal power series $g(t)$ with $g'(0) \neq 0$. □

Proof: Let's take a morphism $h\colon f(x,y) \to g(x,y)$ of formal group laws. Assume that $h'(0) = 0$. Differentiate the identity $f\big(h(x), h(y)\big) = h\big(g(x,y)\big)$ with respect to y, use straightforward differentiation rules, and finally apply $y = 0$ to find the identity

$$h'(0)\frac{\partial}{\partial y}f\big(h(x), 0\big) = h'(x)\frac{\partial}{\partial y}g(x,0)\,.$$

As $h'(0) = 0$, the left-hand side vanishes, hence so must the right-hand side. But $\partial_y g(x,0) = 1 + g_{11}x + g_{21}x^2 + \cdots$, which is always a unit in R, and so it must be the case that $h'(x)$ vanishes too. Thus if $h'(0) = 0$ then also $h'(t) = 0$ for all t.

Next, we make a simple observation. If $h(t)$ is a formal power series over R, say $h(t) = h_0 + h_1 t + \cdots$, and $h'(t) = 0$, then we can write $h(t) = k(t^p)$ for some other formal power series k. Indeed, the vanishing of the derivative tells us that $nh_n = 0$ for all $n \geq 0$, so for all n that is coprime to p, h_n will vanish. We thus find $h(t) = h_0 + h_p t^p + h_{2p}t^{2p} + \cdots$, as desired.

We use the above two observations to prove that, for any morphism $h\colon f(x,y) \to g(x,y)$ of formal group laws, either $h = 0$, or $h(t) = k(t^{p^n})$ for some formal power series $k(t)$ with $k'(0) \neq 0$. To prove this, we argue as follows. First off, let's assume $h'(0) \neq 0$. Then we take $n = 0$ and we're done. So assume $h'(0) = 0$. Then by the above discussion, $h(t) = k(t^p)$ for some formal power series k. In turn, if $k'(0) \neq 0$ then we're done. If $k'(0) = 0$, then we might hope that $k'(t) = 0$ and we could repeat the procedure. But we need to establish k as some morphism of formal group laws in order to do that. We claim that k is indeed such a morphism. Let F^*g be the pullback of the formal group law $g(x,y)$ along the Frobenius endomorphism F of R. Then

$$k(F^*g(x^p, y^p)) = k(g(x,y)^p) = h(g(x,y)) = f(h(x), h(y)) = f(k(x^p), k(y^p)).$$

This may be interpreted as an equality of formal power series in the variables x^p and y^p, so we may as well make the substitution $x^p \mapsto x$ and $y^p \mapsto y$ to find that k is indeed a morphism of formal group laws. We may thus conclude that either $k'(0) \neq 0$ or $k'(t) = 0$, and repeat the procedure.

The statement of the lemma follows by applying the above discussion to the special case of the endomorphism $[p]\colon f(x,y) \to f(x,y)$. To prove that this is indeed an endomorphism, we need the following claim.

Lemma A.3.2. If $h_1, h_2\colon f(x,y) \to f(x,y)$ are two endomorphisms of a formal group law $f(x,y)$, then so is the power series $k(t) = f(h_1(t), h_2(t))$. □

Proof: This is a consequence of the following chain of manipulations:

$$\begin{aligned}
f(k(x), k(y)) &= f\big(f(h_1(x), h_2(x)), f(h_1(y), h_2(y))\big) && \text{by definition}\\
&= f\big(h_1(x), f(h_2(x), f(h_1(y), h_2(y)))\big) && \text{associativity}\\
&= f\big(h_1(x), f(h_2(x), f(h_2(y), h_1(y)))\big) && \text{commutativity}\\
&= f\big(h_1(x), f(f(h_2(x), h_2(y)), h_1(y))\big) && \text{associativity}\\
&= f\big(h_1(x), f(h_1(y), f(h_2(x), h_2(y)))\big) && \text{commutativity}\\
&= f\big(f(h_1(x), h_1(y)), f(h_2(x), h_2(y))\big) && \text{associativity}\\
&= f\big(h_1(f(x,y)), h_2(f(x,y))\big) && h_1 \text{ and } h_2 \text{ are endomorphisms}\\
&= k(f(x,y)) && \text{by definition}
\end{aligned}$$

