These chapters are designed to present the required tools as clearly as possible. They are not intended to be rigorously self contained. Whenever a lengthy proof is available elsewhere in the literature, we will omit it but tell the reader exactly where she can find it. They are also not intended to be comprehensive introductions to the topics in question. Our choice of definitions and results stated, which may strike some readers as idiosyncratic, is dictated by the needs of the latter chapters. We have chosen to ignore some recent developments in these areas, such as the theory of ∞-categories, because we do not need them. On the other hand we have chosen to embrace enriched category theory, the subject of Chapter 3, since it provides the cleanest framework for the definition of equivariant spectra in §9.1.

Equivariant homotopy theory, the arena in which our computation is done, first appears in Chapter 7, and our star player, the spectrum $MU_{\mathbf{R}}$, is constructed in Chapter 12. Our main computational tool, the slice spectral sequence, first appears in Chapter 11.

The inexperienced reader may well wonder why we need to devote over two hundred pages to category theory before we even step into the pool of homotopy theory. The answer is that the tools it provides enable us to proceed with far more elegance and rigor than we could without them. This "categorification of algebraic topology" is most apparent in the twenty-first century approach to **spectra**, the fundamental objects of study in stable homotopy theory.

Spectra were first introduced in print in 1959 by Lima (then a student of Spanier at the University of Chicago and later a prominent mathematical educator in Brazil) in [Lim59]. A spectrum E was defined to be a sequence of spaces E_n for nonnegative integers n, with structure maps

$$\epsilon_n : \Sigma E_n \to E_{n+1}.$$

In the first examples E_n was $(n-1)$-connected, but this was not a formal requirement. The motivation for this definition was the observation that **$(n-1)$-connected spaces behave very nicely in dimensions less than roughly $2n$**. The first theorem in this direction may have been the Freudenthal Suspension Theorem [Fre38] of 1938; see [Rav86, Theorem 1.1.4].

Spectra were defined to create a world where n could be arbitrarily large so we could enjoy this nice behavior in **all** dimensions. Perhaps the first extensive account of this new world was a course given by Adams at UC Berkeley in 1961 and published as [Ada64]. In it (pages 22–23), he said the following.

> I want to go ahead and construct a stable category. Now I should warn you that the proper definitions here are still a matter for much pleasurable argumentation among the experts. The debate is between two attitudes, which I'll personify as the tortoise and the hare. The hare is an idealist: his preferred position is one of elegant and all embracing generality. He wants to build a new heaven and a new earth and no half-measures. If he had to construct the real numbers he'd begin by taking **all** sequences of rationals, and only introduce that tiresome condition about convergence when he was absolutely forced to.

> The tortoise, on the other hand, takes a much more restrictive view. He says that his modest aim is to make a cleaner statement of known theorems, and he'd like to put a lot of restrictions on his stable objects so as to be sure that his category has all the good properties he may need. Of course, the tortoise tends to put on more restrictions than are necessary, but the truth is that the restrictions give him confidence.
>
> You can decide which side you're on by contemplating the Spanier-Whitehead dual of an Eilenberg-Mac Lane object. This is a "complex" with "cells" in all stable dimensions from $-\infty$ to $-n$. According to the hare, Eilenberg-Mac Lane objects are good, Spanier-Whitehead duality is good, therefore this is a good object: And if the negative dimensions worry you, he leaves you to decide whether you are a tortoise or a chicken. According to the tortoise, on the other hand, the first theorem in stable homotopy theory is the Hurewicz Isomorphism Theorem, and this object has no dimension at all where that theorem is applicable, and he doesn't mind the hare introducing this object as long as he is allowed to exclude it. Take the nasty thing away!

The resulting homotopy theoretical paradise was described very nicely by Boardman-Vogt in [BV73] about a decade later, but there were some serious technical problems, especially in connection with smash products. For a further account of the adventures of the hare and the tortoise with an assessment of Boardman's work, see [May99b].

It is safe to say now, over half a century later, that **the hare has prevailed**. The technical problems that vexed stable homotopy theorists for a generation have been vanquished. The increasingly sophisticated use of category theory has been instrumental in this triumph. Many of the advances that led to this happy state of affairs occured in the 1990s, the decade following Adams' untimely death in a car crash. The third author has tried to imagine what it would be like to relate these developments to him.

