IDEALS IN TRIANGULATED CATEGORIES: PHANTOMS, GHOSTS AND SKELETA

J. DANIEL CHRISTENSEN

ABSTRACT. We begin by showing that in a triangulated category, specifying a projective class is equivalent to specifying an ideal $\mathcal{I}$ of morphisms with certain properties, and that if $\mathcal{I}$ has these properties, then so does each of its powers. We show how a projective class leads to an Adams spectral sequence and give some results on the convergence and collapsing of this spectral sequence. We use this to study various ideals. In the stable homotopy category we examine phantom maps, skeletal phantom maps, superphantom maps, and ghosts. (A ghost is a map which induces the zero map of homotopy groups.) We show that ghosts lead to a stable analogue of the Lusternik–Schnirelmann category of a space, and we calculate this stable analogue for low-dimensional real projective spaces. We also give a relation between ghosts and the Hopf and Kervaire invariant problems. In the case of A_∞ modules over an A_∞ ring spectrum, the ghost spectral sequence is a universal coefficient spectral sequence. From the phantom projective class we derive a generalized Milnor sequence for filtered diagrams of finite spectra, and from this it follows that the group of phantom maps from X to Y can always be described as a $\varprojlim^1$ group. The last two sections focus on algebraic examples. In the derived category of an abelian category we study the ideal of maps inducing the zero map of homology groups and find a natural setting for a result of Kelly on the vanishing of composites of such maps. We also explain how pure exact sequences relate to phantom maps in the derived category of a ring and give an example showing that phantoms can compose non-trivially.

CONTENTS

Date: February 25, 2013.

1991 *Mathematics Subject Classification.* Primary 55P42; Secondary 18G35, 18E30, 55U99.

Key words and phrases. Ideal, projective class, injective class, triangulated category, stable homotopy, spectrum, phantom map, superphantom, ghost map, chain complex, derived category, purity, pure extension, Brown representability, real projective space.

The author was partially supported by an NSF grant and an NSERC scholarship.

This version is more up to date than the published version. See the last page for details.

1. Introduction

Let $\mathcal{S}$ be a triangulated category, such as the stable homotopy category or the derived category of a ring. One often tries to study an object X in $\mathcal{S}$ by building it up from a class of better understood pieces. When this is done, there may be maps $X \to Y$ that aren't seen by the pieces (this is made precise below). Obvious questions arise, such as how efficiently X can be built from the given class of pieces, and how the unseen maps behave under composition. This paper presents a systematic way of studying such phenomena, namely by showing that they are captured in the notion of a "projective class". We then apply this formalism to various interesting examples.

If $\mathcal{P}$ is a collection of objects of $\mathcal{S}$, denote by $\mathcal{P}$-null the collection of all maps $X \to Y$ such that the composite $P \to X \to Y$ is zero for all objects P in $\mathcal{P}$ and all maps $P \to X$. These are the maps that the objects of $\mathcal{P}$ fail to "see". If $\mathcal{I}$ is a collection of maps of $\mathcal{S}$, denote by $\mathcal{I}$-proj the collection of all objects P such that the composite $P \to X \to Y$ is zero for all maps $X \to Y$ in $\mathcal{I}$ and all maps $P \to X$. A **projective class** is a pair $(\mathcal{P}, \mathcal{I})$ with $\mathcal{P}$-null $= \mathcal{I}$ and $\mathcal{I}$-proj $= \mathcal{P}$ such that for each object X there is a cofibre sequence $P \to X \to Y$ with P in $\mathcal{P}$ and with $X \to Y$ in $\mathcal{I}$. The objects of $\mathcal{P}$ are referred to as **projectives**.

This definition of "projective class" appears on the surface to be different from other definitions that have appeared, but in the next section we show that it is in fact equivalent. However, when working in a triangulated category, we claim that the above definition is more natural than the others. For example, the collection $\mathcal{I}$ of maps is automatically an **ideal** in $\mathcal{S}$. That is, if f and g are parallel maps in $\mathcal{I}$, then $f + g$ is in $\mathcal{I}$. And if f, g and h are composable and g is in $\mathcal{I}$, then both fg and gh are in $\mathcal{I}$. (All of our ideals will be two-sided.) Many commonly occurring ideals in fact form part of a projective class.

Our definition also has the feature that it leads to a sequence of "derived" projective classes. The powers $\mathcal{I}^n$ of the ideal $\mathcal{I}$ form a decreasing filtration of the class of maps of $\mathcal{S}$. Let $\mathcal{P}_1 = \mathcal{P}$ and inductively define $\mathcal{P}_n$ to be the class of all retracts of objects Y that sit in a cofibre sequence $X \to Y \to P$ with X in $\mathcal{P}_{n-1}$ and P in $\mathcal{P}$. In this way we get an increasing filtration of the class of objects of $\mathcal{S}$, and we have the following result.

Theorem 1.1. *For each n, the pair $(\mathcal{P}_n, \mathcal{I}^n)$ is a projective class.*

This result is a special case of the "product" operation on projective classes. Both are discussed in more detail in Section 3.

Theorem 1.1 forms the cornerstone of our investigations of various ideals in the stable homotopy category and the derived category of an abelian category, investigations which occupy us in the later sections of the paper. But before getting to these examples, we show in Section 4 how a projective class leads to an Adams spectral sequence, and we prove that this spectral sequence is conditionally convergent [5] if the ideal is closed under countable coproducts and the projective class is **generating**, *i.e.,* if for each non-zero X there is a projective P and a non-zero map $P \to X$. The **length** of an object X is the smallest n such that X is in $\mathcal{P}_n$, and is infinite if X is not in any $\mathcal{P}_n$. If X has finite length and we have a generating projective class such that $\mathcal{I}$ is closed under countable coproducts, then the Adams spectral sequence abutting to $[X, Y]$ is strongly convergent [5]. Our next observation is that when dealing with a generating projective class, an upper bound on the length of an object X is given by $1 + \operatorname{proj.dim} X$, where the projective dimension of X is the length of the shortest projective resolution of X with respect to the given projective class.

The remaining sections deal with examples of projective classes. Section 5 studies phantom maps in any category which has a generating phantom projective class (Definition 5.2). A phantom map is a map which is zero when restricted to any "finite" object (Definition 5.1). A necessary tool is the introduction of weaker variants of colimits, such as weak colimits and minimal cones. We conclude that if our category satisfies "Brown Representability" (Definition 5.3), then every object has projective dimension at most one, the composite of any two phantoms is zero, and the Adams spectral sequence degenerates into a generalized Milnor sequence.

In Section 6 we examine various types of phantom maps in the stable homotopy category. We begin with ordinary phantom maps, drawing on the abstract work done in the previous section. As a special case of the generalized Milnor sequence, we get the following result.

Theorem 1.2. *Let X be a CW-spectrum and let $\{X_\alpha\}$ be the filtered diagram of finite CW-subspectra of X. For any spectrum Y there is a short exact sequence*

$$0 \longrightarrow \varprojlim{}^1 [\Sigma X_\alpha, Y] \longrightarrow [X, Y] \longrightarrow \varprojlim [X_\alpha, Y] \longrightarrow 0.$$

The kernel consists precisely of the phantom maps. Moreover, $\varprojlim^i [\Sigma X_\alpha, Y]$ *vanishes for* $i \geq 2$.

This result is straightforward when X has finite skeleta, but is more delicate in general. It was also proved by Tetsusuke Ohkawa [36]. For further historical comments, see the paragraph following Theorem 6.1.

In the next part of Section 6 we study skeletal phantom maps, maps which are zero when restricted to any "skeleton" of the source (Definition 6.2). And in the final part we prove that "superphantoms" exist, answering a question of Margolis [33, p. 81]. A superphantom is a map which is zero when restricted to any (possibly desuspended) suspension spectrum.

We turn to the study of maps which induce the zero map in homotopy groups in Section 7. (We dub these maps "ghosts".) Here things are much more interesting in that our ideal has infinite order. If a spectrum X has length k with respect to this ideal, then any composite of k Steenrod operations in the mod 2 cohomology of X is zero. For this reason, we view the length of a spectrum with respect to the ideal of ghosts as a stable analogue of the Lusternik-Schnirelmann category of a space. In the second part of the section we focus on calculating the lengths of real projective spaces. We give upper and lower bounds on the length of $\mathbb{R}\mathrm{P}^n$ which agree for $2 \leq n \leq 19$, and we show that the filtration of the Steenrod squares is closely related to the Hopf and Kervaire invariant problems. In the third and final part of the section, we show that the Adams spectral sequence with respect to the ghost projective class in the category of A_∞ modules over an A_∞ ring spectrum is a universal coefficient spectral sequence, and we show how this gives another theoretical lower bound on the length of a spectrum.

The last two sections deal with the algebraic analogues of phantoms and ghosts. In the derived category $\mathcal{D}$ of an abelian category, a **ghost** is a map which induces zero in homology, and the ideal of ghosts is a natural ideal to study. We show that if our abelian category has enough projectives and satisfies Grothendieck's AB 5 axiom, then the ideal of ghosts is part of a generating projective class, and we prove the following result.

Theorem 1.3. *Let* X *be a complex such that the projective dimensions of* $B_n X$ *and* $H_n X$ *are less than* k *for each* n*. Then the projective dimension of* X *with respect to the ideal of ghosts is less than* k*. In particular,* X *has length at most* k*, and a* k*-fold composite*

$$X \longrightarrow Y^1 \longrightarrow \cdots \longrightarrow Y^k$$

of maps each zero in homology is zero in $\mathcal{D}$*.*

One can deduce from this a result of Max Kelly [26], whose work provided the inspiration for the general framework presented in this paper. See Section 8 for details.

In the last section we discuss phantom maps in the derived category of an associative ring. A map $X \to Y$ is phantom if and only if the composite $W \to X \to Y$ is zero (in the derived category) for each bounded complex W of finitely generated projectives. Our main result here is a relation between phantom maps and pure extensions of R-modules as well as an analog of Theorem 1.3. We reproduce an example which shows that when $R = \mathbb{C}[x, y]$, Brown representability doesn't hold in the derived category.

I have had the pleasure of discussing this work with many mathematicians, to whom I owe a great debt. I mention in particular Mike Hopkins, for help with the result that $\mathcal{I}^n$ forms part of a projective class; Neil Strickland, for the joint work described in Section 5, which led to the view of a projective class described in this paper; Mark Mahowald, for the efficient construction of $\mathbb{R}\mathrm{P}^\infty$; and my advisor, Haynes Miller, for his constant encouragement, support and good advice.

2. Projective classes

This section has three parts. In the first part we recall the definition of a projective class given by Eilenberg and Moore [10]. This definition focuses on the relation between "projective" objects and three-term "exact" sequences. It appears to be the most general notion in that it allows one to do rudimentary homological algebra in any pointed category. In the second part we show that the more familiar relation between projectives and "epimorphisms" can be used to define "projective class" as long as the pointed category in which we are working has weak kernels. Finally, in the third part, we show that it is equivalent to use the relation between projectives and "null maps" as long as our pointed category has weak kernels and in addition has the property that every map *is* a weak kernel. Any triangulated category satisfies these conditions.

A reader who is more interested in the examples may move on to the next section, using the definition of "projective class" from the introduction.

Everything we say can be dualized to give a discussion of injective classes.

2.1. **Projective classes in pointed categories.** We recall the notion of a projective class in a pointed category. We will be brief; the elegant original paper [10] leaves no room for improvement.

A category $\mathcal{S}$ is **pointed** if it contains an object which is both initial and terminal. If such an object exists, it is unique up to isomorphism and is denoted 0. For any two objects X and

Y of $\mathcal{S}$, there is a unique map from X to Y that factors through 0, and we denote this map by 0 as well. By using the zero map as a basepoint, the hom functor $\mathcal{S}(-,-)$ takes values in the category of pointed sets.

For the rest of this section we assume that $\mathcal{S}$ is pointed.

A composable pair of maps $X \to Y \to Z$ is said to be a **(length two) complex** if the composite is zero.

A complex $X \xrightarrow{f} Y \xrightarrow{g} Z$ in the category of pointed sets is **exact** if $f(X) = g^{-1}(*)$, where $*$ denotes the basepoint in Z.

Definition 2.1. *Let $\mathcal{P}$ be a collection of objects of $\mathcal{S}$. A complex $X \to Y \to Z$ such that*

$$\mathcal{S}(P, X) \longrightarrow \mathcal{S}(P, Y) \longrightarrow \mathcal{S}(P, Z)$$

*is an exact sequence of pointed sets for each P in $\mathcal{P}$ is said to be **$\mathcal{P}$-exact**, and the collection of all such complexes is denoted* $\mathcal{P}$-exact.

Now let $\mathcal{C}$ be any collection of length two complexes. An object P such that

$$\mathcal{S}(P, X) \longrightarrow \mathcal{S}(P, Y) \longrightarrow \mathcal{S}(P, Z)$$

*is an exact sequence of pointed sets for each complex $X \to Y \to Z$ in $\mathcal{C}$ is said to be **$\mathcal{C}$-projective**, and the collection of all such objects is denoted* $\mathcal{C}$-proj.

Note that the class of $\mathcal{C}$-projectives is closed under coproducts and retracts.

Definition 2.2. *Let $\mathcal{P}$ be a collection of objects in $\mathcal{S}$ and let $\mathcal{C}$ be a collection of length two complexes. The pair $(\mathcal{P}, \mathcal{C})$ is **complementary** if* $\mathcal{P}$-exact $= \mathcal{C}$ *and* $\mathcal{C}$-proj $= \mathcal{P}$*, and is a **projective class** if, in addition, for each morphism $X \to Y$ in $\mathcal{S}$ there is a morphism $P \to X$ such that P is in $\mathcal{P}$ and $P \to X \to Y$ is in $\mathcal{C}$.*

It is easily checked that ($\mathcal{P}$-exact-proj, $\mathcal{P}$-exact) is a complementary pair for any collection $\mathcal{P}$ of objects of $\mathcal{S}$. Similarly, ($\mathcal{C}$-proj, $\mathcal{C}$-proj-exact) is complementary for any collection $\mathcal{C}$ of length two complexes.

If $\mathcal{S}$ is an additive category, then all of the usual results about projective resolutions may be proved. See [10] or [23] for details.

2.2. **Projective classes in categories with weak kernels.** We show that in a category with weak kernels we can define "projective class" by using epimorphisms instead of exact sequences. This observation was made in [10] for a category with strict kernels.

Given a map $f : X \to Y$ in a pointed category $\mathcal{S}$, a **weak kernel** for f is a map $W \to X$ such that

$$\mathcal{S}(V, W) \longrightarrow \mathcal{S}(V, X) \longrightarrow \mathcal{S}(V, Y)$$

is an exact sequence of pointed sets for each V in $\mathcal{S}$. This says that a map $V \to X$ factors through W if and only if the composite $V \to X \to Y$ is zero. In particular, the composite $W \to X \to Y$ is zero. We say that $\mathcal{S}$ **has weak kernels** if every map in $\mathcal{S}$ has a weak kernel.

Definition 2.3. *Let $\mathcal{P}$ be a collection of objects of $\mathcal{S}$. A map $X \to Y$ such that*

$$\mathcal{S}(P, X) \longrightarrow \mathcal{S}(P, Y)$$

*is an epimorphism for each P in $\mathcal{P}$ is said to be **$\mathcal{P}$-epic**, and the collection of all such maps is denoted* $\mathcal{P}$-epi.

Now let $\mathcal{E}$ be any collection of maps of $\mathcal{S}$. An object P such that

$$\mathcal{S}(P, X) \longrightarrow \mathcal{S}(P, Y)$$

*is an epimorphism for each map $X \to Y$ in $\mathcal{E}$ is said to be **$\mathcal{E}$-projective**, and the collection of all such objects is denoted* $\mathcal{E}$-proj.

Note that the class of $\mathcal{E}$-projectives is closed under coproducts and retracts.

Proposition 2.4. *Let $\mathcal{S}$ be a pointed category with weak kernels. Let $\mathcal{P}$ be a collection of objects and $\mathcal{E}$ a collection of morphisms such that* $\mathcal{P}$-epi $= \mathcal{E}$ *and* $\mathcal{E}$-proj $= \mathcal{P}$. *Assume also that for each X there is a map $P \to X$ in $\mathcal{E}$ with $P \in \mathcal{P}$. Then* ($\mathcal{P}$, $\mathcal{P}$-exact) *is a projective class. Moreover, every projective class is of this form for a unique pair $(\mathcal{P}, \mathcal{E})$ satisfying the above conditions.*

Proof. Suppose that we are given a pair $(\mathcal{P}, \mathcal{E})$ with $\mathcal{P}$-epi $= \mathcal{E}$ and $\mathcal{E}$-proj $= \mathcal{P}$ such that for each X there is a map $P \to X$ in $\mathcal{E}$ with P in $\mathcal{P}$. Let $\mathcal{C} = \mathcal{P}$-exact. First we will prove that $(\mathcal{P}, \mathcal{C})$ is complementary. By definition $\mathcal{C} = \mathcal{P}$-exact, and it is clear that $\mathcal{P} \subseteq \mathcal{C}$-proj, so all that remains to be shown is that $\mathcal{C}$-proj $\subseteq \mathcal{P}$. Let X be a $\mathcal{C}$-projective object, and choose a map $P \to X$ in $\mathcal{E}$ with P in $\mathcal{P}$. The sequence $P \to X \to 0$ is $\mathcal{P}$-exact and so it is exact under $\mathcal{S}(X, -)$. Thus X is a retract of P and therefore is in $\mathcal{P}$.

To finish the proof of the first part of the proposition, we must show that there are enough projectives. Let $f : Y \to Z$ be a map and choose a weak kernel $X \to Y$ for f. Let $P \to X$ be a map in $\mathcal{E}$ with P in $\mathcal{P}$. Then it is clear that $P \to Y \to Z$ is $\mathcal{P}$-exact, and therefore ($\mathcal{P}$, $\mathcal{P}$-exact) is a projective class.

Now we prove the converse. Suppose $(\mathcal{P}, \mathcal{C})$ is a projective class. If this projective class is obtained from a pair $(\mathcal{P}, \mathcal{E})$ as above, then we have $\mathcal{E} = \mathcal{P}$-epi, so uniqueness is clear. Thus our task is to show that taking this as a definition of $\mathcal{E}$ we have that $(\mathcal{P}, \mathcal{E})$ satisfies the hypotheses of the first part of the proposition. By definition $\mathcal{E} = \mathcal{P}$-epi, and it is clear that $\mathcal{P} \subseteq \mathcal{E}$-proj, so we must show that $\mathcal{E}$-proj $\subseteq \mathcal{P}$. Let X be in $\mathcal{E}$-proj. Choose a map $P \to X$

with P in $\mathcal{P}$ so that $P \to X \to 0$ is in $\mathcal{C}$. It is easy to see that $P \to X$ is in $\mathcal{E}$, and since X is in $\mathcal{E}$-proj, X is a retract of P. Therefore X is in $\mathcal{P}$.

All that is left to be done is to show that for each X there is a map $P \to X$ in $\mathcal{E}$ with P in $\mathcal{P}$. As above, simply choose a map $P \to X$ so that $P \to X \to 0$ is in $\mathcal{C}$. $\square$

2.3. **Projective classes in triangulated categories.** In this section we assume that our pointed category $\mathcal{S}$ has weak kernels and in addition has the property that every map is a weak kernel. Another way to say this is that every map $X \to Y$ lies in a sequence

$$W \to X \to Y \to Z$$

which is exact under $\mathcal{S}(U, -)$ for all objects U. A triangulated category satisfies this condition with W the fibre of $X \to Y$ and Z the cofibre (so $Z \cong \Sigma W$). There are various references for triangulated categories. The reader already looking at Margolis' book [33] will find Appendix 2 to be a handy reference. A standard (and good) reference is Verdier's portion of SGA $4\frac{1}{2}$ [40].