■

We use this lemma in the special case where $h_2(t) = t$ to prove, inductively, that the n-series $[n](t)$ defines an endomorphism of $f(x,y)$. Indeed, this is clearly true for $n = 0$, and if it is true for some $n = k$, then the above claim implies that it holds for $n = k + 1$ as well. ■

Lemma A.3.3. Let $f(x,y)$ and $g(x,y)$ be two formal group laws over a commutative ring R, and fix a prime p. If $f(x,y)$ and $g(x,y)$ are isomorphic, then the heights of $f(x,y)$ and $g(x,y)$ are the same. That is, if $f(x,y)$ is of height $\geq n$, then so is $g(x,y)$, and also if $f(x,y)$ is of height exactly n, then so is $g(x,y)$. □

Proof: If $h\colon f(x,y) \to g(x,y)$ defines an isomorphism, then one can verify that $[p]_f(t) \circ h(t) = h \circ [p]_g(t)$. Now suppose $f(x,y)$ is of height $\geq n$, and g is of height $\geq m$. Expanding both sides of the equality, and invoking Lemma A.3.1 to simplify our lives, we find that the lowest-order term on the left-hand side will be $v_m(g)(h_1 t)^{p^m}$, where $h_1 = h'(0)$, which we recall must be a unit in R, so that $v_m(g)h_1^{p^m} \neq 0$; the lowest-order term on the right-hand side, on the other hand, is $v_n(f)t^{p^n}$. It follows that $m = n$. Moreover, we see that if $v_n(f)$ is a unit in R, then so is $v_m(g)$. ■

Example A.3.4. The p-series of the multiplicative formal group law is $[p](t) = (1+t)^p - 1$. If R is a ring in which $p = 0$, then this simplifies to t^p. Thus for rings of characteristic p, the multiplication formal group law has height exactly 1.

Example A.3.5. The p-series of the additive group law is $[p](t) = pt$. If $p = 0$ in our ring, then the p-series vanishes altogether, in which case we say the formal group law has **infinite height**. From Lemma A.3.3 it follows that over such rings the additive and multiplicative formal group laws are not isomorphic.

We point out that there may well be other rings over which the additive and multiplicative formal group laws are isomorphic. As an example, let R be the ring $\mathbb{Q}$, in which case the additive and multiplicative group law have height 0, regardless of the choice of prime. An explicit isomorphism from the multiplicative formal group law to the additive formal group law is given by the power series associated to $h(t) = e^t - 1$, while the inverse map is given by $h(t) = \log(t+1)$. in fact, this is just one instance of a much more striking result.

Lemma A.3.6. If R is a $\mathbb{Q}$-algebra, any two formal group laws will always be isomorphic. □

Proof: Let $f(x,y)$ be a formal group law over our $\mathbb{Q}$-algebra R. Let $f_2(x,y)$ be the algebraic derivative $\partial_y f(x,y)$, and algebraically define the **logarithm** of $f(x,y)$ to be

$$\log_f(x) = \int_0^x \frac{1}{f_2(t,0)}\,dt\,.$$

(We point out that we critically need R to be a $\mathbb{Q}$-algebra here: the primitive of x^n is $x^{n+1}/(n+1)$, which need not exist if R is not a $\mathbb{Q}$-algebra.) We claim that $\log_f(f(x,y)) = \log_f(x) + \log_f(y)$, which would establish an isomorphism between f and the additive formal group law. To prove this, simply differentiate both sides with respect to x, find that these derivatives coincide, and repeat the story when taking the derivatives with respect to y. As both sides have equal constant terms, they must represent the same formal power series. ■

It will be useful to introduce the following definition, based on our notion of height. The terminology is confusing at first, but will become natural in Section 3.2. Let $f(x,y)$ be a formal group law over a ring R. We say $f(x,y)$ is of **height** $< n$ if the coefficients $v_0, \ldots, v_{n-1}$ occurring in the p-series of $f(x,y)$ generate the ring R. Equivalently, if $f(x,y)$ corresponds to a ring map $g\colon L \to R$, we ask for the ideal $(g(v_0), \ldots, g(v_{n-1}))$ to be all of R.