> Dear Frank,
>
> Stable homotopy theory is in much better shape now than when you left us. The definitions are much cleaner and we have a smash product with all of the nice features you could ask for. As you can probably guess, Peter May has been pounding away at this for decades, but you did not live long enough to see just how much success he and his coauthors have had.
>
> Along with his former student Tony Elmendorf and Igor Kriz, a Czech immigrant (you may also be interested to know that the Berlin Wall came down, the Soviet Union collapsed and the Cold War ended, all within three years of your death), he found a definition of the stable homotopy category that featured a smash product that is **strictly** associative and commutative in 1993. You heard me

right, I said strictly, not just up to homotopy (higher or otherwise) or some other convoluted equivalence relation, but pointwise, on the nose! In 1997 (with a fourth coauthor, Mike Mandell, another former student) they published a book about it, [**EKMM97**].

The construction is complicated and I do not fully understand it. Fortunately May and Mandell found a simpler way to do it a few years later, described in another book, [**MM02**] published in 2002. This one I do understand. It uses a wonderful construction called the **Day convolution**, originally discovered in 1970 [**Day70**] by the Australian category theorist Brian Day (1945-2012). It is a purely categorical result that happens to be exactly what is needed to define the smash product of spectra. This means the proof that said smash product is strictly commutative and associative is "purely formal." Ironically, Day's first job out of graduate school was a postdoctoral position at the University of Chicago, presumably at the behest of Saunders Mac Lane. As far as I can tell, Brian and Peter did not interact mathematically.

So how do Mandell and May do it? As you know, a spectrum E was originally defined to be a sequence of pointed spaces E_n, one for each integer $n \geqslant 0$, along with pointed structure maps $\epsilon_n : \Sigma E_n \to E_{n+1}$. For them **a spectrum is a functor** from a certain small category $\mathscr{J}$ (the Mandell-May category of Definition 7.9.20) to the category $\mathcal{T}$ of pointed topological spaces. Since $\mathscr{J}$ is small, such a functor could be regarded as a **diagram of pointed spaces**, although it would not be a diagram you could actually draw because it would be infinite. This point of view is developed further in the companion paper to [**MM02**], [**MMSS01**].

The objects of $\mathscr{J}$ are finite dimensional real orthogonal vector spaces. Since such a vector space is determined up to isomorphism by its dimension, a $\mathcal{T}$-valued functor E on $\mathscr{J}$ gives us a sequence of spaces E_n, as in the orginal definition, but with some additional structure. In order to spell out the additional structure, I need to tell you about the morphisms in $\mathscr{J}$. **This is where things start to get tricky.**

I said the objects of $\mathscr{J}$ are certain vector spaces, but I did not say that $\mathscr{J}$ is **the category** of such vector spaces and inner product preserving maps as usually defined. In order to describe $\mathscr{J}$ we need to generalize what we mean by a category because $\mathscr{J}$ is not a category in the usual sense. Instead it is an **enriched category**; see Chapter 3. Such things were first studied by Eilenberg and Kelly [**EK66**] and were the subject of Kelly's book [**Kel82**].

In an ordinary category C one has a collection (possibly a set) of objects, and for each pair of objects X and Y a set $C(X, Y)$ (possibly empty) of morphisms $X \to Y$. Of course every object has an identity morphism, and given a third object Z we have a map

$$C(Y,Z) \times C(X,Y) \to C(X,Z) \tag{1.1}$$

that tells us how to compose morphisms. This map is itself a morphism with suitable properties in $\mathcal{S}et$, the category of sets.

In an enriched category, one has objects as before, but $C(X,Y)$ is no longer a set or even a class. Instead it is an object in a second category $\mathcal{V}$, which need not be $\mathcal{S}et$ at all. We say then that C is **enriched over** $\mathcal{V}$. By this definition, an ordinary category is enriched over $\mathcal{S}et$.

This auxiliary category $\mathcal{V}$ has to have a structure that enables to make sense of the source of the morphism in (1.1). In other words it needs a binary operation, analogous to Cartesian product in $\mathcal{S}et$, that allows us to combine two objects into a third. This binary operation must have a unit anolgous to the one element set. A category so endowed is said to be **symmetric monoidal**; see §2.6 for more information. The relevant example for us is $\mathcal{T}$, the category of pointed topological spaces. Its binary operation is the smash product, for which the unit object is S^0.