Definition 2.5. *Let $\mathcal{P}$ be any collection of objects of $\mathcal{S}$. A map $X \to Y$ such that*

$$\mathcal{S}(P, X) \longrightarrow \mathcal{S}(P, Y)$$

*is the zero map for each P in $\mathcal{P}$ is said to be **$\mathcal{P}$-null**, and the collection of all such maps is denoted* $\mathcal{P}$-null.

Now let $\mathcal{I}$ be any collection of maps of $\mathcal{S}$. An object P such that

$$\mathcal{S}(P, X) \longrightarrow \mathcal{S}(P, Y)$$

*is the zero map for each map $X \to Y$ in $\mathcal{I}$ is said to be **$\mathcal{I}$-projective**, and the collection of all such objects is denoted* $\mathcal{I}$-proj.

Note that the class of $\mathcal{I}$-projectives is closed under coproducts and retracts, and that if f, g and h are composable and g is in $\mathcal{P}$-null, then both fg and gh are in $\mathcal{P}$-null. If the category $\mathcal{S}$ is additive, then $\mathcal{P}$-null is an ideal.

Proposition 2.6. *Let $\mathcal{S}$ be a pointed category with weak kernels such that every map is a weak kernel. Let $\mathcal{P}$ be a collection of objects and $\mathcal{I}$ a collection of morphisms such that* $\mathcal{P}$-null $= \mathcal{I}$ *and* $\mathcal{I}$-proj $= \mathcal{P}$. *Assume also that for each X there is a projective P and a map $P \to X$ which is a weak kernel of a map in $\mathcal{I}$. Then* $(\mathcal{P}, \mathcal{P}\text{-exact})$ *is a projective class. Moreover, every projective class is of this form for a unique pair $(\mathcal{P}, \mathcal{I})$ satisfying the above conditions.*

Proof. This can be proved directly, paralleling the proof of the analogous result in the previous part of this section, but it is simpler and more illuminating to show how this formulation relates to the epimorphism formulation and then to apply Proposition 2.4.

Let $\mathcal{P}$ be a class of objects, and let $X \to Y \to Z$ be a sequence exact under $\mathcal{S}(U, -)$ for each U in $\mathcal{P}$. Then $X \to Y$ is $\mathcal{P}$-epic if and only if $Y \to Z$ is $\mathcal{P}$-null. Using this, and the facts that every $\mathcal{P}$-epic map can be detected in this way (since every map *is* a weak kernel) and that every $\mathcal{P}$-null map can be detected in this way (since every map *has* a weak kernel), it is easy to see that pairs $(\mathcal{P}, \mathcal{I})$ as described in the hypotheses correspond bijectively to pairs $(\mathcal{P}, \mathcal{E})$ as described in the previous part of this section. $\square$

For the rest of the paper we will be working in a triangulated category and will freely make use of the equivalent ways of thinking of a projective class. Also, in a triangulated category we can replace the condition

> to any X we can associate a projective P and a map $P \to X$ which is a weak kernel of a map in $\mathcal{I}$

with the condition

> any X lies in a cofibre sequence $P \to X \to Y$ with P in $\mathcal{P}$ and $X \to Y$ in $\mathcal{I}$.

Indeed, the latter clearly implies the former. And given a weak kernel $P \to X$ of a map in $\mathcal{I}$, it is easy to check that the cofibre $X \to Y$ of $P \to X$ is in $\mathcal{I}$.

3. Operations on projective classes

For the rest of this paper, $\mathcal{S}$ will be a triangulated category containing all set-indexed coproducts. We will sometimes slip and call a coproduct a "wedge", and we will write $X \vee Y$ for the coproduct of X and Y. All of our projective classes will be **stable**. That is, both $\mathcal{P}$ and $\mathcal{I}$ are assumed to be closed under suspension and desuspension. If $(\mathcal{P}, \mathcal{I})$ is a projective class in $\mathcal{S}$, we will call the objects of $\mathcal{P}$ **projective**.

3.1. **Meets and products.** There is a natural ordering on the class $\mathcal{PC}(\mathcal{S})$ of projective classes in $\mathcal{S}$. For projective classes $(\mathcal{P}, \mathcal{I})$ and $(\mathcal{Q}, \mathcal{J})$, write $(\mathcal{P}, \mathcal{I}) \leq (\mathcal{Q}, \mathcal{J})$ if $\mathcal{I}$ is contained in $\mathcal{J}$. The projective class $\mathbf{0} = (\operatorname{obj} \mathcal{S}, 0)$, whose ideal contains only the zero maps, is the smallest projective class. The projective class $\mathbf{1} = (0, \operatorname{mor} \mathcal{S})$, whose ideal contains all maps, is the largest projective class.

Proposition 3.1. *Let $\{(\mathcal{P}_\alpha, \mathcal{I}_\alpha)\}$ be a set of projective classes. Then the intersection*

$$\bigcap_\alpha \mathcal{I}_\alpha$$

is an ideal which forms part of a projective class. The projectives are precisely the retracts of wedges of objects from the union

$$\bigcup_\alpha \mathcal{P}_\alpha.$$

Note that not every ideal forms part of a projective class, so there is some content to this proposition.

The following lemma will be used to prove the proposition. In fact, it will be used just about every time we prove that we have an example of a projective class.

Lemma 3.2. *Let $\mathcal{P}$ be a class of objects closed under retracts and let $\mathcal{I}$ be an ideal. Assume that $\mathcal{P}$ and $\mathcal{I}$ are orthogonal,* i.e., *that the composite $P \to X \to Y$ is zero for each P in $\mathcal{P}$, each map $X \to Y$ in $\mathcal{I}$, and each map $P \to X$. Also assume that each object X lies in a cofibre sequence $P \to X \to Y$ with P in $\mathcal{P}$ and $X \to Y$ in $\mathcal{I}$. Then $(\mathcal{P}, \mathcal{I})$ is a projective class.*

Proof of Lemma. All that we have to show is that $\mathcal{I}$-proj $\subseteq \mathcal{P}$ and that $\mathcal{P}$-null $\subseteq \mathcal{I}$. For the former, assume that X is in $\mathcal{I}$-proj and choose a cofibre sequence $P \to X \to Y$ with P in $\mathcal{P}$ and $X \to Y$ in $\mathcal{I}$. Since X is in $\mathcal{I}$-proj, the map $X \to Y$ is zero, and so X is a retract of P. Hence X is in $\mathcal{P}$, as we have assumed that the latter is closed under retracts.

To show that $\mathcal{P}$-null $\subseteq \mathcal{I}$ is equally easy, using that $\mathcal{I}$ is an ideal. □

Proof of Proposition 3.1. Let $\mathcal{I}$ denote the intersection of the ideals $\mathcal{I}_\alpha$ and let $\mathcal{P}$ denote the collection of retracts of wedges of objects each of which lies in some $\mathcal{P}_\alpha$. It is clear that $\mathcal{P}$ and $\mathcal{I}$ are orthogonal, so we must verify that each object X lies in a cofibre sequence $P \to X \to Y$ with P in $\mathcal{P}$ and $X \to Y$ in $\mathcal{I}$. Let X be an object of $\mathcal{S}$. For each α choose a cofibre sequence $P_\alpha \to X \to Y_\alpha$. Consider the map $\vee P_\alpha \to X$. The cofibre of this map is zero when restricted to each P_α and so must lie in each $\mathcal{I}_\alpha$. Thus we have a cofibre sequence $P \to X \to Y$ with P in $\mathcal{P}$ and $X \to Y$ in $\mathcal{I}$. □

The projective class $(\mathcal{P}, \mathcal{I})$ constructed in the proposition is the **meet** of the set $\{(\mathcal{P}_\alpha, \mathcal{I}_\alpha)\}$. That is, it is the greatest lower bound.

I don't know whether an arbitrary set of projective classes has a **join** (*i.e.,* a least upper bound). If A is a set of projective classes and the collection of all projective classes which are upper bounds for A has a meet, then this is the join of A. However, the collection of all projective classes which are upper bounds for A might not be a set, and so might not have a meet. (The proof above involved a coproduct over all α.)

There is also a product on $\mathcal{PC}(\mathcal{S})$.

Proposition 3.3. *If the ideals $\mathcal{I}$ and $\mathcal{J}$ are parts of projective classes $(\mathcal{P}, \mathcal{I})$ and $(\mathcal{Q}, \mathcal{J})$, then so is their product $\mathcal{I}\mathcal{J}$, which consists of all composites fg with f in $\mathcal{I}$ and g in $\mathcal{J}$. The projectives are precisely those objects which are retracts of objects X which lie in cofibre sequences $Q \to X \to P$ with Q in $\mathcal{Q}$ and P in $\mathcal{P}$.*

Proof. Write $\mathcal{R}$ for the collection of retracts of objects X which lie in cofibre sequences $Q \to X \to P$ with Q in $\mathcal{Q}$ and P in $\mathcal{P}$. Then $\mathcal{R}$ is closed under retracts and coproducts. Using the fact that in an additive category binary coproducts are biproducts, one can show that $\mathcal{I}\mathcal{J}$ is in fact an ideal. (The point is that it is automatically closed under sums of parallel maps.) We will show now that $\mathcal{R}$ and $\mathcal{I}\mathcal{J}$ are orthogonal, *i.e.,* that any composite

$$W \longrightarrow X \xrightarrow{f} Y \xrightarrow{g} Z$$

is zero if W is in $\mathcal{R}$, f is in $\mathcal{J}$, and g is in $\mathcal{I}$. We can assume without loss of generality that W lies in a cofibre sequence $Q \to W \to P$ with Q in $\mathcal{Q}$ and P in $\mathcal{P}$. In the following diagram

$$\begin{tikzcd}
Q \arrow[d] & & & \\
W \arrow[r] \arrow[d] & X \arrow[r, "f"] & Y \arrow[r, "g"] & Z \\
P \arrow[urr, dashed] & & &
\end{tikzcd}$$

the dashed arrow exists because $\mathcal{Q}$ and $\mathcal{J}$ are orthogonal. The map $P \to Y \to Z$ is zero because $\mathcal{P}$ and $\mathcal{I}$ are orthogonal, so the map $W \to X \to Y \to Z$ is zero. This shows that $\mathcal{R}$ and $\mathcal{I}\mathcal{J}$ are orthogonal.

It remains to show that any X lies in a cofibre sequence $W \to X \to Z$ with W in $\mathcal{R}$ and $X \to Z$ in $\mathcal{I}\mathcal{J}$. To do this, choose a cofibre sequence $Q \to X \to Y$ with Q in $\mathcal{Q}$ and $X \to Y$ in $\mathcal{J}$. Now choose a cofibre sequence $P \to Y \to Z$ with P in $\mathcal{P}$ and $Y \to Z$ in $\mathcal{I}$, giving a diagram

$$\begin{tikzcd}
X \arrow[rr] & & Y \arrow[rr] \arrow[dl, "\circ" marking] & & Z \arrow[dl, "\circ" marking] \\
& Q \arrow[ul] & & P \arrow[ul] &
\end{tikzcd}$$

(A circle on an arrow $A \to B$ denotes a map $A \to \Sigma B$.) Let W be the fibre of the composite $X \to Y \to Z$. Using the octahedral axiom, one sees that W lies in a cofibre sequence $Q \to W \to P$. Thus $W \to X \to Z$ is the sequence we seek.

With the help of the lemma, we have proved that $(\mathcal{R}, \mathcal{I}\mathcal{J})$ is a projective class. $\square$

Here are some formal properties of the intersection and product operations, all of which are easy to prove.

Proposition 3.4. *Let* $(\mathcal{P}, \mathcal{I})$, $(\mathcal{Q}, \mathcal{J})$ *and* $(\mathcal{R}, \mathcal{K})$ *be projective classes. Then*

(*i*) $\mathbf{0}\mathcal{I} = \mathbf{0} = \mathcal{I}\mathbf{0}$ *and* $\mathbf{0} \cap \mathcal{I} = \mathbf{0} = \mathcal{I} \cap \mathbf{0}$. (*Recall that* $\mathbf{0} = (\operatorname{obj} \mathcal{S}, 0)$.)

(*ii*) $\mathbf{1}\mathcal{I} = \mathcal{I} = \mathcal{I}\mathbf{1}$ *and* $\mathbf{1} \cap \mathcal{I} = \mathcal{I} = \mathcal{I} \cap \mathbf{1}$. (*Recall that* $\mathbf{1} = (0, \operatorname{mor} \mathcal{S})$.)

(*iii*) $\mathcal{I} \cap \mathcal{I} = \mathcal{I}$ *and* $\mathcal{I} \cap \mathcal{J} = \mathcal{J} \cap \mathcal{I}$.

(*iv*) $\mathcal{I}\mathcal{J} \leq \mathcal{I} \cap \mathcal{J}$ *and* $\mathcal{J}\mathcal{I} \leq \mathcal{I} \cap \mathcal{J}$.

(*v*) *If* $\mathcal{J} \leq \mathcal{K}$, *then* $\mathcal{I}\mathcal{J} \leq \mathcal{I}\mathcal{K}$, $\mathcal{J}\mathcal{I} \leq \mathcal{K}\mathcal{I}$, $\mathcal{I} \cap \mathcal{J} \leq \mathcal{I} \cap \mathcal{K}$ *and* $\mathcal{J} \cap \mathcal{I} \leq \mathcal{K} \cap \mathcal{I}$. $\square$

To make the proposition readable, we have blurred the distinction between a projective class and the ideal that it corresponds to. We hope that by now the reader has been convinced that the ideal is the most natural part of a projective class. Indeed, it is usually easier to describe the ideal than the projectives, and it is through the ideal that the operations on projective classes arise naturally.

We next describe the two filtrations that a projective class determines.

3.2. **Two filtrations.** Fix a stable projective class $(\mathcal{P}, \mathcal{I})$. Define $\mathcal{I}^n$ to be the collection of all n-fold composites of maps in $\mathcal{I}$. The ideals $\mathcal{I}^n$ form a decreasing filtration of the class of morphisms of $\mathcal{S}$; write $\mathcal{I}^\infty$ for the intersection. We will use the notation $\mathcal{I}(X, Y)$ for $\mathcal{I} \cap \mathcal{S}(X, Y)$, and more generally $\mathcal{I}^n(X, Y)$ for $\mathcal{I}^n \cap \mathcal{S}(X, Y)$, $1 \leq n \leq \infty$.

By the results of the previous part of this section, each of these ideals forms part of a projective class. To fix notation and terminology, we will explicitly describe the increasing filtration of the class of objects. Let $\mathcal{P}_1 = \mathcal{P}$ and inductively define $\mathcal{P}_n$ to be the class of all retracts of objects Y which sit in cofibre sequences $X \to Y \to P$ with X in $\mathcal{P}_{n-1}$ and P projective. If X is in $\mathcal{P}_n$ but not in $\mathcal{P}_{n-1}$ we say that X has **length** n with respect to the projective class $(\mathcal{P}, \mathcal{I})$. We also say that X can be **built** from n objects of $\mathcal{P}$. Write $\mathcal{P}_\infty$ for the class of all retracts of wedges of objects of finite length. We write $\mathcal{P}_0$ for the collection of zero objects of $\mathcal{S}$, and say that they have length 0. For symmetry, we write I^0 for the collection of all morphisms in $\mathcal{S}$.

The next theorem follows immediately from Propositions 3.1 and 3.3.

Theorem 3.5. *For* $0 \leq n \leq \infty$, *the pair* $(\mathcal{P}_n, \mathcal{I}^n)$ *is a projective class.* $\square$

We say that the projective classes $(\mathcal{P}_n, \mathcal{I}^n)$ are **derived** from $(\mathcal{P}, \mathcal{I})$.

Note 3.6. If there is a cofibre sequence

$$X \to Y \to Z$$

with X of length m and Z of length n, then Y has length at most $m+n$. For if we have a composite

$$Y \longrightarrow Y_1 \longrightarrow \cdots \longrightarrow Y_{m+n}$$

of $m+n$ maps each in $\mathcal{I}$, the composite of the first m maps is zero when restricted to X, and so factors through Z. But then the composite $Y \to Z \to Y_m \to \cdots \to Y_{m+n}$ is zero because Z has length n. So, by the theorem, Y is in $\mathcal{P}_{m+n}$.

Note 3.7. Suppose $\mathcal{I}$ and $\mathcal{J}$ are projective classes with $\mathcal{I} \leq \mathcal{J}$. Since products respect order (Proposition 3.4), we have that $\mathcal{I}^n \leq \mathcal{J}^n$. Therefore the length of an object X with respect to $\mathcal{I}$ is no more than the length of X with respect to $\mathcal{J}$. In contrast, the **projective dimension** of an object X (the length of the shortest projective resolution of X) might not respect the order.

Much of this paper will focus on studying these filtrations, both abstractly and in particular examples.

4. The Adams spectral sequence

Associated to a projective class $(\mathcal{P}, \mathcal{I})$ in a triangulated category $\mathcal{S}$ is an Adams spectral sequence which we now describe. The Adams spectral sequence was discussed in the same generality in [34].

Let X be an object. By repeatedly using the fact that there are enough projectives, one can form a diagram

$$\begin{tikzcd} X = X_0 \arrow[rr] & & X_1 \arrow[rr] \arrow[dl, "\circ"] & & X_2 \arrow[rr] \arrow[dl, "\circ"] & & X_3 \arrow[dl, "\circ"] & \cdots \\ & P_0 \arrow[ul] & & P_1 \arrow[ul] & & P_2 \arrow[ul] & & \end{tikzcd} \tag{4.1}$$

with each P_n projective, each map $X_n \to X_{n+1}$ in $\mathcal{I}$, and each triangle exact. We call such a diagram an **Adams resolution** of X with respect to the projective class $(\mathcal{P}, \mathcal{I})$. Let W_n be the fibre of the composite map $X \to X_n$. Then $W_0 = 0$, $W_1 = P_0$ and W_n sits in a cofibre sequence $W_{n-1} \to W_n \to P_{n-1}$, as one sees using the octahedral axiom. In particular, W_n is in $\mathcal{P}_n$ for each n and so our Adams resolution provides us with choices of cofibre sequences $W_n \to X \to X_n$ with W_n in $\mathcal{P}_n$ and $X \to X_n$ in $\mathcal{I}^n$.

If we apply the functor $\mathcal{S}(-,Y)_*$ for some object Y we get an exact couple, which we display in unraveled form:

$$\begin{tikzcd}
\mathcal{S}(X,Y)_* \arrow[dr] & & \mathcal{S}(X_1,Y)_* \arrow[ll] \arrow[dr] & & \mathcal{S}(X_2,Y)_* \arrow[ll] \arrow[dr] & & \mathcal{S}(X_3,Y)_* \arrow[ll] & \\
& \mathcal{S}(P_0,Y)_* \arrow[ur, "\circ" marking] & & \mathcal{S}(P_1,Y)_* \arrow[ur, "\circ" marking] & & \mathcal{S}(P_2,Y)_* \arrow[ur, "\circ" marking] & & \cdots .
\end{tikzcd} \tag{4.2}$$

(We write $\mathcal{S}(X,Y)_n$ for $\mathcal{S}(\Sigma^n X,Y)$.) This exact couple leads to a spectral sequence that we call the **Adams spectral sequence** associated to the projective class. The filtration on $\mathcal{S}(X,Y)$ is the $\mathcal{I}$-adic filtration, *i.e.,* that given by intersecting with the powers $\mathcal{I}^n$ of the ideal $\mathcal{I}$.