Lemma A.3.7. Let $f(x,y)$ and $g(x,y)$ be two isomorphic formal group laws over R. If $f(x,y)$ is of height $< n$, then so is $g(x,y)$. □

Proof: Let us define R' to be the quotient ring $R/(v_0(g),\dots,v_{n-1}(g))$. Under the obvious quotient map, $g(x,y)$ gets pulled back to a formal group law of height $\geq n$. The isomorphism between $f(x,y)$ and $g(x,y)$ gets pulled back to an isomorphism of formal group laws over R', so that by Lemma A.3.3, $f(x,y)$ is of height $\geq n$ over R'. This means that $v_0(f),\dots,v_{n-1}(f)$ vanish in R', or equivalently, that they are contained within the ideal $(v_0(g),\dots,v_{n-1}(g))$. As $f(x,y)$ was of height $< n$, $(v_0(f),\dots,v_{n-1}(f))$ generates all of R, hence so must $(v_0(g),\dots,v_{n-1}(g))$. ■

In Section 3.4, we will be interested in the following weakening of the notion of height. We remark that the terminology is non-standard. Let $f(x,y)$ be a formal group law over R. We say $f(x,y)$ is **almost of height** n if the coefficients $v_0,\dots,v_{n-1}$ occurring in the p-series of $f(x,y)$ are nilpotent in R, and the element v_n is a unit. We have the following variant of Lemma A.3.3.

Lemma A.3.8. Let $f(x,y)$ and $g(x,y)$ be two formal group laws over R, which we assume to be isomorphic. If $f(x,y)$ is almost of height n, then so is $g(x,y)$. Moreover, if N is the smallest integer for which $v_i(f)^N = 0$ for all $i = 0,\dots,n-1$, then there exists an integer $m(N)$, depending on N but not on $g(x,y)$, such that $v_i(g)^{m(N)} = 0$ for all $i = 0,\dots,n-1$. □

Proof: Since $v_0 = p$ always, the fact that $f(x,y)$ is almost of height n tells us that $p^k = 0$ in R for some fixed k. Write π for the projection $R \to R/(p)$. Notice that, if x is an element in $R/(p)$ such that $x^n = 0$, then any lift of x to an element $\overline{x}$ in R satisfies $\overline{x}^{kn} = 0$ as well.

Denote by F the Frobenius endomorphism on $R/(p)$. The formal group law $f(x,y)$ is almost of height n, so that there's N such that $v_i(f)^N = 0$ for all i. Certainly $(F^N)^*\pi^*f(x,y)$ will then be of height exactly n. (In fact, we need far less compositions of F, but this is irrelevant.) Moreover, if $h\colon f(x,y) \to g(x,y)$ defines our isomorphism of formal group laws, then $(F^N)^*\pi^*h$ defines an isomorphism from $(F^N)^*\pi^*f(x,y)$ to $(F^N)^*\pi^*g(x,y)$. By Lemma A.3.3, we now see that $(F^N)^*\pi^*g(x,y)$ is of height exactly n. In particular, it follows that $v_i(g)^{Np}$ vanishes mod (p), hence $v_i(g)^{Npk} = 0$ in R. As neither N nor p nor k depend on $g(x,y)$, the result follows. ■

The coefficients of the p-series of a formal group law $f(x,y) = \sum_{i,j} c_{ij}x^iy^j$ over R will certainly admit a polynomial expression in terms of the c_{ij}, which the reader can show inductively if he so wishes. In particular, this is true for the p^n-th coefficient, and we denote the corresponding polynomial in the c_{ij} by v_n. Every element v_n corresponds to a fixed element in the Lazard ring L, where we emphasize the dependence on the chosen prime p.

If we use the presentation of L as the polynomial ring $\mathbb{Z}[t_1,t_2,\dots]$, then the v_n admit a surprisingly nice description. We recall that the isomorphism between L and $\mathbb{Z}[t_1,t_2,\dots]$ was not canonical, but we saw in the proof of Theorem A.2.2 that, if I is the ideal of L generated by the elements of positive degree, the isomorphism between $(I/I^2)_{2n}$ and $\mathbb{Z} \cdot t_n$ *is* canonical. It follows that we may expect v_n to correspond to a particular well-defined element in $\mathbb{Z} \cdot t_{p^n-1}$.