Having said what an enriched category is, I can tell you more about the Mandell-May category $\mathscr{J}$, which is enriched over $\mathcal{T}$. This means that for finite dimensional real orthogonal vector spaces V and W, the morphism object $\mathscr{J}(V,W)$ is a pointed topological space, which is defined as follows.

Let $O(V,W)$ denote the (possibly empty) space (also known as a Stiefel manifold) of orthogonal embeddings of V into W. For each such embedding τ, let $W-\tau(V)$ denote the orthogonal compliment of $\tau(V)$ in W. We can regard it as the fiber of a vector bundle over $O(V,W)$, and **we define $\mathscr{J}(V,W)$ to be its Thom space**.

When the dimension of V exceeds that of W, the embedding space $O(V,W)$ is empty, which means the Thom space $\mathscr{J}(V,W)$ is a point. When V and W have the same dimension, the vector bundle has zero dimensional fibers, so $\mathscr{J}(V,W) = O(V)_+$, the orthogonal group with a disjoint base point. When the dimension of W exceeds that of V, we can think of $\mathscr{J}(V,W)$ as a wedge of copies of $S^{W-\tau(V)}$ parametrized by the space of embeddings $O(V,W)$.

Given a third such vector space U, the analog of (1.1) is a suitable map

$$\mathscr{J}(V,W) \wedge \mathscr{J}(U,V) \to \mathscr{J}(U,W). \tag{1.2}$$

It is induced by composition of orthogonal embeddings, i.e., by a map

$$O(V,W) \times O(U,V) \to O(U,W).$$

It does not help to think of points in $\mathscr{J}(V,W)$ as maps from V to W. The space $\mathscr{J}(V,W)$ is not a topologized set of ordinary morphisms, but a replacement of the usual morphism set by a morphism object in $\mathcal{T}$. The map of (1.2) tells us how the replacement of composition works.

Mandell-May define an **orthogonal spectrum** E (Definition 9.1.2) to be a functor from $\mathscr{J}$ to $\mathcal{T}$, which happens to be enriched over itself. Since an object of $\mathscr{J}$ is a finite dimensional vector space, which is determined up to isomorphism by its dimension, we denote the image of the functor on $\mathbf{R}^n$ by E_n as in the orginal definition. Functoriality implies that we have structure maps

$$\epsilon_{n,n+k} : \mathscr{J}(\mathbf{R}^n, \mathbf{R}^{n+k}) \wedge E_n \to E_{n+k}. \tag{1.3}$$

for all $n, k \geqslant 0$.

For $k = 0$ this amounts to a left action on the space E_n of the orthogonal group $O(n)$. That group also acts on $\mathscr{J}(\mathbf{R}^n, \mathbf{R}^{n+k})$ on the right by precomposition. These two actions lead to one on the smash product in (1.3) with $\epsilon_{n,n+k}$ factoring through the orbit space. For $k = 1$ that orbit space is ΣE, so **we have the map** $\epsilon_n = \epsilon_{n,n+1} : \Sigma E_n \to E_{n+1}$ **as in the orginal definition**. The difference is that now the map does not depend on the choice of orthogonal embedding of $\mathbf{R}^n$ into $\mathbf{R}^{n+1}$ as it did in the classical case. **This coordinate free definition is technically convenient.**

We can define the suspension spectra $\Sigma^\infty X$ for a pointed space X by $(\Sigma^\infty X)_n = \Sigma^n X$ with the evident structure maps. More generally we can define the smash product of a pointed space K with a spectrum E by $(K \wedge E)_n = K \wedge E_n$. We can also define a spectrum E^K (maps from K to E) by

$$(E^K)_n = \mathcal{T}(K, E_n).$$

Since spectra are functors, maps between them are natural transformations. This means a map $f : E \to F$ of spectra is a collection of continuous pointed maps $f_n : E_n \to F_n$ compatible with the structure maps. This is analogous to what you called a **function** in [Ada74b, page 140].