A dual construction using an injective class also leads to a spectral sequence. For example, one can obtain the original Adams spectral sequence in this way by taking $X = S^0$, the injectives to be retracts of products of (de)suspensions of mod 2 Eilenberg–Mac Lane spectra, and the maps in the ideal to be those which induce the zero map in mod 2 singular cohomology.

For many results we will need to assume that our projective class "generates", as described in the following definition.

Definition 4.3. *A projective class is said to* ***generate*** *if for each non-zero X there is a projective P such that $\mathcal{S}(P,X) \neq 0$, or, equivalently, if the ideal $\mathcal{I}$ contains no non-zero identity maps. A third equivalent way to say this is that a map $Y \to Z$ is an isomorphism if and only if it is sent to an isomorphism by the functor $\mathcal{S}(P,-)$ for each projective P.*

One can show that a projective class $(\mathcal{P},\mathcal{I})$ generates if and only if one of its derived projective classes $(\mathcal{P}_n,\mathcal{I}^n)$ generates.

The following result uses terminology from [5].

Proposition 4.4. *Let $(\mathcal{P},\mathcal{I})$ be a projective class such that $\mathcal{I}$ is closed under countable coproducts. Then the Adams spectral sequence abutting to $\mathcal{S}(X,Y)$ is conditionally convergent for all X and Y if and only if the projective class generates.*

Proof. Assume that the projectives generate and consider the unraveled exact couple pictured in (4.2). There is an exact sequence

$$0 \longrightarrow \varprojlim \mathcal{S}(X_k,Y)_* \longrightarrow \prod \mathcal{S}(X_k,Y)_* \longrightarrow \prod \mathcal{S}(X_k,Y)_* \longrightarrow \varprojlim{}^1 \mathcal{S}(X_k,Y)_* \longrightarrow 0$$

in which the middle map is induced by the $(1-\text{shift})$ map $\vee X_k \to \vee X_k$. This map induces the identity map under $\mathcal{S}(P,-)$ for each projective P, since the shift map $\vee X_k \to \vee X_{k+1}$ is in $\mathcal{I}$ by assumption. Therefore $\vee X_k \to \vee X_k$ is an isomorphism, since the projectives generate. This shows that

$$\varprojlim \mathcal{S}(X_k, Y)_* = 0 = \varprojlim^1 \mathcal{S}(X_k, Y)_*,$$

and this is what it means for the associated spectral sequence to be conditionally convergent.

It isn't hard to see that if the projective class does not generate, for example if the identity map $X \to X$ is in $\mathcal{I}$ for some non-zero X, then the spectral sequence abutting to $\mathcal{S}(X, X)$ isn't conditionally convergent. $\square$

When our projective class generates, we can characterize the objects in $\mathcal{P}_n$ by the behaviour of the Adams spectral sequence.

Proposition 4.5. *If X has length at most n, then for each Y, the Adams spectral sequence abutting to $\mathcal{S}(X, Y)$ collapses at E_{n+1} and has $E_{n+1} = E_\infty$ concentrated in the first n rows. If the projective class generates and $\mathcal{I}$ is closed under countable coproducts, then the converse holds, and the spectral sequence converges strongly.*

We index our spectral sequence with the "Adams indexing", so that the rows contain groups of the same homological degree and the columns contain groups of the same total degree.

Proof. If X is in $\mathcal{P}_n$, then one can easily see that each X_s appearing in an Adams resolution of X is also in $\mathcal{P}_n$. Therefore, the n-fold composites $X_s \to \cdots \to X_{s+n}$ are each zero. But this implies that for each $r > n$ the differential d_r is zero, and so $E_{n+1} = E_\infty$. Moreover, if X is in $\mathcal{P}_n$, then $\mathcal{I}^n(X, Y) = 0$. But E_∞ is the associated graded of this filtration, and so it must be zero except in the first n rows.

Now we prove the converse. If the projective class generates and $\mathcal{I}$ is closed under countable coproducts, then by Proposition 4.4 the Adams spectral sequence is conditionally convergent. Therefore, if it collapses, it converges strongly. And if E_∞ is concentrated in the first n rows, then the nth stage of the filtration on $\mathcal{S}(X, Y)$ must be zero. That is, $\mathcal{I}^n(X, Y) = 0$. This is true for all Y, so X is in $\mathcal{P}_n$. $\square$

From an Adams resolution, one can form the sequence

$$0 \longleftarrow X \longleftarrow P_0 \longleftarrow\!\!\circ\, P_1 \longleftarrow\!\!\circ\, P_2 \longleftarrow\!\!\circ\, \cdots \;.$$

This is a projective resolution of X with respect to the projective class, *i.e.,* each P_s is projective and the sequence is $\mathcal{P}$-exact at each spot. Therefore the E_2-term of the Adams spectral sequence consists of the derived functors of $\mathcal{S}(-, Y)$ applied to X. By the usual argument, these are independent of the choice of resolution, and we denote them by $\operatorname{Ext}^k(-, Y)$,

$k \geq 0$. In fact, it is easy to see that from the E_2-term onwards the spectral sequence is independent of the choice of Adams resolution.

We note the following facts about the derived functors $\mathrm{Ext}^k(-, Y)$. First, there is no reason to suspect that $\mathrm{Ext}^0(-, Y) = \mathcal{S}(-, Y)$. Indeed, the kernel of the natural map $\mathcal{S}(X, Y) \to \mathrm{Ext}^0(X, Y)$ is $\mathcal{I}(X, Y)$, and differentials in the Adams spectral sequence can prevent this map from being surjective. Second, it is clear that if X has projective dimension n, then the groups $\mathrm{Ext}^k(X, Y)$ vanish for $k > n$. To prove the converse, we need to assume that the projective class generates.

Proposition 4.6. *If the projective class generates, then X has projective dimension at most n if and only if $\mathrm{Ext}^k(X, -)$ vanishes for all $k > n$.*

Proof. Assume that $\mathrm{Ext}^k(X, Y)$ vanishes for each Y and each $k > n$. Consider an Adams resolution

$$\begin{tikzcd}[column sep=small]
X = X_0 \arrow[rr] & & X_1 \arrow[rr] \arrow[dl, "\circ" description] & & X_2 \arrow[dl, "\circ" description] & \cdots & X_n \arrow[rr] & & X_{n+1} \arrow[rr] \arrow[dl, "\circ" description] & & X_{n+2} \arrow[rr] \arrow[dl, "\circ" description] & & X_{n+3} \arrow[dl, "\circ" description] & \cdots \\
& P_0 \arrow[ul] & & P_1 \arrow[ul] & & & & P_n \arrow[ul] & & P_{n+1} \arrow[ul] & & P_{n+2} \arrow[ul] & &
\end{tikzcd}$$

of X. Since the functors $\mathrm{Ext}^k(X, -)$ vanish for $k > n$, the sequence $P_n \leftarrow P_{n+1} \leftarrow P_{n+2}$ is exact at P_{n+1} after applying the functor $\mathcal{S}(-, Y)$ for any Y. In particular, the map $P_{n+1} \to X_{n+1}$ factors through P_n giving a map $P_n \to X_{n+1}$. The composite $P_{n+1} \to X_{n+1} \to P_n \to X_{n+1}$ is equal to the map $P_{n+1} \to X_{n+1}$. This shows that the composite $X_{n+1} \to P_n \to X_{n+1}$ is sent to the identity by the functor $\mathcal{S}(P, -)$ for any projective P. But since the projective class was assumed to generate, this implies that the composite $X_{n+1} \to P_n \to X_{n+1}$ is an isomorphism. Therefore X_{n+1} is a retract of P_n, and so X_n is as well. Thus X_n is projective, and

$$\begin{tikzcd}
0 & X \arrow[l] & P_0 \arrow[l] & P_1 \arrow[l, "\circ" description] & \cdots \arrow[l, "\circ" description] & P_{n-1} \arrow[l, "\circ" description] & X_n \arrow[l, "\circ" description] & 0 \arrow[l]
\end{tikzcd}$$

displays that X has projective dimension at most n.

The converse was noted above. $\square$

So when our projective class generates, saying that X has projective dimension at most n is equivalent to saying that for each Y the E_2-term of the Adams spectral sequence abutting to $\mathcal{S}(X, Y)$ is concentrated in the first $n + 1$ rows.

Here is another result that illustrates the importance of assuming that our projective class generates.

Proposition 4.7. *If the projective class generates, then* $\operatorname{proj.dim} X + 1$ *is an upper bound for the length of* X.

This follows from Proposition 4.5 in the case that $\mathcal{I}$ is closed under countable coproducts.

Proof. If

$$\begin{tikzcd} 0 & X \arrow[l] & P_0 \arrow[l] & P_1 \arrow[l, "\circ" marking] & P_2 \arrow[l, "\circ" marking] & \cdots \arrow[l, "\circ" marking] \end{tikzcd}$$

is a projective resolution of X with respect to the projective class then it can be filled out to an Adams resolution

$$\begin{tikzcd} X = X_0 \arrow[rr] & & X_1 \arrow[rr] \arrow[dl, "\circ" marking] & & X_2 \arrow[rr] \arrow[dl, "\circ" marking] & & X_3 \arrow[dl, "\circ" marking] & \cdots \\ & P_0 \arrow[ul] & & P_1 \arrow[ul] & & P_2 \arrow[ul] & & \end{tikzcd}$$

in which the composites $P_n \to X_n \to P_{n-1}$ equal the maps $P_n \to P_{n-1}$ appearing in the first diagram. If the projective resolution is finite, say with $P_n = 0$ for $n > k$, then $\mathcal{S}(P, X_n) = 0$ for all projectives P and all $n > k$. Because the projective class generates, $X_n = 0$ for $n > k$. Thus X_k has length at most one, X_{k-1} has length at most two, and inductively, $X = X_0$ has length at most $k + 1$, completing the argument. □

In general the projective dimension of X will be larger than its length; the difference comes about because when we measure the length by building up X using projectives, we don't insist that the connecting maps $X_k \to X_{k+1}$ be in $\mathcal{I}$.

5. Abstract phantom maps

In this section we discuss phantom maps in an axiomatic setting. We begin in the first part by defining phantom maps and describing some assumptions that we will need to state our results. The second part is a short study of various flavours of weak colimits. This is essential material for the third part of the section, which gives our results on phantom maps. This section is based on joint work with Neil Strickland [9].

We remind the reader that $\mathcal{S}$ denotes a triangulated category having all set-indexed coproducts.

5.1. **The phantom projective class.** We begin with a definition.

Definition 5.1. *An object* W *in* $\mathcal{S}$ *is **finite** if for any set-indexed collection* $\{X_\alpha\}$ *of objects of* $\mathcal{S}$*, the natural map*

$$\bigoplus_\alpha \mathcal{S}(W, X_\alpha) \longrightarrow \mathcal{S}(W, \bigvee_\alpha X_\alpha)$$

is an isomorphism.

To illustrate that this is a reasonable definition, we describe the finite objects in the categories that we study in the last few sections of the paper. In the stable homotopy category, an object is finite if and only if it is isomorphic to a (possibly desuspended) suspension spectrum of a finite CW-complex. In the derived category of a ring, an object is finite if and only if it is isomorphic to a bounded complex of finitely generated projectives. In both cases, the finite objects are precisely those that can be built from a finite number of copies of the spheres (S^n and $\Sigma^n R$, respectively) using cofibres and retracts.

Write $\mathcal{P}$ for the collection of retracts of wedges of finite objects.

A map $X \to Y$ is a **phantom map** if for each finite W and each map $W \to X$ the composite $W \to X \to Y$ is zero. Write $\mathcal{I}$ for the collection of phantom maps.

Definition 5.2. *We say that $\mathcal{S}$ **has a phantom projective class** if $(\mathcal{P}, \mathcal{I})$ is a projective class. We say that $\mathcal{S}$ **has a generating phantom projective class** if it has a phantom projective class and this projective class generates.*

That the projective class generates says that if $\mathcal{S}(W, X) = 0$ for each finite W, then $X = 0$. In other words, this says that $\mathcal{S}$ is **compactly generated**, in the terminology of Neeman [35].

Assuming that $(\mathcal{P}, \mathcal{I})$ is a projective class is equivalent to assuming that for each X there is a set $\{X_\alpha\}$ of finite objects such that every map $W \to X$ from a finite object to X factors through some X_α. In particular, if there is a set of isomorphism classes of finite objects, then $\mathcal{S}$ has a phantom projective class. Thus we have replaced a set-theoretic condition with the slightly more general and natural condition that $(\mathcal{P}, \mathcal{I})$ be a projective class.

The stable homotopy category and the derived category of a ring are examples of triangulated categories with generating phantom projective classes.

Our strongest results will be possible when $\mathcal{S}$ is a "Brown" category:

Definition 5.3. *We say that $\mathcal{S}$ is a **Brown** category if the following holds for any two objects X and Y. Regarding $\mathcal{S}(-, X)$ and $\mathcal{S}(-, Y)$ as functors from finite objects to abelian groups, any natural transformation $\mathcal{S}(-, X) \to \mathcal{S}(-, Y)$ is induced by a map $X \to Y$.*

In familiar settings, this can be rephrased as the assumption that natural transformations between representable homology theories are representable. The stable homotopy category and the derived category of a countable ring are Brown categories. (See [21, Theorem 4.1.5] or [35, Section 5].)

In order to prove results about the phantom projective class, we need a digression on weak colimits.

5.2. **Weak colimits.** Colimits rarely exist in a triangulated category, so in this section we introduce weaker variants that turn out to be quite useful.

Definition 5.4. *A category $\mathcal{C}$ is **small** if its class of objects forms a set. A **diagram** in $\mathcal{S}$ is a functor F from a small category $\mathcal{C}$ to $\mathcal{S}$. A **cone** from a diagram F to an object X is a natural transformation from F to the constant diagram at X. In other words, for each α in $\mathcal{C}$ we are given a map $i_\alpha : F(\alpha) \to X$ such that for each map $\alpha \to \beta$ in $\mathcal{C}$ the triangle*

$$\begin{tikzcd} F(\alpha) \arrow[d] \arrow[dr] & \\ F(\beta) \arrow[r] & X \end{tikzcd}$$

*commutes. A **weak colimit** of a diagram F is a cone through which every other cone factors. If we require the factorization to be unique, this is the definition of a **colimit**.*

If F is a diagram, then a weak colimit of F always exists. It will not be unique, but there is a distinguished choice of weak colimit defined in the following way. Writing X_α for $F(\alpha)$, there is a natural map

$$\bigvee_{\alpha\to\beta} X_\alpha \longrightarrow \bigvee_{\gamma} X_\gamma.$$

The first coproduct is over the non-identity morphisms of $\mathcal{C}$, and the second is over the objects. (We omit the identity morphisms for reasons explained in Note 5.8.) The restriction of this map to the summand X_α indexed by $\alpha \to \beta$ is

$$X_\alpha \xrightarrow{(1,-F(\alpha\to\beta))} X_\alpha \vee X_\beta \xrightarrow{\text{inclusion}} \bigvee_{\gamma} X_\gamma.$$

Let X be the cofibre of the natural map, so that X sits in a cofibre sequence

$$\bigvee_{\alpha\to\beta} X_\alpha \longrightarrow \bigvee_{\gamma} X_\gamma \longrightarrow X \longrightarrow \bigvee_{\alpha\to\beta} \Sigma X_\alpha.$$

The map $\bigvee_\gamma X_\gamma \to X$ gives a cone i to X, and it is easily checked that X and i form a weak colimit of the diagram F. Also, given another weak colimit (X', i') constructed in the same way, there is an isomorphism $h : X \to X'$ such that $ih = i'$. (This isomorphism might not be unique.) We call X the **standard weak colimit** of F.

The virtue of standard weak colimits is that they always exist and are easy to describe. However, in many cases they are too large, and there are other types of cones that are more useful.

Definition 5.5. *A cone $i : F \to X$ is a **minimal cone** if the natural map*

$$\varinjlim \mathcal{S}(W, X_\alpha) \to \mathcal{S}(W, X)$$

*is an isomorphism for each finite W. We say that a cone to X is a **minimal weak colimit** if it is a minimal cone and a weak colimit.*

When they exist, minimal weak colimits behave well, as we will see in the following proposition. The reason for introducing the weaker notion of a minimal cone is that one is often able to verify that a cone is minimal without being able to prove that it is a weak colimit. And we will find that minimal cones share some of the nice properties of minimal weak colimits.

Proposition 5.6. (i) *Assume that $\mathcal{S}$ has a phantom projective class. Then any object X is a minimal cone on a filtered diagram of finite objects.*

(ii) *Assume that $\mathcal{S}$ is a Brown category. Then any minimal cone on a diagram of projective objects is a weak colimit.*

(iii) *Assume that the finite objects generate,* i.e., *that for each non-zero X there is a finite object W and a non-zero map $W \to X$. Then a minimal weak colimit is unique up to (non-unique) isomorphism and is a retract of any other weak colimit.*

Proof. We begin by proving (i). Choose a set $\{X_\alpha\}$ of finite objects such that every map $W \to X$ from a finite object to X factors through some X_α. Consider the thick subcategory $\mathcal{C}$ generated by the X_α. In other words, consider the smallest full subcategory that contains each X_α and is closed under taking cofibres, desuspensions and retracts. Since the collection of X_α's is a set, one can show that there is a set $\mathcal{C}'$ containing a representative of each isomorphism class of objects in $\mathcal{C}$. Let $\Lambda(X)$ denote the category whose objects are the maps $W \to X$ with W in $\mathcal{C}'$ and whose morphisms are the obvious commutative triangles. One can check that this is a filtered category, using the fact that $\mathcal{C}$ is a thick subcategory. There is a natural functor $\Lambda(X) \to \mathcal{S}$ sending an object $W \to X$ to W, and there is a natural cone from this diagram to X. Consider a finite object V. We must show that the natural map

$$\varinjlim_{W \to X \in \Lambda(X)} \mathcal{S}(V, W) \to \mathcal{S}(V, X)$$

is an isomorphism. Surjectivity is easy, since any map from V to X factors through some X_α. Now we prove injectivity. Since the diagram is filtered, a general element of the colimit can be represented by a map $V \to W$ for some $W \to X \in \Lambda(X)$. Suppose that such an element is sent to zero, *i.e.,* that the composite $V \to W \to X$ is zero. Define V' to be the cofibre of the map $V \to W$. The map from W to X factors through V'. Now V' might not be in $\mathcal{C}$, but the map from it to X factors through an object W' in $\mathcal{C}$. The composite $V \to W \to V' \to W'$ is zero, since the first three terms form a cofibre sequence, so the element of $\mathcal{S}(V, W)$ given by the map $V \to W$ goes to zero under the map $W \to W'$. Thus $V \to W$ represents the zero element of the colimit, and we have proved injectivity.