Lemma A.3.9. With the notation as above, the image of v_n in $(I/I^2)_{2(p^n-1)}$ is $p^{p^n-1} - 1$. □

Proof: The proof becomes transparent once we look more closely into the construction of the (non-unique) isomorphism between L and $\mathbb{Z}[t_1, t_2, \ldots]$ — something we did not do, as we will not need it at any point in the future. We therefore refer the reader to [5, Prop. 1 of Lecture 13] instead. ■

Corollary A.3.10. If we localize at p, we may choose a particular isomorphism $L_{(p)} \cong \mathbb{Z}_{(p)}[t_1, t_2, \ldots]$, where each t_{p^n-1} is given by v_n. □

Proof: For any p-local ring R the integer $p^{p^n-1} - 1$ is invertible in R, and so, up to a unit, we may as well say that the image of v_n in $\mathbb{Z}t_{p^n-1}$ is a generator. ■

Corollary A.3.11. Let R be a ring of characteristic p. Then, for every positive integer n, as well as $n = \infty$, there exists a formal group law of height exactly n. □

Proof: A formal group law over R corresponds to a ring map $L \cong \mathbb{Z}[t_1, t_2, \ldots] \to R$. Consider the map which sends t_{p^n-1} to 1, but which takes the other t_i to 0. From the above lemma, the element v_n gets sent to -1, while the other v_i vanish. Whatever the formal group law looks like, it is clearly of height n. Finally, an example of a formal group law of infinite height is given by the additive formal group law by Example A.3.5. ■

We might ask ourselves if we can somehow classify the formal group laws up to isomorphism. In general, we can expect this to be an intractable problem: Given that formal group laws correspond to infinitely many choices of elements in our ring, there are simply too many to list them all.

But let's look at the very simplest scenarios only. We've already solved one case. In the proof of Lemma A.3.6, we saw that the ring being a $\mathbb{Q}$-algebra allowed for sufficiently many degrees of freedom (in the sense of enough having multiplicative inverses) to construct many isomorphisms of formal group laws. It is reasonable to expect a similar such freedom in the case of fields of characteristic p. Over such fields, we know that every formal group law is either of some exact height n, or of infinite height. By Lemma A.3.3, we can restrict attention to one height at a time, and by Corollary A.3.11 we know that there is at least one such formal group law for every such height.

Lazard proved that, under the additional constraint that our field is separably closed, we have a tractable classification.

Theorem A.3.12 (Lazard's Theorem). Let k be a separably closed field of characteristic p. Then any two formal group laws of height exactly n are isomorphic. □

For a proof, see [5, Lecture 14]. In Theorem 3.2.5, we will revisit this result, where we will find that it is manifestly geometric.

A.4 Endomorphisms of formal group laws

Let $f(x, y)$ be a formal group law over a ring R. Recall that an endomorphism of $f(x, y)$ is a morphism $h\colon f(x, y) \to f(x, y)$ of formal group laws. The set of endomorphisms becomes a group under the operation $(h_1, h_2) \mapsto f(h_1, h_2)$ (see Lemma A.3.2), and it becomes a (possibly non-commutative) ring under composition. Endomorphisms of formal group laws are a complicated matter, but they have been well studied. For instance, the work of Dieudonné and Lubin, among others, has resulted in the following theorem. See also [9, Thm. A2.2.18] and the reference therein.