As you pointed out on [Ada74b, page 141], there is no function $f : \Sigma^\infty S^1 \to \Sigma^\infty S^0$ for which $f_2 : S^3 \to S^2$ is the Hopf map η. Since we all love the Hopf map, we would like to have such a function. The fix you suggested is to replace the source spectrum $E = \Sigma^\infty S^1$ by a spectrum E' defined by

$$E'_n = \begin{cases} * & \text{for } n = 0, 1 \\ S^{n+1} & \text{otherwise} \end{cases}$$

Then there is an obvious function $g : E' \to E$ for which g_n is an isomorphism for $n \geqslant 2$, and a function $f' : E' \to \Sigma^\infty S^0$ with

$f_2' = \eta$. You defined a **map** $E \to F$ ([**Ada74b**, page 142]) to be an equivalence class of composites of the form $f' = fg$ as above.

You also defined a homotopy between two functions $E \to F$ ([**Ada74b**, page 144]) in terms of a map $I_+ \wedge E \to F$, a homotopy between maps, in similar terms. Finally, you defined a **morphism** in your category ([**Ada74b**, page 143]) to be a homotopy class of such maps.

Thus you made a distinction between functions, maps and morphisms. Subsequent experience has led us to approach these issues a little differently. We have learned that the framework provided by Quillen's theory of model categories, the subject of Chapters 4–6 of this book, is very helpful. Among other things, it tells us there are two categories one should consider here. The first is the category of spectra $\mathcal{S}p$ in which the objects are the functors $\mathscr{J} \to \mathcal{T}$ described above, and the morphisms are natural transformations between them, what you called "functions."

Before describing the second category, we need to define stable homotopy groups of spectra and weak equivalences between spectra. This can be done as you did in [**Ada74b**, §III.3]. Then one gets a **homotopy category** $\mathrm{Ho}\,\mathcal{S}p$ (see Definition 4.3.16) having the same objects as $\mathcal{S}p$ in which weak equivalences are invertible. Your "morphisms" are morphisms in this category. Your "maps" are equivalence classes of "functions" precomposed with weak equivalences.

9/14/18. We may need to update some forward references here, such as that to Definition 9.2.41.

Now, at last, I can tell you about smash products. You defined the smash product of two spectra in [**Ada74b**, §III.4] and spent 30 pages showing that it has the desired properties (commutativity and associativity with the sphere spectrum as unit) **up to homotopy**, that is up to coherent natural weak equivalence. Another way of saying this is that we get a symmetric monoidal structure in the homotopy category $\mathrm{Ho}\,\mathcal{S}p$. The Mandell-May smash product (Definition 9.2.41), which is based on a very insightful observation by Jeff Smith, leads to such a structure in $\mathcal{S}p$ itself. This smash product has the desired properties up to coherent natural weak **isomorphism**. Not only is this a huge improvement, it has a much shorter and more elegant proof.

If we have two spectra X and Y, each of which is a functor $\mathscr{J} \to \mathcal{T}$, then together they give us a functor from $\mathscr{J} \times \mathscr{J}$ to

$\mathcal{T} \times \mathcal{T}$. Now consider the diagram

(1.4)

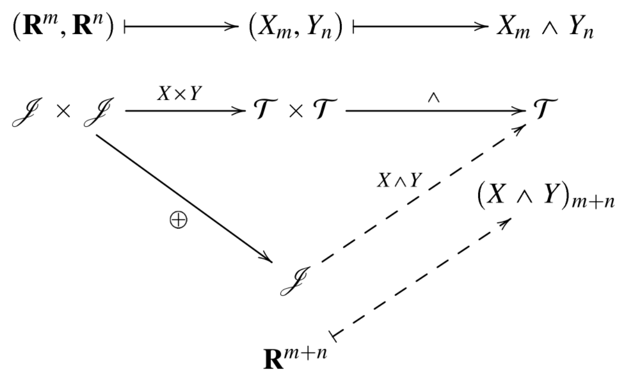

The smash product we are looking for is a yet to be defined functor

$$X \wedge Y : \mathscr{J} \to \mathcal{T}$$

with suitable properties. We are **not** hoping for the diagram to commute. That would mean

$$X_m \wedge Y_n \cong (X \wedge Y)_{m+n}$$

in all cases, which is not a reasonable thing to expect. On the other hand, we do expect to have maps

$$\eta_{m,n} : X_m \wedge Y_n \to (X \wedge Y)_{m+n}. \tag{1.5}$$

They should be induced by a natural transformation η from the composite functor $\mathscr{J} \times \mathscr{J} \to \mathcal{T}$ on the top of the triangle in (1.4) to the one on the bottom.