Next we prove (*ii*). Let X be a minimal cone on a diagram $\{X_\alpha\}$ of projective objects. Given another cone Y we must show that there exists a map $X \to Y$ commuting with the cone maps. To construct this map, we will produce a natural transformation from $\mathcal{S}(-, X)$ to $\mathcal{S}(-, Y)$, regarded as functors on finite objects. Since there is a cone to Y, there is a natural transformation $\varinjlim \mathcal{S}(-, X_\alpha) \to \mathcal{S}(-, Y)$. And since X is a minimal cone, we have a natural isomorphism $\varinjlim \mathcal{S}(-, X_\alpha) \to \mathcal{S}(-, X)$. The composite of the inverse of the isomorphism with the map to $\mathcal{S}(-, Y)$ is the natural transformation we said that we would produce. Thus there is a map $X \to Y$ inducing this natural transformation, as we have assumed that $\mathcal{S}$ is a Brown category. This is the map we seek. By construction, the triangles commute up to phantom maps. But since we assumed that each X_α is projective, the triangles actually commute.

Finally, we prove (*iii*). Let $\{X_\alpha\}$ be a diagram with minimal weak colimit X and minimal cone Y. Because X is a weak colimit, there is a map $X \to Y$ commuting with the cone maps. Because both cones are minimal, this map is an isomorphism under $\mathcal{S}(W, -)$ for each finite W. Since the finite objects generate, we can conclude that the map $X \to Y$ is an isomorphism. In particular, if X and Y are both minimal weak colimits, then they are isomorphic.

For the second part of (*iii*), assume that X is a minimal weak colimit and that Y is a weak colimit of a diagram $\{X_\alpha\}$. Since both cones are weak colimits, there are maps $X \to Y$ and $Y \to X$ commuting with the cone maps. Because X is minimal, the composite $X \to Y \to X$ is an isomorphism, again using that the finite objects generate. $\square$

Minimal weak colimits earned their name by way of part (*iii*) of this Proposition. As far as I know, they were first defined by Margolis [33, Section 3.1].

Next we consider some examples. Let $\mathcal{C}$ be the non-negative integers, where there is one map $m \to n$ when $m \leq n$, and no maps otherwise. A functor $F : \mathcal{C} \to \mathcal{S}$ is a diagram of the form

$$X_0 \longrightarrow X_1 \longrightarrow X_2 \longrightarrow \cdots .$$

The minimal weak colimit is the cofibre of the usual $(1 - \text{shift})$ map

$$\bigvee X_k \longrightarrow \bigvee X_k.$$

Thus the minimal weak colimit is what is usually called the **telescope** of the sequence. The standard weak colimit is in this case much less manageable.

Our second example concerns the weak pushout.

Lemma 5.7. *Given a commutative diagram*

$$\begin{tikzcd}
 & V \arrow[r, equal] \arrow[d] & V \arrow[d] & \\
H \arrow[r] \arrow[d, equal] & W \arrow[r] \arrow[d] & X \arrow[r] \arrow[d] & \Sigma H \arrow[d, equal] \\
H \arrow[r] & Y \arrow[r] \arrow[d] & Z \arrow[r] \arrow[d] & \Sigma H \\
 & \Sigma V \arrow[r, equal] & \Sigma V &
\end{tikzcd}$$

with exact rows and columns, the centre square is both a weak pushout and a weak pullback. Moreover, the standard weak pushout of the diagram

$$\begin{tikzcd}
W \arrow[r] \arrow[d] & X \\
Y &
\end{tikzcd}$$

fits into a diagram of the above form.

The proof is omitted, as we will not make use of this result.

Note 5.8. In the definition of the standard weak colimit, the coproduct

$$\bigvee_{\alpha \to \beta} X_\alpha$$

was taken over the *non-identity* morphisms of the indexing category. Had we instead taken the coproduct over all of the morphisms, we would have obtained a different distinguished weak colimit, but the second half of Lemma 5.7 would no longer be true.

Note 5.9. For a weak pushout it turns out that the minimal weak colimit is less useful than the standard weak colimit. For example, the standard weak colimit of the diagram

$$\begin{tikzcd}
W \arrow[r] \arrow[d] & 0 \\
0 &
\end{tikzcd}$$

is ΣW, while the minimal weak colimit is 0.

5.3. **Consequences.** With the work we have done, we can immediately prove the following theorem.

Theorem 5.10. *Assume that $\mathcal{S}$ has a phantom projective class. If X is a minimal cone on a filtered diagram $\{X_\alpha\}$ of finite objects, then the E_2-term of the phantom Adams spectral sequence abutting to $\mathcal{S}(X,Y)$ is given by*

$$E_2^s = \varprojlim{}^s \mathcal{S}(X_\alpha, Y),$$

where $\varprojlim^s$ denotes the sth derived functor of the inverse limit functor. Here s is the homological degree and we have suppressed the internal degree.

Proof. We prove this by constructing a specific Adams resolution of X with respect to the phantom projective class. Consider the sequence

$$0 \longleftarrow X \longleftarrow \bigvee_{\alpha} X_\alpha \longleftarrow \bigvee_{\alpha\to\beta} X_\alpha \longleftarrow \bigvee_{\alpha\to\beta\to\gamma} X_\alpha \longleftarrow \cdots .$$

The wedges are over sequences of morphisms in the category over which the diagram $\{X_\alpha\}$ is indexed (and here identity morphisms are included). Because X is a weak colimit of the X_α, there are given maps $i_\alpha : X_\alpha \to X$. The map $\vee_\alpha X_\alpha \to X$ is equal to i_α on the α summand. The map $\vee_{\alpha\to\beta} X_\alpha \to \vee_\alpha X_\alpha$ sends the summand X_α indexed by the map $\alpha \to \beta$ to the α summand of the target using the identity map and to the β summand of the target using the negative of the map $X_\alpha \to X_\beta$. In general, one gets an alternating sum. When we apply the functor $\mathcal{S}(W, -)$ for finite W we get the sequence used for computing the derived functors of $\varinjlim$ [13, Appendix II, Section 3]. Since these vanish (because we have a filtered colimit) and since $\varinjlim \mathcal{S}(W, X_\alpha) = \mathcal{S}(W, X)$ (because we have a minimal cone), the sequence obtained is exact. That is, the sequence above is a phantom projective resolution of X. Therefore it is part of an Adams tower for X. The E_2-term of the Adams spectral sequence obtained by applying $\mathcal{S}(-, Y)$ is the cohomology of the sequence

$$0 \longrightarrow \bigoplus_{\alpha} \mathcal{S}(X_\alpha, Y) \longrightarrow \bigoplus_{\alpha\to\beta} \mathcal{S}(X_\alpha, Y) \longrightarrow \cdots .$$

But the cohomology at the sth place is just $\varprojlim^s \mathcal{S}(X_\alpha, Y)$, again by [13, App. II]. $\square$

Under the assumptions described in the first part of the section, the phantom projective class is very well behaved.

Theorem 5.11. *If $\mathcal{S}$ is a Brown category with a generating phantom projective class, then any object X has projective dimension at most one.*

This result was proved independently by several people [36, 35, 9]. The proofs by Neeman and Christensen-Strickland are essentially the same, and both are phrased in an axiomatic setting similar to that presented here. On the other hand, Ohkawa's proof was written in the context of the stable homotopy category and makes use of CW-structures, so it is not clear whether it goes through in the same generality. Ohkawa also noticed the consequence for the Adams spectral sequence, which is our Corollary 5.12.

Proof. Let X be an object of $\mathcal{S}$. We saw in Proposition 5.6 (*i*) that X is a minimal cone on a filtered diagram $\{X_\alpha\}$ of finite objects. The standard weak colimit Y of this diagram lies in a cofibre sequence

$$\bigvee_{\alpha\to\beta} X_\alpha \longrightarrow \bigvee_{\gamma} X_\gamma \longrightarrow Y \longrightarrow \bigvee_{\alpha\to\beta} \Sigma X_\alpha.$$

By Proposition 5.6 (*ii*) and (*iii*), X is a retract of Y. This implies that the fibre P of the natural map $\vee_\gamma X_\gamma \to X$ is a retract of $\vee_{\alpha\to\beta} X_\alpha$ and thus is projective. Moreover, the connecting map $X \to \Sigma P$ is phantom because X is a minimal cone, so we have constructed a projective resolution

$$0 \longrightarrow P \longrightarrow Q \longrightarrow X \longrightarrow 0.$$

Therefore, X has projective dimension at most one. □

Corollary 5.12. *If $\mathcal{S}$ is a Brown category with a generating phantom projective class, then the Adams spectral sequence collapses at the E_2-term and the composite of two phantom maps is zero. Moreover, if $\{X_\alpha\}$ is a filtered diagram of finite objects with minimal cone X, then there is a short exact sequence*

$$0 \longrightarrow \varprojlim{}^1 \mathcal{S}(\Sigma X_\alpha, Y) \longrightarrow \mathcal{S}(X,Y) \longrightarrow \varprojlim \mathcal{S}(X_\alpha, Y) \longrightarrow 0$$

natural in Y. The kernel consists of the phantom maps from X to Y, and $\varprojlim{}^i \mathcal{S}(\Sigma X_\alpha, Y)$ is zero for $i \geq 2$.

Proof. Consider the Adams spectral sequence abutting to $\mathcal{S}(X,Y)$. By Theorem 5.11, X has projective dimension at most one. Since the E_2-term of the Adams spectral sequence consists of the derived functors of $\mathcal{S}(-,Y)$ applied to X, it must vanish in all but the first two rows. There being no room for differentials, it collapses at E_2, degenerating into the displayed Milnor sequence.

By Proposition 4.7, we see that each object has length at most two, and so the composite of two phantoms must be zero. □

6. Topological phantom maps

In this section $\mathcal{S}$ denotes the stable homotopy category and we usually write $[X, Y]$ for the set of morphisms from X to Y. There are many descriptions of this category; a good one can be found in the book by Adams [2].

There are three parts to this section. In the first we discuss phantom maps, *i.e.,* maps which are zero when restricted to any finite spectrum, and we use the results of the previous section to conclude that there is a generalized Milnor sequence. In the second part of this section we discuss skeletal phantom maps, *i.e.,* maps which are zero when restricted to each skeleton of the source. And in the third we discuss superphantom maps, *i.e.,* maps which are zero when restricted to any (possibly desuspended) suspension spectrum.

6.1. **Phantom maps and a generalized Milnor sequence.** In the stable homotopy category, a finite spectrum is one isomorphic to a (possibly desuspended) suspension spectrum of a finite CW-complex. As in the previous section, a map $f : X \to Y$ is said to be **phantom** if the composite $W \to X \to Y$ is zero for each finite spectrum W and each map $W \to X$. Phantom maps form an ideal which we denote $\mathcal{I}$. Write $\mathcal{P}$ for the collection of all retracts of wedges of finite spectra. One can show that there exists a set $\mathcal{F}'$ of finite spectra containing a representative of each isomorphism class [33, Prop. 3.2.11], and it follows that $(\mathcal{P}, \mathcal{I})$ is a projective class. Since $\mathcal{P}$ contains the spheres, it is in fact a generating projective class. Also, it is well-known that the stable homotopy category is a Brown category (Definition 5.3).

We should point out that there do exist non-zero phantom maps. For example, if G is any non-zero divisible abelian group, then the Moore spectrum $S(G)$ is not a retract of a wedge of finite spectra. If it was, then applying integral homology would show that G is a retract of a sum of finitely generated abelian groups, which is impossible. So there is a non-zero phantom map with source $S(G)$.

For another example, consider the natural map

$$\bigvee X_\alpha \longrightarrow \prod X_\alpha$$

for some indexed collection of spectra. For finite W, $[W, \vee X_\alpha] = \oplus [W, X_\alpha]$, so we get a monomorphism

$$[W, \bigvee X_\alpha] \longrightarrow [W, \prod X_\alpha].$$

Thus the fibre of the map $\vee X_\alpha \to \prod X_\alpha$ is phantom. It is non-zero if and only if the map from the wedge to the product is not split. This is the case for the map

$$\bigvee H\mathbb{Z} \longrightarrow \prod H\mathbb{Z}$$

from the countable wedge of integral Eilenberg–Mac Lane spectra to the countable product. Indeed, if this map splits, then so does the map

$$\bigoplus \mathbb{Z} \longrightarrow \prod \mathbb{Z}$$

of abelian groups. But the cokernel of this map contains the element $[1, 2, 4, 8, \dots]$. This element is divisible by all powers of 2, so the cokernel is not a subgroup of the product.

Other examples will appear in the second and third parts of this section.

The stable homotopy category is a Brown category with a generating phantom projective class. Thus, the results of Section 5.3 hold and we find that every spectrum has projective dimension at most one and length at most two, and that the composite of two phantom maps is zero. Also, a minimal cone on a diagram of projective objects is automatically a minimal weak colimit and leads to a generalized Milnor sequence.

Minimal cones arise in practice. If a diagram is filtered, then to check that a cone is minimal, it suffices to check that it becomes a colimiting cone in homotopy groups. This is proved using the fact that a finite spectrum is built from a finite number of spheres using cofibres. One uses induction on the number of cells, that filtered colimits are exact, and the five-lemma. For this reason, all of our examples will involve filtered diagrams.

First of all, if X is a CW-spectrum (in the sense of Adams [2]) and $\{X_\alpha\}$ is a filtered collection of finite CW-subspectra whose union is X, then X is the minimal weak colimit of the X_α.

For another example, let $\{G_\alpha\}$ be a filtered diagram of abelian groups with colimit G. Then the Eilenberg–Mac Lane spectrum HG is a minimal cone on the diagram $\{HG_\alpha\}$. (By construction, this cone becomes a colimiting cone in homotopy groups.) I don't believe that such a minimal cone is always a minimal weak colimit.

More generally, a filtered homotopy colimit of spectra (taken in some geometric category of spectra) is a minimal cone, because homotopy groups commute with filtered homotopy colimits. If the spectra in the diagram are projective, then the homotopy colimit is a weak colimit.

From the first example and Corollary 5.12, we can state the following theorem.

Theorem 6.1. *Let X be a CW-spectrum and let $\{X_\alpha\}$ be a filtered diagram of finite CW-subspectra whose union is X. For any spectrum Y there is a short exact sequence*

$$0 \longrightarrow \varprojlim{}^1 [\Sigma X_\alpha, Y] \longrightarrow [X, Y] \longrightarrow \varprojlim [X_\alpha, Y] \longrightarrow 0.$$

The kernel consists precisely of the phantom maps. Moreover, $\varprojlim^i [\Sigma X_\alpha, Y]$ vanishes for $i \geq 2$. $\square$

This generalizes results of Pezennec [37], Huber and Meier [22], and Yosimura [43]. Pezennec makes the assumption that Y has finite type. Huber and Meier make the weaker assumption that the cohomology theory represented by Y is related by a universal coefficient sequence to a homology theory of finite type, while Yosimura drops the assumption that the homology theory has finite type. Our point is that no restriction on Y is necessary; this was proved independently, and earlier, by Ohkawa [36].

Using a slightly more elaborate proof (and a slightly different projective class), one can show that the assumption that each X_α is finite can be replaced by the assumption that there are no phantom maps from X_α to Y for each α.

6.2. **Skeletal phantom maps.** There is a related but smaller ideal of maps which have also been called phantom maps in the literature. We begin with some background on cellular towers. We treat cellular towers in this abstract manner because we want to define and use them without stepping outside of the homotopy category. One reason for this is that we want to make it clear that our results do not depend on a particular choice of model for the category of spectra. But more importantly, we would like our arguments to go through in any nice enough triangulated category.

Definition 6.2. *Let X be a spectrum. A **cellular tower** for X is a diagram*

$$\begin{tikzcd} \cdots \arrow[r] & X^{(n)} \arrow[r] \arrow[drrr] & X^{(n+1)} \arrow[r] \arrow[drr] & X^{(n+2)} \arrow[r] \arrow[dr] & \cdots \\ & & & & X \end{tikzcd} \tag{6.3}$$

satisfying:

(*i*) *X is the telescope of the sequence $\cdots \to X^{(n)} \to X^{(n+1)} \to \cdots$.*

(*ii*) *The fibre of the map $X^{(n)} \to X^{(n+1)}$ is a wedge of n-spheres.*

(*iii*) *The inverse limit of abelian groups $\varprojlim H_*(X^{(n)})$ is zero, where H_* denotes integral homology.*

*We say that $X^{(n)}$ is an **n-skeleton** of X.*

The first condition says that the sequence

$$\vee X^{(n)} \longrightarrow \vee X^{(n)} \longrightarrow X$$

is a cofibre sequence, where the first map is the $(1 - \text{shift})$ map.

The above definition is taken from Margolis' book [33, Section 6.3], which is also the source of the results below whose proofs are omitted.

Proposition 6.4. *A diagram* (6.3) *is a cellular tower for X if and only if all of the following conditions hold:*

(i) *The map $\pi_i(X^{(n)}) \to \pi_i(X)$ is an isomorphism for each $i < n$.*

(ii) *Each $H_n(X^{(n)})$ is a free abelian group and $H_n(X^{(n)}) \to H_n(X^{(n+1)})$ is an epimorphism.*

(iii) *$H_i(X^{(n)}) = 0$ for $i > n$.* □

Proposition 6.5. *Each spectrum X has a cellular tower.* □

Proposition 6.6. *Let $\cdots \to X^{(n)} \to X^{(n+1)} \to \cdots \to X$ be a cellular tower for X and let $\cdots \to Y^{(n)} \to Y^{(n+1)} \to \cdots \to Y$ be a cellular tower for Y. Given any map $X \to Y$, there exist maps $X^{(n)} \to Y^{(n)}$ making the following diagram commute:*

$$\begin{tikzcd}
\cdots \arrow[r] & X^{(n)} \arrow[r] \arrow[d] & X^{(n+1)} \arrow[r] \arrow[d] & \cdots \arrow[r] & X \arrow[d] \\
\cdots \arrow[r] & Y^{(n)} \arrow[r] & Y^{(n+1)} \arrow[r] & \cdots \arrow[r] & Y \; .
\end{tikzcd}$$

□

Definition 6.7. *We say that a spectrum X is an* ***n-skeleton*** *if it is an n-skeleton of some spectrum Y. We say that X is a* ***skeleton****, or is* ***skeletal****, if X is an n-skeleton for some n.*

It is easy to see that if X is an n-skeleton, then X is an n-skeleton of itself.

Proposition 6.8. *A spectrum X is an n-skeleton if and only if $H_i(X) = 0$ for $i > n$ and $H_n(X)$ is a free abelian group (possibly zero). In particular, X is skeletal if and only if it has bounded above integral homology.* □

We call a map $f : X \to Y$ a **skeletal phantom map** if the composite $W \to X \to Y$ is zero for each skeleton W and each map $W \to X$. By Proposition 6.6, it suffices to test this for the skeleta $X^{(n)}$ of a fixed cellular tower for X. This shows that skeletal phantoms form part of a projective class, because it allows us to restrict to a *set* of test objects. In this case the projectives are retracts of wedges of skeletal spectra. In the cofibre sequence $\vee X^{(n)} \to X \to \vee \Sigma X^{(n)}$, $\vee X^{(n)}$ is projective and $X \to \vee \Sigma X^{(n)}$ is a skeletal phantom.

Since skeletal phantoms are phantoms, it follows from the previous part of this section that the composite of two skeletal phantoms is zero. In fact, this is obvious because the cofibre sequence $\vee X^{(n)} \to X \to \vee \Sigma X^{(n)}$ shows that every X has length at most two. That the composite of two skeletal phantoms is zero has been known for some time. (See [14] and [16].)