Theorem A.4.1. Let $f(x, y)$ be a formal group law of height d over a perfect field k of characteristic p containing $\mathbb{F}_{p^d}$, and assume that $[p]_f(t) = t^{p^d}$ (or, in more modern language, take $f(x, y)$ to be the Honda formal group law of height d). Then we have an isomorphism of non-commutative rings

$$\operatorname{End} f(x, y) \cong W(\mathbb{F}_{p^d})\langle x\rangle/(x^d - p, xw - \operatorname{Frob}(w)x)$$

where $W(\mathbb{F}_{p^d})$ is the ring of Witt vectors, w ranges over the Witt vectors in $W(\mathbb{F}_{p^d})$, and $\operatorname{Frob}(\,\cdot\,)$ is a fixed lift of the Frobenius endomorphism of $\mathbb{F}_{p^d}$ to $W(\mathbb{F}_{p^d})$. □

Let us consider the special case of $d = 1$. By Example A.3.4, the multiplicative formal group law $f(x, y)$ fits within the framework for the above theorem. Viewing $f(x, y)$ as a formal group law in $\mathbb{F}_p$, the above theorem tells us that the endomorphism ring of the multiplicative formal group law must be isomorphic to the ring $\mathbb{Z}_p$; in particular, the automorphism group is given by the unit group $\mathbb{Z}_p^*$. The main goal in this section is to give an elementary proof of this fact. For convenience, we state it as a separate result.

Theorem A.4.2. The endomorphism ring of the multiplicative formal group law over the ring $\mathbb{F}_p$ is isomorphic to $\mathbb{Z}_p$, hence the automorphism group is the unit group $\mathbb{Z}_p^*$. □

Proof: We begin by establishing a map $\mathbb{Z}_p \to \operatorname{End} f(x, y)$. The map is described as follows. Let n be a p-adic integer, and express it as a sum $n = a_0 + a_1 p + a_2 p^2 + \cdots$. Let's now define the formal power series

$$h_n(t) \coloneqq (1 + t)^n - 1\,.$$

If n is an integer, or equivalently if the series $a_0 + a_1 p + a_2 p^2 + \cdots$ terminates after finitely many steps, then this power series is just a polynomial. If not, we should interpret the formal power series as

$$h_n(t) = (1 + t)^{a_0}(1 + t)^{a_1 p}(1 + t)^{a_2 p^2} \cdots - 1\,,$$

which can be expanded in the usual way, yielding possibly infinitely many well-defined coefficients. Let's take a closer look at these coefficients. As we are working over characteristic p, we may as well write

$$h_n(t) = (1+t)^{a_0}(1+t^p)^{a_1}(1+t^{p^2})^{a_2}\cdots - 1\,.$$

We can now make the following observation. As the a_i are strictly less than p, the coefficient in front of t^{p^k} is given by a_k, and while the coefficient in front of t^c, $p^k < c < p^{k+1}$, is a finite expression in the variables $a_0, a_1, \ldots, a_k$. In particular, we find that any endomorphism coming from our map is uniquely determined by, and uniquely determines, the coefficients in front of the t_{p^k}, for $k = 0, 1, \ldots$.

We first verify that $h_n(t)$ defines an endomorphism of the multiplicative formal group law. To verify this, we should check that

$$h_n(x+y+xy) = h_n(x) + h_n(y) + h_n(x)h_n(y)\,.$$

In order to prove this identity, we need to check that the coefficients in front of $x^k y^\ell$ coincide, for all finite k and ℓ. But by the observation above, for every pair (k, ℓ), the coefficient in front of $x^k y^\ell$ depends on only finitely many of the a_i. Thus, to prove that the coefficients in front of $x^k y^\ell$ coincide, we may as well throw away the other a_i, effectively replacing n by an integer. This tells us that we are done once we have verified the identity when n is an integer. To do this, just write

$$\begin{aligned}
h_n(x+y+xy) &= (1+x+y+xy)^n - 1 \\
&= (1+x)^n(1+y)^n - 1 \\
&= ((1+x)^n - 1) + ((1+y)^n - 1) + ((1+x)^n - 1)((1+y)^n - 1) \\
&= h_n(x) + h_n(y) + h_n(x)h_n(y)
\end{aligned}$$

as was to be shown.