It turns out that the right way to define $X \wedge Y$ involves a universal property of this natural transformation. In order to state it, we will replace (1.4) with the following diagram, which will be discussed further in §2.5. Suppose we have categories $\mathcal{C}$, $\mathcal{D}$ and $\mathcal{E}$, with functors F and K as in

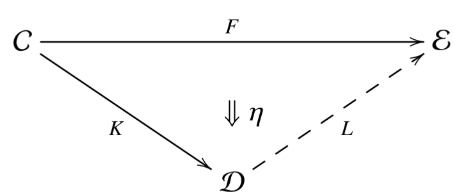

We wish to extend the functor F along K to a new functor

$$L : \mathcal{D} \to \mathcal{E}$$

with a natural transformation $\eta : F \Rightarrow LK$. The composite functor LK need not be the same as F. Instead we want L and η to have the following universal property: given another such extension G with a natural transformation $\gamma : F \Rightarrow GK$ as in the diagram

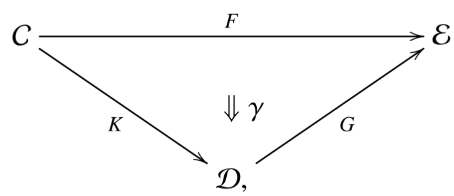

there is a unique natural transformation $\alpha : L \Rightarrow G$ with

$$\gamma = (\alpha K)\eta$$

as in the following diagram.

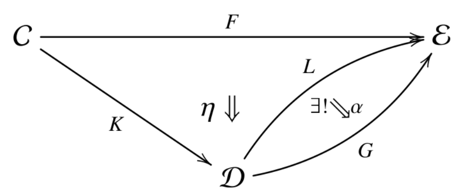

If such an L exists, it is unique and is called the **left Kan extension of F along K**. It is so named because such functors were first studied by Dan Kan in [Kan58]; see §2.5. The bottom line is that such an L exists when the categories $\mathcal{C}$ and $\mathcal{D}$ are small and the category $\mathcal{E}$ is closed under colimits. These conditions are met by the categories of (1.4).

There is also an explicit formula for L under these conditions that is described below in § 2.5B. In the case at hand, where $L = X \wedge Y$, it is as follows. Define pointed spaces

$$W_n = \bigvee_{0 \leqslant i \leqslant n} X_i \wedge Y_{n-i}$$

and

$$W'_n = \bigvee_{0 \leqslant i \leqslant n-1} X_i \wedge S^1 \wedge Y_{n-1-i}.$$

Then the maps $\eta_{i,n-i}$ of (1.5) determine a map

$$W_n \to (X \wedge Y)_n.$$

The two maps

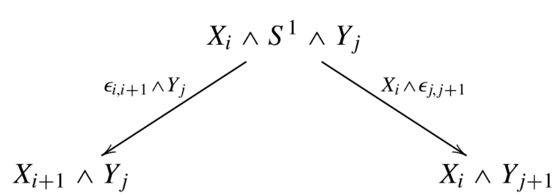

lead to two maps $\alpha, \beta : W'_n \to W_n$. **Then $(X \wedge Y)_n$ is the coequalizer of these two maps**, meaning the quotient of the space

W_n obtained by identifying the two images of W_n' with each other. **This is similar in spirit but not identical to the double telescope you described in** [Ada74b, pages 173–176]**.**

How do we know that this smash product has the properties advertized? This is the subject of the Day Convolution Theorem 3.3.5. Suppose that $\mathcal{D}$ is a small symmetric monoidal category (such as $\mathscr{J}$) enriched over a cocomplete (Definition 2.3.24) closed symmetric monoidal category (Definition 2.6.37) $\mathcal{V}$ such as $\mathcal{T}$. Then we can define a binary operation on the category $[\mathcal{D}, \mathcal{V}]$ (Definition 3.2.15) of functors $\mathcal{D} \to \mathcal{V}$ (the category $\mathcal{S}p$ of orthogonal spectra in our case) using a left Kan extension as in (1.4). The theorem says that this binary operation makes the functor category itself a closed symmetric monoidal category. Its unit is defined in a certain way in terms of the unit objects of $\mathcal{D}$ and $\mathcal{V}$. In the present case this unit is the sphere spectrum as expected.