Applying the functor $[-, Y]$ to the cofibre sequence $\vee X^{(n)} \to X \to \vee \Sigma X^{(n)}$ immediately reveals another Milnor sequence

$$0 \longrightarrow \varprojlim{}^1 [\Sigma X^{(n)}, Y] \longrightarrow [X, Y] \longrightarrow \varprojlim [X^{(n)}, Y] \longrightarrow 0.$$

It is clear that the kernel consists precisely of the skeletal phantoms, and that, for $i \geq 2$, $\varprojlim^i [X^{(n)}, Y] = 0$ (since, for $i \geq 2$, $\varprojlim^i$ is zero for a diagram indexed by the integers).

If X is a spectrum with finite skeleta, then a map $X \to Y$ is phantom if and only if it is a skeletal phantom. But in general the ideal of skeletal phantoms is strictly smaller than the ideal of phantom maps. For example, we saw above that for G a non-zero divisible abelian group, the Moore spectrum $S(G)$ is the source of a non-zero phantom map. But $S(G)$ is skeletal, since it has bounded above integral homology, and so $S(G)$ is not the source of a non-zero skeletal phantom.

On the other hand, there are non-zero skeletal phantoms. Consider $H\mathbb{Z}/p$, the mod p Eilenberg–Mac Lane spectrum. This spectrum has finite skeleta, so it suffices to show that there is a non-zero phantom with source $H\mathbb{Z}/p$. This is equivalent to showing that $H\mathbb{Z}/p$ is not a retract of a wedge of finite spectra. If it was and $H\mathbb{Z}/p \to \vee W_\alpha$ was the inclusion, then we would look at the composite $H\mathbb{Z}/p \to \vee W_\alpha \to \prod W_\alpha$, which is a monomorphism in homotopy groups. But any such map is zero because there are no maps from $H\mathbb{Z}/p$ to a finite spectrum (see [32], [30], or [38]) and so we would conclude that $H\mathbb{Z}/p$ has no homotopy. This type of argument will appear repeatedly in what follows.

A non-zero phantom from $H\mathbb{Z}/p$ is also an example of a phantom map which is not divisible by p, since p kills $[H\mathbb{Z}/p, Y]$ for any Y.

6.3. **Superphantom maps.** There is another special class of phantom maps. We call a map $f : X \to Y$ a **superphantom map** if the composite $W \to X \to Y$ is zero for each (possibly desuspended) suspension spectrum W and each map $W \to X$. Again the superphantoms form an ideal which is part of a projective class. The projectives here are all retracts of wedges of (possibly desuspended) suspension spectra. To see that we have a projective class, one uses the following lemma, which allows us to use a *set* of objects in order to test whether a map $X \to Y$ is superphantom.

Lemma 6.9. *A map $X \to Y$ is a superphantom if and only if for each n the composite*

$$\Sigma^{-n} \Sigma^\infty \Omega^\infty \Sigma^n X \longrightarrow X \longrightarrow Y$$

is zero.

Recall that Σ^∞ is left adjoint to Ω^∞. The map $\Sigma^{-n}\Sigma^\infty\Omega^\infty\Sigma^n X \to X$ is the nth desuspension of the counit of the adjunction.

Proof. The "only if" direction is clear. So suppose $f : X \to Y$ is a map such that each composite $\Sigma^{-n}\Sigma^\infty\Omega^\infty\Sigma^n X \to X \to Y$ is zero, and let W be a space. Consider a map $\Sigma^{-n}\Sigma^\infty W \to X$. This map factors through $\Sigma^{-n}\Sigma^\infty\Omega^\infty\Sigma^n X \to X$, as one sees by suspending everything n times and using that Σ^∞ is left adjoint to Ω^∞. Thus the composite $\Sigma^{-n}\Sigma^\infty W \to X \to Y$ is zero, and we have shown that $f : X \to Y$ is a superphantom. □

Margolis states in his book [33, p. 81] that whether there exist non-zero superphantoms is an open question. We answer this question now.

Proposition 6.10. *The mod p Eilenberg–Mac Lane spectrum $H\mathbb{Z}/p$ is the source of a non-zero superphantom map.*

Proof. We note that by the main result of [19] there are no maps from $H\mathbb{Z}/p$ to a suspension spectrum. Therefore, by the argument in the previous part of this section, $H\mathbb{Z}/p$ is not a retract of a wedge of suspension spectra. Therefore $H\mathbb{Z}/p$ is the source of a non-zero superphantom map. □

There are skeletal phantoms which are not superphantoms. For example, following Gray [15] one can show that there are uncountably many skeletal phantoms from $\mathbb{C}\mathrm{P}^\infty$ to S^3. But $\mathbb{C}\mathrm{P}^\infty$ is a suspension spectrum, and so none of these maps is a superphantom. We don't know if there is an example of a superphantom which is not a skeletal phantom.

7. Topological ghosts

In this section we continue to work in the stable homotopy category. Again, there are three parts. In the first, we describe the ghost projective class and give its elementary properties. In the second part, which is a little bit longer, but lots of fun, we calculate the length of $\mathbb{R}\mathrm{P}^n$ for small n. And in the third part, we show that the Adams spectral sequence with respect to the ghost projective class in a category of A_∞ modules over an A_∞ ring is in fact a universal coefficient spectral sequence, and we explain how this gives new lower bounds on the ghost-length of a spectrum.

7.1. **The ghost projective class.** A map $X \to Y$ is called a **ghost** if the induced map $\pi_*(X) \to \pi_*(Y)$ of homotopy groups is zero. Let $\mathcal{I}$ denote the ideal of ghosts. Let $\mathcal{P}$ denote the class of all retracts of wedges of spheres. It is easy to see that $(\mathcal{P}, \mathcal{I})$ is a projective class (use Lemma 3.2) and that it generates (as $\mathcal{P}$ contains the spheres).

Let's begin by noticing that there are spectra of arbitrarily high length with respect to this projective class. For example, the length of $\mathbb{R}\mathrm{P}^{2^k}$ is at least $k+1$. One sees this by noticing that if $u \in H^1(\mathbb{R}\mathrm{P}^{2^k}; \mathbb{Z}/2)$ is the non-zero class, then $\mathrm{Sq}^{2^{k-1}} \cdots \mathrm{Sq}^4 \mathrm{Sq}^2 \mathrm{Sq}^1 u$ is non-zero. But the composite

$$\mathbb{R}\mathrm{P}^{2^k} \xrightarrow{u} \Sigma H\mathbb{Z}/2 \xrightarrow{\mathrm{Sq}^1} \Sigma^2 H\mathbb{Z}/2 \xrightarrow{\mathrm{Sq}^2} \cdots \xrightarrow{\mathrm{Sq}^{2^{k-1}}} \Sigma^{2^k} H\mathbb{Z}/2$$

is in $\mathcal{I}^k$ and thus would be zero if $\mathbb{R}\mathrm{P}^{2^k}$ had length k or less.

On the other hand, by constructing $\mathbb{R}\mathrm{P}^n$ one cell at a time, it is clear that the length of $\mathbb{R}\mathrm{P}^n$ is no more than n. We will see in the next part of this section that it is possible to improve on both of these bounds.

The filtration of the morphisms of $\mathcal{S}$ is also interesting. Again by looking at composites of Steenrod operations, one sees that the powers $\mathcal{I}^k$ are all non-trivial. Also, every phantom map is a ghost. The ghost-filtration of a phantom map is analogous to what Gray called the "index" of a phantom map, so we'll use that terminology here. Every non-zero phantom map from the Moore spectrum $S(G)$ has index 1, since $S(G)$ has length 2 with respect to the ghost ideal. And we saw in Section 6.1 that for G non-zero and divisible, such phantom maps exist.

The Lusternik–Schnirelmann category of a space X is an upper bound for the "cup length" of the reduced cohomology H^*X. That is, if the cup product $u_1 \cdots u_n$ is non-zero for some $u_i \in H^*X$, then the Lusternik–Schnirelmann category of X is at least n. Stably there are no products in cohomology, but we have instead the action of the Steenrod algebra. And we saw above that if there is a chain of Steenrod operations acting non-trivially on the mod 2 cohomology of a spectrum X, say $\mathrm{Sq}^{i_1} \cdots \mathrm{Sq}^{i_n} u \neq 0$, then the ghost-length of X is at least n. Thus we think of ghost-length as a stable analogue of Lusternik–Schnirelmann category.

If for a pair X and Y of spectra the Adams spectral sequence abutting to $[X, Y]$ is strongly convergent, then $\mathcal{I}^\infty(X, Y)$ is zero. Call a non-zero map in $\mathcal{I}^\infty$ a **persistent ghost**. We saw in Theorem 3.5 that $(\mathcal{P}_\infty, \mathcal{I}^\infty)$ is a projective class, but for all we know at this point, it could be that $\mathcal{P}_\infty$ contains all of the objects of the stable homotopy category and that $\mathcal{I}^\infty$ is zero. Our first goal is to show that this is not the case.

Proposition 7.1. *Let X be a non-zero connective spectrum such that there are no maps from X to a connective wedge of spheres. Then there is a persistent ghost $X \to Y$ for some Y.*

If X is a dissonant spectrum, such as $H\mathbb{Z}/p$, then there are no maps from X to a connective wedge of spheres. Indeed, a connective wedge of spheres is a (possibly desuspended) suspension spectrum, and by a result of Hopkins and Ravenel [19] suspension spectra are harmonic.

Proof. If X is connective, the projectives P_n in a ghost Adams resolution for X can be chosen to be connective wedges of spheres. Let W_n be the fibre of the map $X \to X_n$. There is a natural map $\vee W_n \to X$ whose cofibre is a persistent ghost. So if there is no persistent ghost with X as its source, then X is a retract of $\vee W_n$. We saw at the beginning of Section 4 that W_n lies in the cofibre sequence $W_{n-1} \to W_n \to P_{n-1}$. It follows inductively that if there are no maps from X to a connective wedge of spheres, then there are no maps from X to W_n for each n. Thus the map $X \to \vee W_n \to \prod W_n$ is zero. But this map is also monic in homotopy groups, and so we conclude that X is zero. $\square$

Example 2.6 in the paper [29] also implies that $\mathcal{I}^\infty$ is non-zero, but I have been unable to follow the argument.

Along the same lines, we can also obtain the next result.

Proposition 7.2. *If X is a connective spectrum of length n, then X can be built using n connective wedges of spheres.*

We use the word "built" to mean "built using cofibres and retracts", as in the definition of $\mathcal{P}_n$.

Proof. Form an Adams resolution (4.1) of X with the P_n chosen to be connective wedges of spheres. Since X has length n, the composite $X = X_0 \to \cdots \to X_n$ is zero. Thus X is a retract of W_n, the fibre of this composite. But we saw at the beginning of Section 4 that W_n can be built from $P_0, \ldots, P_{n-1}$, n connective wedges of spheres. $\square$

Similarly, one can show that a spectrum of finite type ($\pi_i X$ a finitely generated abelian group for each i) and length n can be built using n wedges of spheres with only a finite number of spheres of each dimension.

Corollary 7.3. *If X is a connective spectrum of finite length, then X is harmonic.*

Proof. We saw in the proof of Proposition 7.1 that a spectrum built from a finite number of connective wedges of spheres is harmonic, so the result follows from Proposition 7.2. $\square$

Question 7.4. Is every spectrum of finite length harmonic? It would suffice to show that any wedge of spheres is harmonic.

Another related question is whether every finite spectrum of finite length can be built from a finite number of *finite* wedges of spheres. We don't know the answer to either question.

7.2. **The ghost-length of real projective spaces.** In this part of the section we give upper and lower bounds on the ghost-length of $\mathbb{R}\mathrm{P}^n$. The upper bound is obtained by building $\mathbb{R}\mathrm{P}^n$ carefully using a cofibre sequence involving $\mathbb{C}\mathrm{P}^\infty$ and a Thom spectrum. Our first lower bound is simply the length of the longest chain of non-zero Steenrod operations acting on the mod 2 cohomology of $\mathbb{R}\mathrm{P}^n$. This bound agrees with the upper bound for $n < 20$, providing us with a calculation of the length of $\mathbb{R}\mathrm{P}^n$ in this range. However, we will show that while the squares Sq^1, Sq^2, Sq^4 and Sq^8 have ghost filtration 1, the squares Sq^{2^k} for $k \geq 4$ have ghost filtration at least 2, and using this we obtain a significantly better lower bound. The ghost filtration of the Steenrod squares is closely connected with the Hopf and Kervaire invariant problems, and we give theorems explaining this relation.

We work localized at the prime 2. The 2-local category is triangulated, so all of our general theory applies. We will write H^*X for the mod 2 cohomology of X.

We begin by recalling the action of the Steenrod algebra on $H^*\mathbb{R}\mathrm{P}^\infty = \mathbb{F}_2[x]$, with $|x| = 1$. The Steenrod square Sq^{2^k} acts non-trivially on x^n if and only if the kth bit in the binary expansion of n is 1. Figure 1 illustrates.

Now we describe a construction of $\mathbb{R}\mathrm{P}^\infty$ that was explained to us by Mahowald. The double cover map $S^1 \to S^1$, with fibre $\mathbb{Z}/2$, can be extended to a fibre sequence of spaces

$$S^1 \xrightarrow{2} S^1 \longrightarrow \mathbb{R}\mathrm{P}^\infty \longrightarrow \mathbb{C}\mathrm{P}^\infty \xrightarrow{2} \mathbb{C}\mathrm{P}^\infty$$

by applying the classifying space functor. Thus $\mathbb{R}\mathrm{P}^\infty$ is the circle bundle of the complex line bundle $\eta \otimes \eta$ over $\mathbb{C}\mathrm{P}^\infty$. (Here η denotes the tautological line bundle.) The Thom space is the disk bundle modulo the circle bundle, and the disk bundle is homotopy equivalent to $\mathbb{C}\mathrm{P}^\infty$, so there is a cofibre sequence

$$\mathbb{R}\mathrm{P}^\infty \longrightarrow \mathbb{C}\mathrm{P}^\infty \longrightarrow \mathrm{Th}\,.$$

Writing T for the desuspension of the Thom space, we get a stable cofibre sequence

$$T \longrightarrow \mathbb{R}\mathrm{P}^\infty \longrightarrow \mathbb{C}\mathrm{P}^\infty$$

which we will use to build $\mathbb{R}\mathrm{P}^n$ efficiently. The Thom isomorphism tells us that T can be built with a single cell in each odd non-negative dimension, and no other cells. So $H^*T = H^*\Sigma\mathbb{C}\mathrm{P}_0^\infty$, where $\mathbb{C}\mathrm{P}_0^\infty$ denotes $\mathbb{C}\mathrm{P}^\infty \vee S^0$. In fact, the map $H^*T \leftarrow H^*\mathbb{R}\mathrm{P}^\infty$ is surjective, so we can deduce the action of the Steenrod algebra on H^*T. Figure 1 displays

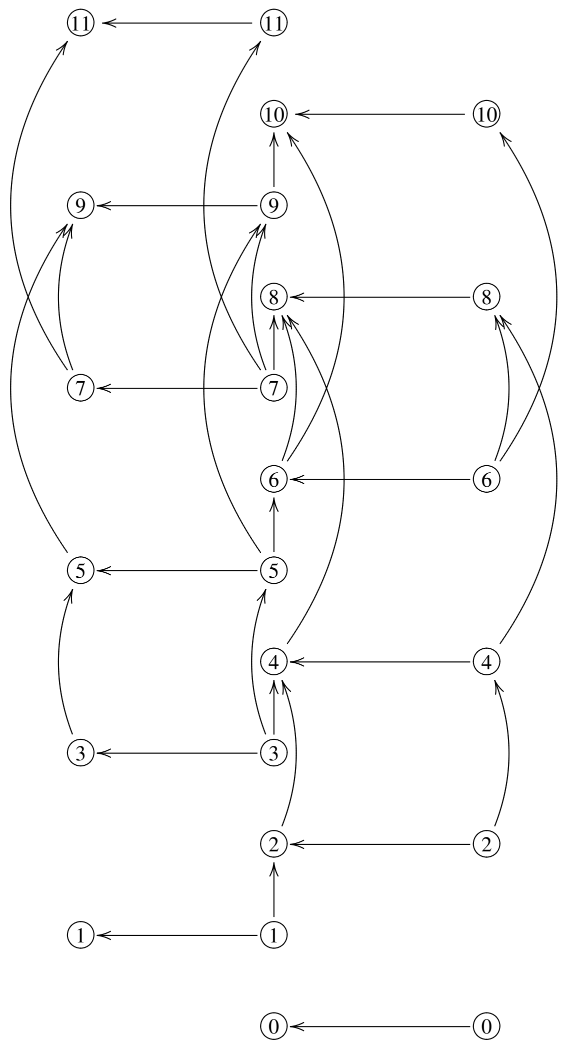

$$0 \longleftarrow H^*T \longleftarrow H^*\mathbb{RP}_0^\infty \longleftarrow H^*\mathbb{CP}_0^\infty \longleftarrow 0$$

FIGURE 1. The action of the Steenrod algebra

the low degree part of the short exact sequence of modules over the Steenrod algebra that we obtain. Each circle represents a basis element over $\mathbb{F}_2$ and the vertical arrows give the action of the Sq^{2^k}'s. To make the pattern as nice as possible, we have replaced $\mathbb{RP}^\infty$ with $\mathbb{RP}_0^\infty = \mathbb{RP}^\infty \vee S^0$, and similarly with $\mathbb{CP}^\infty$.

The essential point is that $\mathbb{RP}^\infty$ can be built up odd cells first. That is, we can first build T completely, and then attach the even cells. By looking at the cofibre sequence of skeleta, we see that $\mathbb{R}\mathrm{P}^n$ can also be built by building up the odd part, and then the even part. We will use this to deduce an upper bound on the length of $\mathbb{R}\mathrm{P}^n$.

Let us begin by considering $\mathbb{R}\mathrm{P}^3$. The length of $\mathbb{R}\mathrm{P}^3$ is at least two, because $\mathrm{Sq}^1 x = x^2$. And since $\mathbb{R}\mathrm{P}^3$ has only three cells, it has length at most three. In fact, it has length two. One way to see this is to notice that the 3-skeleton of T is $S^1 \vee S^3$, and so $\mathbb{R}\mathrm{P}^3$ is formed by attaching a 2-cell to this wedge.

Both $\mathbb{R}\mathrm{P}^4$ and $\mathbb{R}\mathrm{P}^5$ can similarly be seen to have length three because the 4- and 5-cells can to added to $\mathbb{R}\mathrm{P}^3$ simultaneously. (A connection between them would be detected by a non-zero Sq^1 on x^4.) But what about $\mathbb{R}\mathrm{P}^6$? This is trickier and will reveal the power of the decomposition into odd and even cells. We just mentioned that the 5-cell can be added after the 1-, 2- and 3-cells have been added. But even cells are never needed for the attachment of odd cells, so the 5-cell can actually be attached at the *same time* as the 2-cell. And this means that the 6-cell can be attached at the same time as the 4-cell. So $\mathbb{R}\mathrm{P}^6$ also has length three.

With this under our belt, we now prove the following theorem.

Theorem 7.5. *The length of* $\mathbb{R}\mathrm{P}^n$ *is no more than* $\lfloor n/4 \rfloor + 2$. *Here* $\lfloor x \rfloor$ *denotes the greatest integer less than or equal to* x.