We now verify that the map $\mathbb{Z}_p \to \operatorname{End} f(x,y)$ is a homomorphism of rings. Let us write $n = a_0 + a_1 p + a_2 p^2 + \cdots$, and $m = b_0 + b_1 p + b_2 p^2 + \cdots$. Then we get

$$\begin{aligned}
f(h_n(t), h_m(t)) &= h_n(t) + h_m(t) + h_n(t)h_m(t) \\
&= ((1+t)^n - 1) + ((1+t)^m - 1) + ((1+t)^n - 1)((1+t)^m - 1) \\
&= (1+t)^n(1+t)^m - 1 \\
&= (1+t)^{a_0+b_0}(1+t)^{(a_1+b_1)p}(1+t)^{(a_2+b_2)p^2}\cdots - 1
\end{aligned}$$

and

$$\begin{aligned}
h_n(h_m(t)) &= \Big(1 + ((1+t)^m - 1)\Big)^n - 1 \\
&= \Big((1+t)^{b_0}(1+t)^{b_1 p}\cdots\Big)^{a_0}\Big((1+t)^{b_0}(1+t)^{b_1 p}\cdots\Big)^{a_1 p}\cdots - 1 \\
&= \Big((1+t)^{b_0}(1+t^p)^{b_1}\cdots\Big)^{a_0}\Big((1+t^p)^{b_0}(1+t^{p^2})^{b_1}\cdots\Big)^{a_1}\cdots - 1 \\
&= (1+t)^{a_0 b_0}(1+t^p)^{a_0 b_1 + a_1 b_0}(1+t^{p^2})^{a_0 b_2 + a_1 b_1 + a_2 b_0}\cdots - 1
\end{aligned}$$

which shows the desired result.

The final step is to verify that any endomorphism is in the image of our map. To this end we start with an arbitrary endomorphism $h(t) = h_1 t + h_2 t^2 + \cdots$. Expand the equation

$$h(x + y + xy) - h(x) - h(y) - h(x)h(y) = \sum_{i,j\geq 0} c_{ij} x^i y^j . \qquad (*)$$

All of the c_{ij} are polynomial expressions in the h_k. In fact, with some elementary combinatorics, one can verify that

$$c_{ij} = \sum_{k=0}^{\min(i,j)} \left(\frac{(i+j-k)!}{(i-k)!(j-k)!k!} h_{i+j-k} \right) - h_i h_j .$$

Notice in particular that $c_{ij} = 0$ if either $i = 0$ or $j = 0$ — something which is easily seen directly from the functional equations. Our goal is to consider special cases of (∗) to show that, if we know the values h_1, h_p, h_{p^2}, h_{p^3}, and so on, then we can uniquely determine the values of all other h_k. By our previous considerations, we conversely know that any choice of values h_1, h_p, h_{p^2}, and so on, actually yields a unique endomorphism.

Consider first the special case of (∗) where $j = 1$. If $i \geq 1$, we find that

$$c_{i1} = (i + 1)h_{i+1} + ih_i - h_i h_1 .$$

Filling in $i = 1$, we find that h_2 is determined by h_1; filling in $i = 2$, h_3 ends up being determined by h_1 and h_2. This pattern continues until $i = p - 1$, in which case the characteristic of our field forces the equation to become trivial, thus yielding no relation between h_p and $h_1, \ldots, h_{p-1}$. Going on, we find that $h_{p+1}, \ldots, h_{2p-1}$ are determined by $h_1, \ldots, h_p$, but when filling in $i = 2p-1$, we find no relation between h_{2p} and $h_1, \ldots, h_{2p-1}$. This pattern continues, and we find that it suffices to determine the values of h_1, h_p, h_{2p}, h_{3p}, and so on.

At this point, we consider the special case of (∗) where $j = p$, and $i = np$ for some $n \geq 1$. We find the equation

$$\begin{aligned} c_{ij} &= \frac{(np+p-0)!}{(np-0)!(p-0)!0!} h_{np+p-0} + \frac{(np+p-1)!}{(np-1)!(p-1)!1!} h_{np+p-1} + \cdots \\ &\quad + \frac{(np+p-(p-1))!}{(np-(p-1))!(p-(p-1))!(p-1)!} h_{np-(p-1)} + \frac{(np+p-p)!}{(np-p)!(p-p)!p!} h_{np+p-p} - h_{np} h_p \\ &= \frac{(np+p)!}{(np)!p!} h_{np+p} + \frac{(np+p-1)!}{(np-1)!(p-1)!1!} h_{np+p-1} + \cdots \\ &\quad + \frac{(np+1)!}{(np-(p-1))!1!(p-1)!} h_{np+1} + \frac{(np)!}{(np-p)!p!} h_{np} - h_{np} h_p \end{aligned}$$