I hope you agree this is a big improvement over the state of affairs of forty years ago.

In closing I have two additional comments for you.

(i) It is not difficult to adapt this setup to the equivariant case. This is the main point of [MM02]. For a finite group G, let $\mathcal{T}_G$ be the category of pointed G-spaces and continuous (but not necessarily equivariant) pointed maps. Then the mapping space $\mathcal{T}_G(X, Y)$ has a pointed G-action of its own, for which the fixed point set, $\mathcal{T}_G(X, Y)^G$, is the space of all equivariant maps. Hence $\mathcal{T}_G$ is enriched over itself. See Chapter 7 for more discussion.

However, if we want to do homotopy theory, we must limit ourselves to equivariant maps. The reason is that the fiber or cofiber of a map between G-spaces has a well defined G-action only when the map is equivariant. We denote the corresponding category, which is enriched over $\mathcal{T}$, by $\mathcal{T}^G$.

The category $\mathscr{J}_G$ (Definition 7.9.20) has finite dimensional orthogonal representations V of G as objects. The morphism space $\mathscr{J}_G(V, W)$ is the same Thom space as in the nonequivariant case, but now it has a G-action based on the ones on V and W. Hence $\mathscr{J}_G$ is enriched over $\mathcal{T}_G$. **We define a G-spectrum E to be an enriched functor** $\mathscr{J}_G \to \mathcal{T}_G$, and we denote the image of V by E_V and the resulting category by $\mathcal{S}p_G$. The Day Convolution Theorem still applies, so we get a nice smash product as before.

As in the case of spaces, in order to do homotopy theory we must limit ourselves to equivariant maps. We denote the corresponding category by $\mathcal{S}p^G$.

(ii) You might worry that orthogonal spectra are rarer than spectra as originally defined since they appear to have more structure. Fortunately this is not the case. It was shown in [MMSS01] (see **??**) that every ordinary (meaning as defined by Lima) spectrum can be described as an orthogonal one with the help of a left Kan extension. Better yet, all of the computations done with ordinary spectra, in particular everything you did in [Ada64], are still valid in the new category of orthogonal spectra, as well as in various others that have been proposed and studied in recent years. Remarkably, the shifting theoretical foundations of our subject have had no impact on the calculations we actually want to do. **Computation precedes theory. Our intuition about spectra was right all along.**

Thanks for reading, and best of luck in your future travels,

Doug

We wish to thank Phil Hirschhorn, Stefan Schwede, David White, David Barnes, Peter May, Eric Peterson (who suggested the use of the symbol よ for the Yoneda embedding), John McCleary (who told me about the difference between hiragana and kanji in Japanese), Carolyn Yarnall, Yexin Qu, Yan Zou, Mingcong Zeng, Carl McTague, Irina Bobkova, Brooke Shipley, Saul Glasman, Clark Barwick, Ugur Yigit, Anthony Lanese, John Greenlees, Peter Webb, Dan Dugger, Pete Bousfield, Jiri Adamek, Enrico Vitale, Dino Lorenzini, Tomer Schlank, Constanze Roitzheim, Lars Hesselholt, Scott Balchin, Christopher Deninger, Kevin Carlson, George Raptis, Todd Trimble and . . . for helpful conversations.

****

Outline:

History of problem: Pontryagin and Kervaire-Milnor, Toda's work and odd primary case, Mahowald's EHP conjecture

Chromatic backgound and Morava stabilizer group.

Equivariant tools: 4 flavors of fixed points, slice filtration, induction and norm, Mackey functors

Define slice SS.

Construct $MU_{\mathbf{R}}$ and describe its homotopy.

Construct Φ^G and norm.

Determine slices for $MU_{\mathbf{R}}$ and its norms.

Describe slice SS for $MU_{\mathbf{R}}$ and prove periodicity.

Discuss slice SS in C_4 case