Proof. We prove this inductively by showing that we can add four cells at a time, if we are careful about the order. To make the pattern work from the start, we build $\mathbb{R}\mathrm{P}^n_0 = \mathbb{R}\mathrm{P}^n \vee S^0$ instead of $\mathbb{R}\mathrm{P}^n$. This makes little difference as both have the same length. We start with $S^1 \vee S^3$. To this we can add the 5-cell, since an odd cell only needs the odd cells below it. Now in T there is no Sq^2 from the 5-cell to the 7-cell, and this implies that the 7-cell is not attached to the 5-cell. So we can attach the 7-cell to $S^1 \vee S^3$. And we can of course attach the 0- and 2-cells. Call the resulting complex W. W has cells in dimensions 0, 1, 2, 3, 5 and 7, and has length 2. To W we can attach the 9- and 11-cells, since they only require the odd cells below them and are not connected in T. At the same time we can attach the 4- and 6-cells, because they are not connected in $\mathbb{C}\mathrm{P}^\infty$. (Again, because there is no Sq^2.) Thus we can add the 4-, 6-, 9- and 11-cells to W to get a length 3 complex X. In a similar way we see that we can add the 8-, 10-, 13- and 15-cells to X. Thus $\mathbb{R}\mathrm{P}^8$, $\mathbb{R}\mathrm{P}^9$, $\mathbb{R}\mathrm{P}^{10}$ and $\mathbb{R}\mathrm{P}^{11}$ all have length at most 4. This pattern continues, proving the theorem. $\square$

We saw in the previous part of this section that if there is a chain of Steenrod operations acting non-trivially on the cohomology of a spectrum X, say $\mathrm{Sq}^{i_1} \cdots \mathrm{Sq}^{i_n} u \neq 0$, then the ghost-length of X is at least n. Letting $\mathrm{St}(\mathbb{R}\mathrm{P}^n_0)$ denote one more than the length of the longest such chain in the cohomology of $\mathbb{R}\mathrm{P}^n_0$, one obtains the following sequence of numbers,

n	-1	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$\mathrm{St}(\mathbb{R}\mathrm{P}_0^n)$	0	1	1	2	2	3	3	3	3	4	4	4	4	5	5	5	5	6	6	6	6	6 ,

where we regard $\mathbb{R}\mathrm{P}_0^{-1}$ as the zero object, which has length zero. If we count the number of consecutive 0's, then the number of consecutive 1's, and so on, we obtain the sequence 1, 2, 2, 4, 4, 4, 8, 8, 8, 8, ...

Theorem 7.6 (Vakil [39]). *The sequence obtained in this way consists of the powers of* 2*, in order, with* $k+1$ *repetitions of* 2^k*.* □

This theorem completely determines the sequence $\mathrm{St}(\mathbb{R}\mathrm{P}_0^n)$. The proof of the theorem has a striking feature. Vakil studies a more fundamental sequence defined in the following way. The nth term of the sequence is one more than the length of the longest chain of non-zero Steenrod operations in $\mathbb{R}\mathrm{P}^\infty$ which *ends at the nth cell.* Starting with $n = 1$, this sequence begins $1, 2, 1, 3, 2, 3, 1, \ldots$. The nth term of the sequence displayed in the table above is obtained by taking the supremum of the first n terms of the fundamental sequence. The method of proof that Vakil uses to determine where the jumps in the values of the suprema occur is to explicitly define certain "canonical moves". In more detail, given an n, Vakil determines a non-zero Steenrod operation Sq^{2^k} which is the last step in a longest chain ending at the n-cell (if there are any Steenrod operations hitting the n-cell). Using this, one can quickly compute the nth term in the fundamental sequence by following the canonical moves downwards until one reaches a cell not hit by a Steenrod operation. And one can also use the canonical moves to prove Theorem 7.6.

Note 7.7. For $2 \leq n \leq 19$, $\mathrm{St}(\mathbb{R}\mathrm{P}^n) = \lfloor n/4 \rfloor + 2$. Thus we know the length of $\mathbb{R}\mathrm{P}^n$ for such n. But $\mathrm{St}(\mathbb{R}\mathrm{P}^{20}) = 6$ and $\lfloor 20/4 \rfloor + 2 = 7$, and for larger n this just get worse. For example, $\mathrm{St}(\mathbb{R}\mathrm{P}^{2^{20}}) = 136$ and $\lfloor 2^{20}/4 \rfloor + 2 = 2^{18} + 2$.

One might wonder whether the lower bound is correct. It is not. For example, the length of $\mathbb{R}\mathrm{P}^{2^{20}}$ is actually at least 264. The first case where I know that the Steenrod length gives the wrong answer is for $n = 56$. The Steenrod length of $\mathbb{R}\mathrm{P}^{56}$ is 10 but the length is at least 11. These facts are deduced from the following result.

Theorem 7.8. *The Steenrod operations* Sq^1*,* Sq^2*,* Sq^4 *and* Sq^8 *have ghost-filtration exactly one. The Steenrod operations* Sq^{16}*,* Sq^{32}*, ... have ghost-filtration at least two.*

This allows us to count a higher Steenrod square occurring in the cohomology of $\mathbb{R}\mathrm{P}^n$ as two maps. Using a computer to do the computation, this is how we obtained the improved lower bound on the length of $\mathbb{R}\mathrm{P}^{2^{20}}$. It turns out that the improved lower bound is still wrong in general. For example, the improved lower bound tells us that the length of $\mathbb{R}\mathrm{P}^{127}$ is at least 17. But it also tells us that the length of $\mathbb{R}\mathrm{P}^{128}$ is at least 19, hence the length of $\mathbb{R}\mathrm{P}^{127}$ must be at least 18.

Proof of Theorem 7.8. The (non-identity) Steenrod squares all have filtration at least one. It is well-known that Sq^1, Sq^2, Sq^4 and Sq^8 act non-trivially on the length two complexes $\mathbb{R}\mathrm{P}^2$, $\mathbb{C}\mathrm{P}^2$, $\mathbb{H}\mathrm{P}^2$ and $\mathbb{O}\mathrm{P}^2$ respectively, so these operations must have filtration exactly one.

To show that the higher squares can be factored into two pieces, each zero in homotopy, we make use of Adams' result on the Hopf invariant one problem [1]. Adams shows that no complex with only two cells supports a non-zero Sq^{2^k} with $k \geq 4$. Consider the beginning of a ghost Adams resolution of H, the mod 2 Eilenberg–Mac Lane spectrum:

$$\begin{tikzcd}
H \arrow[rr] & & \bar{H} \arrow[rr] \arrow[dl, "\circ"] & & \bar{\bar{H}} \arrow[dl, "\circ"] \\
& S^0 \arrow[ul] & & P \arrow[ul] &
\end{tikzcd} .$$

Here P is a large wedge of spheres. The fibre W of the map $H \to \bar{\bar{H}}$ has length two—it lies in a cofibre sequence $S^0 \to W \to P$. If the composite

$$W \longrightarrow H \xrightarrow{\mathrm{Sq}^{2^k}} \Sigma^{2^k} H$$

is zero, then Sq^{2^k} factors through $H \to \bar{H} \to \bar{\bar{H}}$ and thus has filtration at least two. So we have reduced the problem from checking that Sq^{2^k} vanishes on *all* cohomology classes of *all* length two spectra to checking that it is zero on a *particular* cohomology class in a *particular* length two spectrum. To do this, notice that the composite $W \to H \to \Sigma^{2^k} H$ factors (uniquely) through the map $W \to P$. To show that the map $P \to \Sigma^{2^k} H$ is zero, it suffices to check this on each 2^k-dimensional sphere appearing as a summand of P. Choose such a summand, and consider the following diagram

$$\begin{tikzcd}
& \Sigma^{2^k} H & & \\
& H \arrow[u, "\mathrm{Sq}^{2^k}"] & & \\
S^0 \arrow[r] & W \arrow[r] \arrow[u] & P \arrow[r] \arrow[uul] & S^1 \\
S^0 \arrow[r] \arrow[u, equal] & W' \arrow[r] \arrow[u] & S^{2^k} \arrow[r] \arrow[u] & S^1 \arrow[u, equal]
\end{tikzcd} ,$$

in which W' is defined to be the fibre of the map from S^{2^k} to S^1 and the map $W' \to W$ is some choice of fill-in map. We must show that the composite from S^{2^k} to $\Sigma^{2^k} H$ is zero. Well, by Adams' result, it is zero when restricted to W'; so it factors through S^1; but there are no maps from S^1 to $\Sigma^{2^k} H$; so it must be null. $\square$

Now we quote a theorem which relates the filtration of the Steenrod squares to the Kervaire invariant problem.

Theorem 7.9 (W.-H. Lin [31]). *If, for $k \geq 4$, the Kervaire class θ_{k-1} exists and has order 2, then Sq^{2^k} has filtration exactly 2.* $\square$

We end this section with a conjecture.

Conjecture 7.10. *The length of $\mathbb{R}\mathrm{P}^n$ is an increasing function of n.*[1]

7.3. **A universal coefficient spectral sequence.** In this part of the section we need to briefly step outside of the homotopy category. Given a ring spectrum R, we would like to have a triangulated category of R-modules. Unfortunately, this isn't possible if R is simply a monoid object in the homotopy category. So by an "A_∞ ring spectrum" R we mean any notion of structured ring spectrum such that the homotopy category R-Mod of "A_∞ module spectra" is triangulated and satisfies the following formal properties. There is a "free module" functor $F : \mathcal{S} \to R$-Mod which is left adjoint to a "forgetful" functor $U : R\text{-Mod} \to \mathcal{S}$. Both F and U preserve triangles, commute with suspension, and commute with coproducts, and the composite UF is naturally isomorphic to the functor sending X to $R \wedge X$. We will usually omit writing U, and will write $R \wedge X$ for both FX and UFX, with the context making clear which is intended.

There are various notions of structured ring spectra available to us today [12, 20, 28]. Unfortunately, we know of no published proof that the formal properties hold in these settings. It is certainly expected that they do.

Fix an A_∞ ring spectrum R and write R for FS^0. R is the "sphere" in the category of A_∞ R-modules. Indeed, by adjointness, $[R, M]_R = [S^0, M] = \pi_0 M$, where we write $[M, N]_R$ for maps from M to N in R-Mod.

Because F preserves triangles, commutes with suspension, and preserves retracts, it is clear that if a spectrum X can be built from n wedges of spheres, then FX can be built from n wedges of suspensions of R. To make this more precise, we note that in R-Mod there is a projective class $(\mathcal{P}_R, \mathcal{I}_R)$, where $\mathcal{P}_R$ is the collection of retracts of wedges of suspensions

[1] This has now been proven by the author.

of R and $\mathcal{I}_R$ is the collection of maps zero in homotopy groups. To put it another way, $\mathcal{P}_R$ is the image of $\mathcal{P}$ under F (with retracts thrown in), and $\mathcal{I}_R$ is $U^{-1}\mathcal{I}$. And our claim is that the length of FX with respect to $(\mathcal{P}_R, \mathcal{I}_R)$ is no more than the length of X with respect to $(\mathcal{P}, \mathcal{I})$.

So it would be useful to give a lower bound for the length of an A_∞ R-module. This is accomplished in the remainder of the section.

Theorem 7.11. *Let M and N be A_∞ modules over the A_∞ ring spectrum R. Then there is a conditionally convergent spectral sequence*

$$E_2^{*,*} = \mathrm{Ext}_{R_*}^{*,*}(M_*, N_*) \implies [M, N]_R^*.$$

If M has length at most n with respect to $(\mathcal{P}_R, \mathcal{I}_R)$, then $E_{n+1} = E_\infty$.

By taking $M = R \wedge X$ we get the following consequence.

Corollary 7.12. *If X is a spectrum and N is an A_∞ module over an A_∞ ring spectrum R, then there is a conditionally convergent spectral sequence*

$$E_2^{*,*} = \mathrm{Ext}_{R_*}^{*,*}(R_*X, N_*) \implies N^*X.$$

If X has length at most n with respect to $(\mathcal{P}, \mathcal{I})$, then $E_{n+1} = E_\infty$. □

This is called the **universal coefficient spectral sequence**. For another account, see [11].

Proof of Theorem. The spectral sequence is simply the Adams spectral sequence with respect to the projective class $(\mathcal{P}_R, \mathcal{I}_R)$. The E_2-term consists of the derived functors of $[-, N]_R$ applied to M. A projective resolution of M is a sequence

$$0 \longleftarrow M \longleftarrow P_0 \longleftarrow P_1 \longleftarrow \cdots$$

of A_∞ R-modules with each P_s in $\mathcal{P}_R$ and which is exact in homotopy. Thus applying $\pi_*(-)$ gives a projective resolution of M_*. One can check that $[P_s, N]_R = \mathrm{Hom}_{R_*}(\pi_* P_s, N_*)$ and thus that the E_2-term is $\mathrm{Ext}_{R_*}^{*,*}(M_*, N_*)$. The projective class generates and $\mathcal{I}_R$ is closed under coproducts, so by Proposition 4.4, the spectral sequence is conditionally convergent.

That the spectral sequence collapses at E_{n+1} when M has length at most n is Proposition 4.5. □

Thus the existence of a non-zero differential d_n implies that X has length at least n. We suspect that for $X = \mathbb{R}\mathrm{P}^n$, $R = J$ or KO, and $N = KO$, this gives a very good lower bound for the ghost-length of $\mathbb{R}\mathrm{P}^n$. However, while in some cases we have been able to compute the E_2- and E_∞-terms, we haven't been able to conclude anything about the differentials.

We finish this section by mentioning the following example of Theorem 7.11. Take $R = S^0$ and $M = N = H$, the mod 2 Eilenberg–Mac Lane spectrum. Then we get a spectral sequence with E_2-term

$$E_2^{*,*} = \operatorname{Ext}^{*,*}_{\pi_* S^0}(\mathbb{Z}/2, \mathbb{Z}/2)$$

converging to the Steenrod algebra. This has been called the dual or reverse Adams spectral sequence and has been studied by various authors [31, 29]. An unstable version is described in [3] and [4].

8. Algebraic ghosts

We now discuss a projective class which provided the motivation for this work. Indeed, an old result of Kelly [26] (presented here as Theorem 8.5) concludes that under certain conditions a composite of maps vanishes. We wondered whether there was more than just a superficial similarity between this result and the fact that in the stable homotopy category a composite of two phantoms is zero. It turns out that the arguments can be arranged to have a common part, by proving that the ideals in question are parts of projective classes and then applying Theorem 3.5.

We work in the derived category of an abelian category in this section, and so we begin with a brief overview of the derived category. Good references are [42] and [24].

Let $\mathcal{A}$ be an abelian category with enough projectives. We mean this in the usual sense—that is, we are assuming that the categorical projectives and the categorical epimorphisms form a projective class. We also assume that $\mathcal{A}$ satisfies Grothendieck's AB 5 axiom which says that set-indexed colimits exist and filtered colimits are exact [17]. We write Ch for the category of $\mathbb{Z}$-graded chain complexes of objects of $\mathcal{A}$ and degree 0 chain maps. To fix notation, assume the differentials have degree -1. For an object X of Ch, define $Z_n X := \ker(d : X_n \to X_{n-1})$ and $B_n X := \operatorname{im}(d : X_{n+1} \to X_n)$, and write $H_n X$ for the quotient.

Write $\mathcal{K}$ for the category in which we identify chain homotopic maps. It is well-known that the category $\mathcal{K}$ is triangulated, so we will only briefly describe the triangulation. There is an automorphism Σ of Ch which is defined on objects by $(\Sigma X)_n = X_{n-1}$ and $d_{\Sigma X} = -d_X$, and on morphisms by $(\Sigma f)_n = f_{n-1}$; this induces an automorphism Σ of $\mathcal{K}$ which serves as the suspension for the triangulated structure. A short exact sequence

$$0 \longrightarrow W \xrightarrow{i} X \xrightarrow{p} Y \longrightarrow 0$$

of chain complexes is **weakly split** if for each n the sequence

$$0 \longrightarrow W_n \xrightarrow{i_n} X_n \xrightarrow{p_n} Y_n \longrightarrow 0$$

is split. Given a weakly split short exact sequence of chain complexes as above, choose for each n a splitting of the nth level, *i.e.,* choose maps $q_n : X_n \to W_n$ and $j_n : Y_n \to X_n$ such that $p_n j_n = 1$, $q_n i_n = 1$ and $i_n q_n + j_n p_n = 1$. Define $h_n : Y_n \to W_{n-1}$ to be $q_{n-1} d_X j_n$. One easily checks that h is a chain map $Y \to \Sigma W$ and that up to homotopy h is independent of the choice of splittings. A triangle in $\mathcal{K}$ is a sequence isomorphic (in $\mathcal{K}$) to one of the form $W \to X \to Y \to \Sigma W$ constructed in this way from a weakly split short exact sequence. See [24] for details.

One fact we use about the triangulation is that the homology functors $H_n : \mathcal{K} \to \mathcal{A}$ send triangles to long exact sequences.

A chain map $f : X \to Y$ is a **quasi-isomorphism** if it induces an isomorphism in homology. The derived category $\mathcal{D}$ is the category obtained from Ch by formally inverting the quasi-isomorphisms. (It is equivalent to invert the quasi-isomorphisms in $\mathcal{K}$.) With our hypotheses on $\mathcal{A}$, this category of fractions exists [42, Exercise 10.4.5]. In fact, it is equivalent to the full subcategory of $\mathcal{K}$ containing the "cofibrant" complexes. A complex X is **cofibrant** if it can be written as an increasing union $X = \cup_{n \geq 0} C^n$ of subcomplexes C^n with $C^0 = 0$ and C^n / C^{n-1} a complex of projectives with zero differential. To prove the equivalence of categories, one shows that for any X there is a cofibrant complex W and a quasi-isomorphism $W \to X$, and that when X is cofibrant, the natural map $\mathcal{K}(X, Y) \to \mathcal{D}(X, Y)$ is an isomorphism for all Y.

The derived category is a triangulated category. The automorphism Σ of Ch induces an automorphism Σ of $\mathcal{D}$. There is a natural functor $\mathcal{K} \to \mathcal{D}$, and a sequence $X \to Y \to Z \to \Sigma X$ is a triangle in $\mathcal{D}$ if and only if it is isomorphic (in $\mathcal{D}$) to the image of a triangle in $\mathcal{K}$. One important fact about the triangulation is that if $f : X \to Y$ is a chain map which is an epimorphism in each degree, then the fibre of f in $\mathcal{D}$ is given by the degreewise kernel.

We record the following lemma, whose proof is straightforward.

Lemma 8.1. *Let X be an object of $\mathcal{K}$. Then the following are equivalent:*

(*i*) *X is isomorphic in $\mathcal{K}$ to a complex of projectives with zero differential.*

(*ii*) *X is isomorphic in $\mathcal{K}$ to a complex Y with Y_n, $Z_n Y$, $B_n Y$ and $H_n Y$ projective for each n.*

(*iii*) *X is isomorphic in $\mathcal{K}$ to a complex Y with $B_n Y$ and $H_n Y$ projective for each n.* □

Since the homology of a complex is analogous to the homotopy of a spectrum, we call a map which is zero in homology a **ghost**. Let $\mathcal{I}$ denote the ideal of ghosts in $\mathcal{D}$. Call an object P of $\mathcal{D}$ **ghost projective** if it is isomorphic (in $\mathcal{D}$) to an object satisfying the equivalent conditions of Lemma 8.1. Write $\mathcal{P}$ for the collection of ghost projective complexes. One can check that $\mathcal{P}$ is closed under retracts.

As the reader has no doubt guessed, we have the following result.