Barring the potential vanishing of coefficients for a moment, we notice that, if $n = 1$, we find a relation between h_{2p} and h_p, thus telling us that h_{2p} is determined by h_p; if $n = 2$, we find that h_{3p} is determined by h_{2p}, and so on. The only situation in which this reasoning could fail, is when the coefficient in fron of $h_{(n+1)p}$ in the equation in the case $n = p$ is divisible by p. Let's take a look at this coefficient. We write

$$\frac{(np+p)!}{(np)!p!} = \frac{(np+1)(np+2)\cdots(np+p)}{p!}.$$

The denominator is a multiple of p, but not of p^2. As for the numerator, the only term divisible by p is $np+p$. Thus, the coefficient vanishes modulo p if and only if $np+p$ is a multiple of p^2. That is, when $n-1$ is a multiple of p. Thus we now find that $h(t)$ is entirely determined by the values of h_1, h_p, h_{p^2}, h_{2p^2}, h_{3p^2}, and so on.

At the next step, the pattern will hopefully become clear. Look at (∗) in the case $j = p^2$ and $i = np^2$ for some $n \geq 1$. We find the equation

$$\begin{aligned}
c_{ij} = {} & \frac{(np^2+p^2-0)!}{(np^2-0)!(p^2-0)!0!} h_{np^2+p^2-0} + \frac{(np^2+p^2-1)!}{(np^2-1)!(p^2-1)!1!} h_{np^2+p^21} + \cdots \\
& + \frac{(np^2+p^2-(p^2-1))!}{(np^2+p^2-(p^2-1))!(p^2-(p^2-1))!(p^2-1)!} h_{np^2+p^2-(p^2-1)} \\
& + \frac{(np^2+p^2-p^2)!}{(np^2-p^2)!(p^2-p^2)!(p^2)!} h_{np^2+p^2-p^2} - h_{np^2} h_{p^2} \\
= {} & \frac{(np^2+p^2)!}{(np^2)!(p^2)!} h_{np^2+p^2} + \frac{(np^2+p^2-1)!}{(np^2-1)!(p^2-1)!1!} h_{np^2+p^2-1} + \cdots \\
& + \frac{(np^2+1)!}{(np^2+1)!(p^2+1)!(p^2-1)!} h_{np^2+1} + \frac{(np^2)!}{(np^2-p^2)!(p^2)!} h_{np^2} - h_{np^2} h_{p^2}
\end{aligned}$$

This equation shows that $h_{(n+1)p^2}$ is determined by h_{np^2}, except in the case where the coefficient in front of $h_{(n+1)p^2}$ in the equation in the case $i = np^2$ is divisible by p. The coefficient is

$$\frac{(np^2+p^2)!}{(np^2)!(p^2)!} = \frac{(np^2+1)(np^2+2)\cdots(np^2+p^2)}{(p^2)!}.$$

The denominator is divisible by p^{p+1}, but not by p^{p+2}. On the other hand, the numerator can only be divisible by p^{p+2} if np^2+p^2 is divisible by p^3, which happens only when $n = p-1$. Thus we find that $h(t)$ is uniquely determined by the values h_1, h_p, h_{p^2}, h_{p^3}, h_{2p^3}, h_{3p^3}, and so on.

This procedure continues indefinitely, with the same general pattern. We consider the case of (∗) where $j = p^3$, $i = np^3$ for some $n \geq 1$, then the case where $j = p^3$, $i = np^4$, and so on. The conclusion is as desired: our equation is uniquely by a choice of h_1, h_p, h_{p^2}, h_{p^3}, and so on, and by the existence of the map $\mathbb{Z}_p \to \operatorname{End} f(x,y)$, any such choice yields an endomorphism. ■

We will find that it will prove to be useful to also compute the truncated analogue of the above theorem. In Section 3.1 we'll define a k-bud to be essentially a formal group law $f(x,y)$, but defined only modulo $(x,y)^{k+1}$; likewise, a morphism of k-buds is described by a series $h(t)$ which is defined modulo (t^{k+1}). For all $k \geq 1$, we may interpret the multiplicative formal group law $f(x,y) = x+y+xy$ as a k-bud, so that we may ask what its endomorphisms are.