Proposition 8.2. *The pair* $(\mathcal{P}, \mathcal{I})$ *forms a projective class.*

Proof. We begin by showing that $\mathcal{P}$ and $\mathcal{I}$ are orthogonal. Let P be a ghost projective complex. Without loss of generality, we may assume that P is a complex of projectives with zero differential. Since a complex with zero differential is a coproduct of complexes concentrated in a single degree, we may even assume that P is a projective object concentrated in degree zero, say. Such a complex is cofibrant, so $\mathcal{D}(P, Y) = \mathcal{K}(P, Y)$ for any Y. Now suppose that $f : P \to Y$ is a map. That is, we have a map $f : P_0 \to Y_0$ such that the composite $P_0 \to Y_0 \to Y_{-1}$ is zero; so f factors through the kernel to give a map $P \to Z_0 Y$. If f is zero in homology, then the composite $P_0 \to Z_0 Y \to H_0 Y$ is zero; so f factors through the inclusion of $B_0 Y$ into $Z_0 Y$. And because P_0 is projective, f lifts over the epimorphism $Y_1 \to B_0 Y$. That is, f is null homotopic. We conclude that if P is ghost projective and $g : P \to X$ and $h : X \to Y$ are maps in $\mathcal{D}$ with h zero in homology, then the composite is zero in $\mathcal{D}$.

Now, given a chain complex X, we construct a cofibre sequence $P \to X \to Y$ with P ghost projective and with $X \to Y$ zero in homology. First we choose projectives P^{B_n} and P^{H_n} and epimorphisms $P^{B_n} \to B_n X$ and $P^{H_n} \to H_n X$. It is easy to see that we can now choose a projective P^{Z_n} and an epimorphism $P^{Z_n} \to Z_n X$ which fit into a diagram

$$\begin{tikzcd} 0 \arrow[r] & P^{B_n} \arrow[r] \arrow[d, two heads] & P^{Z_n} \arrow[r] \arrow[d, two heads] & P^{H_n} \arrow[r] \arrow[d, two heads] & 0 \\ 0 \arrow[r] & B_n X \arrow[r] & Z_n X \arrow[r] & H_n X \arrow[r] & 0 \end{tikzcd}$$

with exact rows. Similarly, one can form a diagram

$$\begin{tikzcd} 0 \arrow[r] & P^{Z_n} \arrow[r] \arrow[d, two heads] & P^{X_n} \arrow[r] \arrow[d, two heads] & P^{B_n} \arrow[r] \arrow[d, two heads] & 0 \\ 0 \arrow[r] & Z_n X \arrow[r] & X_n \arrow[r] & B_n X \arrow[r] & 0 \end{tikzcd}$$

with exact rows and with P^{X_n} projective. Defining $P_n := P^{X_n}$ and using the composite $P^{X_{n+1}} \to P^{B_n} \to P^{Z_n} \to P^{X_n}$ as a differential, we get a chain complex P. By definition,

$Z_nP = P^{Z_n}$. The same holds for B_n and H_n, so, by Lemma 8.1, P is ghost projective. The maps $P^{X_n} \to X_n$ piece together to give a chain map $P \to X$. Under the functor H_n this map induces the chosen epimorphism $P^{H_n} \to H_nX$, and since H_n is an exact functor, the cofibre map $X \to Y$ is zero in homology.

Thus by Lemma 3.2 we have a projective class. □

The main result of this section is the following theorem.

Theorem 8.3. *Let X be a complex such that the projective dimensions of B_nX and H_nX are less than k for each n. Then the projective dimension of X with respect to the ideal of ghosts is less than k. In particular, X has length at most k, and a k-fold composite*

$$X \longrightarrow Y^1 \longrightarrow \cdots \longrightarrow Y^k$$

of maps each zero in homology is zero in $\mathcal{D}$.

Proof. Let $X^0 = X$. As in the proof of the previous proposition, one can construct a map $P^0 \to X^0$ such that each of the maps $P_n^0 \to X_n^0$, $H_nP^0 \to H_nX^0$, $B_nP^0 \to B_nX^0$ and $Z_nP^0 \to Z_nX^0$ is an epimorphism from a projective. Let X^1 be the suspension of the degreewise kernel, which is a choice of cofibre, and inductively continue this process. For any exact sequence

$$0 \longrightarrow A \longrightarrow Q_{k-2} \longrightarrow \cdots \longrightarrow Q_1 \longrightarrow H_nX \longrightarrow 0$$

in $\mathcal{A}$ with each Q_i projective, the object A is projective because of the assumption on H_nX. The same holds with H_nX replaced with B_nX, Z_nX and X_n, for each n, so applying Lemma 8.1 one finds that X^{k-1} is ghost projective. Thus X^{k-1} has length at most one, X^{k-2} has length at most two, and inductively, $X = X^0$ has length at most k. □

Corollary 8.4. *If every object in $\mathcal{A}$ has projective dimension less than k, then every object of $\mathcal{D}$ has projective dimension less than k with respect to the ideal of ghosts.* □

Note that the projective dimension of H_nX is a lower bound for the projective dimension of X.

By assuming that X is a complex of projectives, we can strengthen the conclusion of the theorem and obtain the following result of Kelly [26].

Theorem 8.5. *Let X be a complex of projectives such that the projective dimensions of B_nX and H_nX are less than k for each n. Then the projective dimension of X with respect* *to the ideal of ghosts in $\mathcal{K}$ is less than k. In particular, X has length at most k, and a composite*

$$X \longrightarrow Y^1 \longrightarrow \cdots \longrightarrow Y^k$$

of k maps in Ch*, each zero in homology, is null homotopic.*

We emphasize that we are claiming that the composite is null homotopic, not just zero in the derived category.

Sketch of proof. One begins by showing that the collection of retracts (in $\mathcal{K}$) of complexes satisfying the conditions of Lemma 8.1 along with the ideal of maps in $\mathcal{K}$ which are zero in homology is a projective class. Then one imitates the proof of Theorem 8.3, making use of the fact that if Z is a complex of projectives and $Y \to Z$ is a map in Ch which is degreewise surjective, then the complex X of degreewise kernels is the fibre (since the sequence $X \to Y \to Z$ is degreewise split). $\square$

Corollary 8.6. *Let X be a complex of projectives such that the projective dimensions of $B_n X$ and $H_n X$ are less than k for each n. Then X has the homotopy type of a cofibrant complex. That is, X is isomorphic in $\mathcal{K}$ to a cofibrant complex.*

Proof. This follows from Theorem 8.5 and the following lemma. $\square$

Lemma 8.7. *A complex in $\mathcal{K}$ of finite length has the homotopy type of a cofibrant complex.*

Proof. A complex X has the homotopy type of a cofibrant complex if and only if the natural map $\mathcal{K}(X,Y) \to \mathcal{D}(X,Y)$ is an isomorphism for all Y. A complex of projectives with zero differential is cofibrant, and so a retract in $\mathcal{K}$ of such a complex has the homotopy type of a cofibrant complex. The functors $\mathcal{K}(-,Y)$ and $\mathcal{D}(-,Y)$ are exact and send coproducts to products, so the collection of complexes of the homotopy type of a cofibrant complex is closed under coproducts and cofibre sequences. Thus this collection contains all complexes of finite length. $\square$

One can also prove Corollary 8.6 directly and then deduce Theorem 8.5 from Theorem 8.3.

9. Algebraic phantom maps

In this section we study phantom maps in the derived category of an associative ring R. We restrict attention from a general abelian category to the category of R-modules because it is in this setting that one can easily discuss the notion of purity. We provide such a discussion in the first part of this section. In the second part we describe the finite objects in the derived category of R and the phantom projective class that results. Under some assumptions on R we show that there is a relation between pure extensions and phantom maps. We end by recounting an example of Neeman's that shows that phantom maps can compose non-trivially and hence that Brown representability can fail in the derived category of R-modules.

In this section Ch denotes the category of chain complexes of R-modules, $\mathcal{K}$ denotes the category obtained from Ch by identifying chain homotopic maps, and $\mathcal{D}$ denotes the derived category obtained from Ch by inverting quasi-isomorphisms. See Section 8 for descriptions of these categories.

We write Hom for Hom_R and $\otimes$ for $\otimes_R$, and, unless otherwise stated, we take our modules to be left R-modules.

9.1. **Purity.** A module P is **pure projective** if it is a summand of a coproduct of finitely presented modules. A sequence $K \to L \to M$ is **pure exact** if it is exact under $\operatorname{Hom}(P, -)$ for each pure projective P. A longer sequence is **pure exact** if each three term subsequence is pure exact. A map $L \to M$ is a **pure epimorphism** if each map $P \to M$ from a pure projective factors through L.

We recall some standard facts about purity. A good reference here is [41].

Proposition 9.1. (*i*) *The pure projectives, pure exact sequences and pure epimorphisms form a projective class as described in Sections 2.1 and 2.2.*

(*ii*) *Every projective is pure projective; every pure exact sequence is exact; and every pure epimorphism is epic.*

(*iii*) *An infinite sequence*

$$\cdots \longrightarrow M_{-1} \longrightarrow M_0 \longrightarrow M_1 \longrightarrow \cdots$$

is pure exact if and only if it is exact after tensoring with each right module E. In particular, a finite sequence

$$0 \longrightarrow M_0 \longrightarrow M_1 \longrightarrow \cdots \longrightarrow M_k \longrightarrow 0$$

beginning and ending with 0 *is pure exact if and only if it is exact after tensoring with each right module E. A similar statement holds for semi-infinite sequences beginning or ending with* 0. □

Note 9.2. The sequence $0 \longrightarrow \mathbb{Z} \xrightarrow{2} \mathbb{Z}$ of abelian groups is pure exact, but fails to be exact after tensoring with $\mathbb{Z}/2$.

For $k \geq 1$, a **pure extension of length** k is a pure exact sequence

$$0 \longrightarrow N \longrightarrow A_1 \longrightarrow \cdots \longrightarrow A_k \longrightarrow M \longrightarrow 0.$$

A morphism of extensions is a commutative diagram of the form

$$\begin{tikzcd}
E: & 0 \arrow[r] & N \arrow[r] \arrow[d, equal] & A_1 \arrow[r] \arrow[d] & \cdots \arrow[r] & A_k \arrow[r] \arrow[d] & M \arrow[r] \arrow[d, equal] & 0 \\
E': & 0 \arrow[r] & N \arrow[r] & A'_1 \arrow[r] & \cdots \arrow[r] & A'_k \arrow[r] & M \arrow[r] & 0 .
\end{tikzcd}$$

We say that two pure extensions E and E' are equivalent if they are connected by a chain of morphisms of pure extensions, with the morphisms going in either direction. Under the operation of Baer sum, the collection $\mathrm{PExt}^k(M, N)$ of equivalence classes of pure extensions of length k forms an abelian group which is a functor of M and N: the induced maps are given by pullback and pushforward of extensions. The functors Ext^k can be defined in the same way. (See [27] for details.) By forgetting pure exactness, one gets a natural transformation $\mathrm{PExt}^k \to \mathrm{Ext}^k$ which is not always a monomorphism. We set $\mathrm{PExt}^0(M, N) = \mathrm{Ext}^0(M, N) = \mathrm{Hom}(M, N)$.

Extensions $\alpha \in \mathrm{PExt}^k(K, L)$ and $\beta \in \mathrm{PExt}^l(L, M)$ can be spliced together to give their **Yoneda product**, an element of $\mathrm{PExt}^{k+l}(K, M)$ which we denote $\beta\alpha$. We also use this composition notation when one or both of k and l are zero. Similarly, one can compose extensions in Ext, and the natural transformation $\mathrm{PExt} \to \mathrm{Ext}$ respects composition.

As one might expect, $\mathrm{PExt}^*(M, N)$ can be calculated by forming a pure projective resolution of M, applying $\mathrm{Hom}(-, N)$, and taking homology. In fact, associated to each filtered diagram $\{M_\alpha\}$ of finitely presented modules with colimit M there is a natural pure projective resolution of M. (That every M is in fact a filtered colimit of finitely presented modules is proved below.) Consider the following sequence:

$$\cdots \longrightarrow \bigoplus_{\alpha \to \beta \to \gamma} M_\alpha \longrightarrow \bigoplus_{\alpha \to \beta} M_\alpha \longrightarrow \bigoplus_{\alpha} M_\alpha \longrightarrow M \longrightarrow 0. \tag{9.3}$$

The sums are over sequences of morphisms in the filtered diagram. Write $i_\alpha : M_\alpha \to M$ for the colimiting cone to M. The map $\oplus_\alpha M_\alpha \to M$ is equal to i_α on the α summand. A summand of $\oplus_{\alpha \to \beta} M_\alpha$ is indexed by a triple (α, β, u), where u is a map $M_\alpha \to M_\beta$ such that $i_\alpha = i_\beta u$. The map $\oplus_{\alpha \to \beta} M_\alpha \to \oplus_\alpha M_\alpha$ sends the summand M_α indexed by such a triple to the M_α summand using the identity map and to the M_β summand using the map $-u$. In general, one gets an alternating sum. Taking cohomology gives the derived functors of colimit (see [13, App. II]) and because colimits of filtered diagrams are exact the sequence is exact. Since tensor products commute with colimits, it is in fact pure exact and hence can be used to compute $\mathrm{PExt}^*(M,-)$.

As promised in the previous paragraph, we now show that every module M is a filtered colimit of finitely presented modules. To avoid set theoretic problems, fix a set of finitely presented modules containing a representative from each isomorphism class. Let $\Lambda(M)$ be the category whose objects are maps $P \to M$ where P is in our set of finitely presented modules. The morphisms are the obvious commutative triangles. This category is filtered, and there is a natural functor $\Lambda(M) \to R$-Mod sending $P \to M$ to P. The colimit of this diagram is M. A smaller but less canonical filtered diagram of finitely presented modules with colimit M is described in [7, Exercise I.2.10].

The exact sequence (9.3) leads to a spectral sequence

$$E_2^{p,q} = \varprojlim{}^p \operatorname{Ext}^q(M_\alpha, N) \implies \operatorname{Ext}^{p+q}(M,N) \tag{9.4}$$

involving the derived functors of the inverse limit functor. One way to construct this spectral sequence is as follows. Break the exact sequence displayed above into short exact sequences, defining modules M_i in the process:

$$\begin{tikzcd}
M = M_0 & & M_1 \arrow[dl, tail] & & M_2 \arrow[dl, tail] & & M_3 \arrow[dl, tail] & \\
& \bigoplus_{\alpha} M_\alpha \arrow[ul, two heads] & & \bigoplus_{\alpha\to\beta} M_\alpha \arrow[ul, two heads] & & \bigoplus_{\alpha\to\beta\to\gamma} M_\alpha \arrow[ul, two heads] & & \cdots .
\end{tikzcd}$$

Applying $\operatorname{Ext}^*(-,N)$ produces an unraveled exact couple

$$\begin{tikzcd}
\operatorname{Ext}^*(M,N) \arrow[dr] & & \operatorname{Ext}^*(M_1,N) \arrow[ll, "\circ" marking] \arrow[dr] & & \operatorname{Ext}^*(M_2,N) \arrow[ll, "\circ" marking] \arrow[dr] & & \operatorname{Ext}^*(M_3,N) \arrow[ll, "\circ" marking] & \\
& \prod_{\alpha} \operatorname{Ext}^*(M_\alpha,N) \arrow[ur] & & \prod_{\alpha\to\beta} \operatorname{Ext}^*(M_\alpha,N) \arrow[ur] & & \prod_{\alpha\to\beta\to\gamma} \operatorname{Ext}^*(M_\alpha,N) \arrow[ur] & & \cdots ,
\end{tikzcd}$$

in which the horizontal maps are the connecting maps in the long exact sequence of Ext groups. This exact couple leads to a spectral sequence abutting to $\operatorname{Ext}^*(M,N)$, and the E_2-term is the cohomology of the bottom row which, by [13, App. II], is $\varprojlim{}^p \operatorname{Ext}^q(M_\alpha,N)$. The same construction works if the sequence (9.3) is replaced by any pure projective resolution of M. The spectral sequences produced in this way agree from the E_2-term onwards. The E_2-term consists of the derived functors of $\operatorname{Ext}^*(-,N)$ with respect to the pure projective class.

The spectral sequence determines a decreasing filtration of $\operatorname{Ext}^k(M,N)$. We write $\mathrm{P}^l\operatorname{Ext}^k(M,N)$ for the lth stage, so $\mathrm{P}^0\operatorname{Ext}^k(M,N) = \operatorname{Ext}^k(M,N)$. The next stage,

$\mathrm{P}^1\mathrm{Ext}^k(M,N)$, consists of those extensions α in $\mathrm{Ext}^k(M,N)$ which can be factored into a product $\beta\gamma$ with β in $\mathrm{Ext}^{k-1}(K,N)$ and γ in $\mathrm{PExt}^1(M,K)$ for some K. Indeed, the map $\mathrm{Ext}^{k-1}(M_1,N) \to \mathrm{Ext}^k(M,N)$, whose image is $\mathrm{P}^1\mathrm{Ext}^k(M,N)$, is given by composition with the pure extension $0 \to M_1 \to \oplus_\alpha M_\alpha \to M \to 0$. In general, for $0 \leq l \leq k$, $\mathrm{P}^l\mathrm{Ext}^k(M,N)$ consists of those extensions α in $\mathrm{Ext}^k(M,N)$ which can be factored into a product $\beta\gamma$ with β in $\mathrm{Ext}^{k-l}(K,N)$ and γ in $\mathrm{PExt}^l(M,K)$ for some K. Note that $\mathrm{P}^k\mathrm{Ext}^k(M,N)$ is exactly the image of $\mathrm{PExt}^k(M,N)$ in $\mathrm{Ext}^k(M,N)$, and that $\mathrm{P}^l\mathrm{Ext}^k(M,N)$ is zero for $l > k$.

Question 9.5. From the exact couple it is clear that $\mathrm{P}^1\mathrm{Ext}^k(M,N)$ can also be described as those extensions which pullback to zero under any map from a finitely presented module to M. Is it true that $\mathrm{P}^l\mathrm{Ext}^k(M,N)$ consists of those extensions which pullback to zero under any map from a module of pure projective dimension at most $l-1$ to M?

9.2. **The relation between phantom maps and pure extensions.** Recall that R is an associative ring and that $\mathcal{D}$ denotes the derived category of (left) R-modules. We begin by characterizing the finite objects in $\mathcal{D}$, *i.e.,* those objects X such that $\mathcal{D}(X, \oplus Y_\alpha) = \oplus\mathcal{D}(X, Y_\alpha)$. We first note that R, regarded as a complex concentrated in degree 0, is finite. Indeed, R is cofibrant, and so $\mathcal{D}(R,X) \cong \mathcal{K}(R,X)$. It is easy to see that $\mathcal{K}(R,X) \cong H_0X$, and since $H_0(\oplus Y_\alpha) = \oplus H_0 Y_\alpha$, we see that R is finite, as claimed. This also shows that R is a weak graded generator for $\mathcal{D}$, *i.e.,* that a complex X is (isomorphic to) zero if and only if $\mathcal{D}(\Sigma^n R, X) = 0$ for all n.

The following concept will be of use to us. An R-module M is **FP** if there is a finite resolution of M by finitely generated projectives.

We will also need the following terminology. A full subcategory $\mathcal{T}$ of a triangulated category $\mathcal{S}$ is a **thick subcategory** if it is closed under cofibres, retracts and desuspensions. The thick subcategory **generated** by a collection $\mathcal{U}$ of objects is the full subcategory determined by all objects which can be built from a finite number of objects of $\mathcal{U}$ using cofibres, retracts and desuspensions. It is the smallest thick subcategory containing $\mathcal{U}$.