Theorem A.4.3. Let the notation be as above. Write r for $\lfloor \log_p(k) \rfloor$. Then the endomorphism ring of the multiplicative k-bud over the ring $\mathbb{F}_p$ is isomorphic to $\mathbb{Z}/p^r\mathbb{Z}$, hence the automorphism group is $(\mathbb{Z}/p^r\mathbb{Z})^*$. □

Proof sketch: On the other hand, every automorphism of $f(x, y)$, when viewed as a formal group law, truncates to an automorphism of k-buds. This allows us to find a map $\mathbb{Z}_p \to \operatorname{End}(f(x, y))$, where the endomorphism ring is the ring of endomorphisms of k-buds. In fact, as k-buds are determined by finitely many coefficients, and each of the coefficients in the power series $h(t) = (1 + t)^n - 1$, $n \in \mathbb{Z}_p$, are determined by finitely many of the coefficients in the decomposition $n = a_0 + a_1 p + a_2 p^2 + \cdots$, it's perhaps better to say that we have a map $\mathbb{Z}_{(p)} \to \operatorname{End}(f(x, y))$.

But more is true than that. Looking at the second part of the proof of the previous theorem, we find that any endomorphism of the formal group law $f(x, y) = x + y + xy$ was determined by the coefficients h_1, h_p, h_{p^2}, and so on. This carries over to the case of k-buds, except this time, we have only finitely many coefficients to deal with, thanks to the truncation. In fact, we have precisely $\lfloor \log_p(k) \rfloor$ of them. As every endomorphism of k-buds is determined by $\lfloor \log_p(k) \rfloor$ choices of coefficients in $\mathbb{F}_p$, and conversely all possible choices are known to actually give an endomorphism, there are no further possible choices, and the desired result thus follows. ■

Bibliography

[1] Alridge Bousfield. The localization of spectra with respect to homology. *Topology*, 18:257–281, 1979.

[2] Alridge Bousfield. On the homotopy theory of K-local spectra at an odd prime. *American Journal of Mathematics*, 107(4):895–932, 1985.

[3] Alridge Bousfield. A classification of K-local spectra. *Journal of Pure and Applied Algebra*, 66(2):121–163, 1990.

[4] Paul G. Goerss. Quasi-coherent shaves on the moduli stack of formal groups. *arXiv preprint arXiv:0802.0996*, 2008.

[5] Jacob Lurie. Chromatic Homotopy Theory. Lecture notes for the 2010 Harvard course Math 252x; available at http://www.math.harvard.edu/~lurie/.

[6] Cary Malkiewich. The stable homotopy category. Expository notes; available at https://faculty.math.illinois.edu/~cmalkiew/.

[7] Lennart Meier. Elliptic Homology and Topological Modular Forms. Lecture notes; version 0.12; available at https://www.staff.science.uu.nl/~meier007/.

[8] Niko Naumann. The stack of formal groups in homotopy theory. *Advances in Mathematics*, 215(2):569–600, 2007.

[9] Douglas Ravanel. *Complex cobordism and stable homotopy groups of spheres*. American Mathematical Society, 2003.

[10] Michael Schlessinger. Functors of Artin rings. *Transactions of the American Mathematical Society*, 130(2):208–222, 1968.

[11] The Stacks Project Authors. Stacks project. Open-source textbook; currently in progress; available at https://stacks.math.columbia.edu/.

[12] Neil Strickland. An introduction to the category of spectra. Expository notes; available at http://neil-strickland.staff.shef.ac.uk/ under 'Research'.

[13] Robert M. Switzer. *Algebraic Topology — Homotopy and Homology*, volume 212 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1975.

[14] Angelo Vistoli. Notes on Grothendieck topologies, fibered categories and descent theory. *arXiv preprint math/0412512*, 2004.

[15] Qiaochu Yuan. Operations, pro-objects, and grothendieck's galois theory. Blog post; available at https://qchu.wordpress.com/2013/06/23/operations-pro-objects-and-grothendiecks-galois-theory/.

Index