Proposition 9.6. *Let X be an object of $\mathcal{D}$. Then the following are equivalent:*

(*i*) *X is finite.*

(*ii*) *X is in the thick subcategory generated by R.*

(*iii*) *X is isomorphic to a bounded complex of finitely generated projectives.*

Moreover, if H_nX is FP for all n and is zero for all but finitely many n, then X is finite.

Note 9.7. The converse of the last statement doesn't hold in general. For example, if $R = k[x]/x^2$ for some field k, then the chain complex

$$\cdots \longrightarrow 0 \longrightarrow k[x]/x^2 \xrightarrow{x} k[x]/x^2 \longrightarrow 0 \longrightarrow \cdots$$

is finite, but has homology modules of infinite projective dimension.

Proof of Proposition 9.6. Since R is a finite weak graded generator of $\mathcal{D}$, it follows from [21, Corollary 2.3.12] that (*i*) and (*ii*) are equivalent. (That (*ii*) implies (*i*) is straightforward, but the other direction is less so.)

We next prove that (*iii*) implies (*ii*): The thick subcategory generated by R is closed under finite coproducts and retracts, and so it contains all finitely generated projective modules (considered as complexes concentrated in one degree). A bounded complex of finitely generated projectives can be built from such complexes using a finite number of cofibres and thus is also contained in the thick subcategory generated by R.

Now we prove that (*ii*) implies (*iii*): Let $\mathcal{T}$ be the collection of complexes isomorphic to a bounded complex of finitely generated projectives. Since R is in $\mathcal{T}$, it suffices to prove that $\mathcal{T}$ is a thick subcategory. Given objects X and Y in $\mathcal{T}$ and a map $f : X \to Y$ we must show that the cofibre of f is in $\mathcal{T}$. We can assume without loss of generality that X and Y are bounded complexes of finitely generated projectives. A choice of cofibre has as its nth module the direct sum $Y_n \oplus X_{n-1}$. Thus the cofibre is again in $\mathcal{T}$. The subcategory $\mathcal{T}$ is clearly closed under desuspension, so it remains to show that $\mathcal{T}$ is closed under retracts. This is proved as Proposition 3.4 of [6].

Finally, we prove that if X is a complex such that each $H_n X$ is FP and only finitely many are non-zero, then X is finite. If X has no homology, then $X \cong 0$ in $\mathcal{D}$, so X is finite. If X has homology concentrated in one degree, and this module has a finite resolution by finitely generated projectives, then X is isomorphic to this finite resolution and is thus finite (since (*iii*) implies (*i*)). Assume now that X has non-zero homology only in a range of k degrees, with $k > 1$. Without loss of generality, assume that $H_n X = 0$ for $n < 0$ and $n \geq k$. Choose a finitely generated projective P and a map $P \to X$ inducing an epimorphism $P \to H_0 X$ with FP kernel K. Let X' be the cofibre of the map $P \to X$. Then one finds that $H_n X'$ is zero for $n \leq 0$ and $n \geq k$. Moreover, we have that $H_n X' = H_n X$ for $1 < n < k$ and that there is a short exact sequence $0 \to H_1 X \to H_1 X' \to K \to 0$. Since $H_1 X$ and K are FP, so is $H_1 X'$, and so by induction we can conclude that X' is finite. The complex P is certainly finite; therefore X is finite as well (since (*iii*) implies (*i*)). $\square$

Note 9.8. If R is a coherent ring over which every finitely presented module has finite projective dimension, then the above result simplifies. First, over such a ring, a module is finitely presented if and only if it is FP. Second, over a coherent ring, finitely presented modules form an abelian subcategory of the category of all modules, and this subcategory is closed under retracts and extensions. This allows one to show that a complex X is finite if and *only if* each $H_n X$ is finitely presented and only finitely many are non-zero. In addition, in this situation, the reliance on the result of Bökstedt and Neeman can be removed from the above proof. (See [7, Exercise I.2.11] for a brief discussion of coherence. Note that Noetherian rings are coherent.)

Proposition 9.6 implies that there is a set of isomorphism classes of finite objects and therefore that $\mathcal{D}$ has a phantom projective class (see Definition 5.2 and the subsequent discussion). The class $\mathcal{P}$ of projectives consists of all retracts of coproducts of finite objects, and we write $\mathcal{I}$ for the ideal of phantom maps. Since R is finite, the phantom projective class generates.

We recall a standard fact which is easily proved.

Proposition 9.9. *Let M and N be R-modules. Then*

$$\mathcal{D}(M, \Sigma^k N) \cong \mathrm{Ext}_R^k(M, N). \qquad \square$$

We can now prove one of the main results of this section.

Theorem 9.10. *Let M be a filtered colimit of FP modules and let N be any R-module. Then the phantom spectral sequence abutting to $\mathcal{D}(M, N)$ is the same as the spectral sequence* (9.4) *described in the previous part of this section. In particular, the filtrations agree:*

$$\mathcal{I}^l(M, \Sigma^k N) \cong \mathrm{P}^l \mathrm{Ext}^k(M, N).$$

When we say that the spectral sequences are the same, we mean that they are naturally isomorphic from the E_2-term onwards.

Proof. Let $\{M_\alpha\}$ be a filtered diagram of FP modules with a colimiting cone to M. Then, regarding these modules as complexes concentrated in degree zero, the cone is a minimal cone. Indeed, the cone from the complexes $\{M_\alpha\}$ to the complex M becomes a colimiting cone under $\mathcal{D}(\Sigma^n R, -) = H_n(-)$ for each n. And since filtered colimits are exact, one can use the five-lemma to show that it becomes a colimiting cone under $\mathcal{D}(W, -)$ for each finite W. This is what it means for the cone to be a minimal cone.

By Proposition 9.6, the complexes M_α are finite.

We saw in Theorem 5.10 that from a minimal cone on a filtered diagram of finite objects one can construct a phantom resolution. In fact, the construction corresponds exactly to the construction of a pure resolution of M in the previous part of this section. Moreover, to get the spectral sequence (9.4) we apply the functor $\mathrm{Ext}^*(-, N)$. To get the phantom spectral sequence we apply the functor $\mathcal{D}(-, N)_*$. By Proposition 9.9, these agree. Thus the spectral sequences agree. $\square$

The following lemma will allow us to find modules of pure projective dimension greater than one.

Lemma 9.11. *Let M be a flat R-module. Then any projective resolution of M is a pure projective resolution. Moreover, the projective dimension of M equals the pure projective dimension of M and for any N the natural map $\mathrm{PExt}^*(M, N) \to \mathrm{Ext}^*(M, N)$ is an isomorphism.*

Proof. A projective resolution of M is an exact sequence

$$0 \longleftarrow M \longleftarrow P_0 \longleftarrow P_1 \longleftarrow \cdots$$

with each P_i projective (and hence pure projective). We must show that this sequence is pure exact. Write M_i for the image of the map $P_{i-1} \leftarrow P_i$. For each right module E, the sequence $0 \leftarrow M \leftarrow P_0 \leftarrow M_1 \leftarrow 0$ is exact under $E \otimes -$ since M is flat. And since M and P_0 are flat, so is M_1. Thus we can inductively conclude that each sequence $M_i \leftarrow P_i \leftarrow M_{i+1}$ is pure exact. Therefore, the resolution is pure exact and we have proved the first part of the proposition.

From what we have proved, it follows that a projective resolution of M can be used to calculate both $\mathrm{PExt}^*(M, N)$ and $\mathrm{Ext}^*(M, N)$. Thus the natural map $\mathrm{PExt}^*(M, N) \to \mathrm{Ext}^*(M, N)$ is an isomorphism, and the projective and pure projective dimensions of M agree. $\square$

For the following theorem, we need to define the notion of a "pure complex".

Definition 9.12. *A complex X is **pure** if for each n the short exact sequences*

$$0 \longrightarrow B_n X \longrightarrow Z_n X \longrightarrow H_n X \longrightarrow 0$$

and

$$0 \longrightarrow Z_n X \longrightarrow X_n \longrightarrow B_{n-1} X \longrightarrow 0$$

are pure exact.

Theorem 9.13. *Let R be a coherent ring over which finitely presented modules have finite projective dimension.*

(*i*) *If every R-module has pure projective dimension less than n, then every pure complex has length at most n.*

(*ii*) *If every pure complex has length at most n, then every* flat *R-module has pure projective dimension less than n.*

(*iii*) *If every pure complex has phantom projective dimension at most one, then every R-module has pure projective dimension at most one.*

Proof. The proof of (*i*) is similar to the proof of Theorem 8.3. Note, however, that we do not claim that each pure complex has phantom projective dimension less than n. The correct statement along these lines must be made in the "pure" derived category, and will be described in a subsequent paper.

We proceed to prove (*ii*). Assume that there exists a flat module M of projective dimension at least n. (Recall that for flat modules, the pure projective and projective dimensions agree.) Then $\mathrm{Ext}^n(M, N)$ is non-zero for some N. But $\mathcal{I}^n(M, \Sigma^n N) = \mathrm{P}^n\mathrm{Ext}^n(M, N) = \mathrm{Ext}^n(M, N)$. Indeed, the first equality is Theorem 9.10 and the second is Lemma 9.11. (Note that the hypothesis of Theorem 9.10 is satisfied because we have assumed that every finitely presented R-module is FP.) Thus there are n phantoms in $\mathcal{D}$ which compose non-trivially, and the source is M, a pure complex.

Finally, we prove (*iii*). Let M be an R-module. By assumption there is a phantom exact sequence

$$0 \longrightarrow P \longrightarrow Q \longrightarrow M \longrightarrow 0$$

with P and Q retracts of sums of finite complexes. By Note 9.8, H_0P and H_0Q are finitely presented. And since a finitely presented module is finite when regarded as a complex concentrated in degree zero, it is easy to see that

$$0 \longrightarrow H_0P \longrightarrow H_0Q \longrightarrow M \longrightarrow 0$$

is a pure projective resolution of M. Thus, in this setting, every R-module has pure projective dimension at most one. □

Corollary 9.14. *If R is a coherent ring over which finitely presented modules are FP and whose derived category $\mathcal{D}$ is a Brown category (see Definition 5.3), then every R-module has pure projective dimension at most one.*

Proof. Suppose that $\mathcal{D}$ is a Brown category. Then, by Theorem 5.11, it follows that every complex has phantom projective dimension at most one. Thus, by Theorem 9.13 (*iii*), every R-module has pure projective dimension at most one. $\square$

It is proved in [21, Theorem 4.1.5] and in [35, Section 5] that when R is a countable ring, $\mathcal{D}$ is a Brown category. So we have shown in particular that a countable coherent ring over which finitely presented modules are FP has pure global dimension at most one. In fact, *any* countable ring has pure global dimension at most one [18, Proposition 10.5].

Example 9.15. Let k be an uncountable field. Write $k[x,y]$ for the polynomial ring and $k(x,y)$ for its field of fractions. A theorem of Kaplansky [25] states that the projective dimension of $k(x,y)$ as a $k[x,y]$-module is at least two. Moreover, $k(x,y)$ is a flat $k[x,y]$-module and $k[x,y]$ is a coherent ring of global dimension 2. Thus, by Theorem 9.13 (*ii*), in the derived category of $k[x,y]$-modules, $k(x,y)$ is the source of a non-zero composite of two phantoms. In particular, the derived category of $k[x,y]$-modules is not a Brown category. I learned this example from Neeman [35] and Neeman credits it to Bernhard Keller.

The phantom maps which compose non-trivially when there is a flat module of projective dimension greater than one can be made more explicit. Let M be a flat R-module of projective dimension at least n and let

$$0 \longrightarrow N \longrightarrow A_1 \longrightarrow \cdots \longrightarrow A_n \longrightarrow M \longrightarrow 0$$

be a pure exact sequence which is non-zero as an element of $\operatorname{Ext}^n(M,N)$. (This is possible by Lemma 9.11.) This sequence factors into n short exact sequences which are also pure exact. By Theorem 9.10 each of these represents a phantom map; their composite is the given non-zero element of $\operatorname{Ext}^n(M,N) = \mathcal{D}(M, \Sigma^n N)$.

References

[1] J. F. Adams. On the non-existence of elements of Hopf invariant one. *Ann. of Math. (2)*, 72:20–104, 1960.

[2] J. F. Adams. *Stable Homotopy and Generalised Homology*. University of Chicago Press, 1974.

[3] David Blanc. A Hurewicz spectral sequence for homology. *Trans. Amer. Math. Soc.*, 318:335–354, 1990.

[4] David Blanc. Operations on resolutions and the reverse Adams spectral sequence. *Trans. Amer. Math. Soc.*, 342:197–213, 1994.

[5] J. M. Boardman. Conditionally convergent spectral sequences. Homotopy invariant algebraic structures (Baltimore, MD, 1998), 49–84, *Contemp. Math.*, 239, Amer. Math. Soc., Providence, RI, 1999.

[6] Marcel Bökstedt and Amnon Neeman. Homotopy limits in triangulated categories. *Compositio Math.*, 86:209–234, 1993.

[7] Nicolas Bourbaki. *Commutative Algebra*. Hermann, 1972.

[8] Hans-Berndt Brinkmann. Relative homological algebra and the Adams spectral sequence. *Arch. Math.*, 19:137–155, 1968.

[9] J. Daniel Christensen and Neil P. Strickland. Phantom maps and homology theories. *Topology*, 37:339–364, 1998, math.AT/9802008.

[10] Samuel Eilenberg and J. C. Moore. Foundations of relative homological algebra. *Mem. Amer. Math. Soc.*, 55, 1965.

[11] A. D. Elmendorf, J. P. Greenless, I. Kriz, and J. P. May. Commutative algebra in stable homotopy theory and a completion theorem. *Math. Res. Lett.*, 1:225–239, 1994.
[12] A. D. Elmendorf, I. Kriz, M. A. Mandell, and J. P. May. *Rings, Modules, and Algebras in Stable Homotopy Theory*, volume 47 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1997. With an appendix by M. Cole.
[13] P. Gabriel and M. Zisman. *Calculus of Fractions and Homotopy Theory*. Springer-Verlag, 1967.
[14] Brayton Gray. Operations and a problem of Heller. Ph.D. thesis, University of Chicago, 1965.
[15] Brayton Gray. Spaces of the same n-type, for all n. *Topology*, 5:241–243, 1966.
[16] Brayton Gray and C. A. McGibbon. Universal phantom maps. *Topology*, 32:371–394, 1993.
[17] A. Grothendieck. Sur quelque points d'algèbre homologique. *Tôhoku Math. J.*, 9:119–221, 1957.
[18] L. Gruson and C. U. Jensen. Dimensions cohomologiques relieés aux foncteurs $\varprojlim^{(i)}$. In *Paul Dubreil and Marie-Paule Malliavin Algebra Seminar*, pages 234–294, 1981. Lecture Notes in Mathematics, vol. 867, Springer-Verlag.
[19] Michael J. Hopkins and Douglas C. Ravenel. Suspension spectra are harmonic. *Boletin de la Sociedad Matematica Mexicana*, 37:271–280, 1992. This is a special volume in memory of José Adem and is really a book. The editor is Enrique Ramírez de Arellano.
[20] M. Hovey, B. Shipley, and J. Smith. Symmetric spectra. Preprint, 1996.
[21] Mark Hovey, John H. Palmieri, and Neil P. Strickland. *Axiomatic Stable Homotopy Theory*, volume 610 of *Mem. Amer. Math. Soc.* 1997.
[22] Martin Huber and Willi Meier. Cohomology theories and infinite CW-complexes. *Comment. Math. Helv.*, 53:239–257, 1978.
[23] Dale Husemoller and John C. Moore. Differential graded homological algebra of several variables. *Symposia Mathematica*, IV:397–429, 1970.
[24] Birger Iversen. *Cohomology of Sheaves*. Universitext. Springer-Verlag, 1986.
[25] I. Kaplansky. Homological dimension of a quotient field. *Nagoya Math. J.*, 27:139–142, 1966.
[26] G. M. Kelly. Chain maps inducing zero homology maps. *Proc. Cambridge Philos. Soc.*, 61:847–854, 1965.
[27] S. Mac Lane. *Homology*. Springer-Verlag, 1975.
[28] L. Gaunce Lewis, J. Peter May, and M. Steinberger. *Equivariant Stable Homotopy Theory*, volume 1213 of *Lecture Notes in Mathematics*. Springer–Verlag, New York, 1986. With contributions by Jim E. McClure.
[29] T. Y. Lin. Adams type spectral sequences and stable homotopy modules. *Indiana Univ. Math. J.*, 25:135–158, 1976.
[30] T. Y. Lin. Duality and Eilenberg–Mac Lane spectra. *Proc. Amer. Math. Soc.*, 56:291–299, 1976.
[31] W.-H. Lin. The dual Adams spectral sequence. *Res. Rep. Nat. Sci. Council Math. Res. Center*, 9:249–264, 1980/81.
[32] Harvey R. Margolis. Eilenberg–Mac Lane spectra. *Proc. Amer. Math. Soc.*, 43:409–415, 1974.
[33] Harvey R. Margolis. *Spectra and the Steenrod Algebra*. North-Holland, 1983.
[34] Haynes R. Miller. Some algebraic aspects of the Adams-Novikov spectral sequence. Ph.D. thesis, Princeton University, 1974.
[35] Amnon Neeman. On a theorem of Brown and Adams. *Topology*, 36:619–645, 1997.
[36] Tetsusuke Ohkawa. A vanishing theorem of Araki-Yosimura-Bousfield-Kan spectral sequences. *Hiroshima Math. J.*, 23:1–14, 1993.
[37] Jean-Paul Pezennec. Propriétés topologiques de $[X, Y]$ et fantômes de finitude. *Bull. Soc. Math. France*, 107:113–126, 1979.
[38] Douglas C. Ravenel. Localization with respect to certain periodic homology theories. *Amer. J. Math.*, 106:351–414, 1984.
[39] Ravi Vakil. On the Steenrod length of real projective spaces. Submitted to the Gould Anniversary Volume of Discrete Mathematics, 1998.
[40] J.-L. Verdier. Catégories dérivées. In P. Deligne, editor, *Cohomologie Etale (SGA 4$\frac{1}{2}$)*, pages 262–311, 1977. Lecture Notes in Mathematics, vol. 569, Springer-Verlag.
[41] R. B. Warfield, Jr. Purity and compactness for modules. *Pacific J. Math.*, 28:699–719, 1969.
[42] C. A. Weibel. *An Introduction to Homological Algebra*, volume 38 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1994.
[43] Zen-ichi Yosimura. Hausdorff condition for Brown-Peterson cohomologies. *Osaka J. Math.*, 25:881–890, 1988.

This version has been updated compared to the published version:

The published version appeared in *Advances in Mathematics* 136 (1998), 284-339. Since then, the following changes have been made:

- In Proposition 5.6 (ii), the assumption that the objects in the diagram are projective was added, and two sentences were added at the end of the proof using this assumption. The discussion before Theorem 6.1 is adjusted accordingly.
- In Section 4, the assumption that the ideal $\mathcal{J}$ is closed under countable coproducts was added to Propositions 4.4 and 4.5. And in Section 7.3, the assumption that the forgetful functor U commutes with coproducts was added. The introduction was modified accordingly. Thanks to Ciprian Modoi and Ralf Meyer for pointing out this problem.
- My address and e-mail address have been updated.
- Reference [5] has been updated.
- Reference [8], which I was unaware of and should have cited, has been added.
- Various minor typos fixed.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WESTERN ONTARIO, LONDON, ONTARIO, CANADA
E-mail address: jdc@uwo.ca