EXTRAORDINARY HOMOTOPY GROUPS

ERIC PETERSON

Abstract. In this talk, we'll introduce the field of chromatic homotopy theory, which is where all the major advancements on the π_*^S problem have come from in the past 30+ years. Our express goal will be to study the Picard groups of the $K(n)$-localized stable categories, but to even make sense of that we'll have to go on a fairly deep safari through chromatic homotopy. Technical aspects will be mostly kept to a minimum, but some passing familiarity with algebraic geometry will help.

1. Introduction

In algebraic topology, we compute algebraic invariants of topological spaces. The most routinely useful of these invariants, called homology theories, satisfy a kind of locality: given a decomposition of the space in question, one should be able to reconstruct the invariant assigned to the total space by studying the invariants assigned to the pieces in the decomposition, together with the structure of the decomposition itself. It turns out that all invariants assigned in this way are "stable" in the following sense: a classical decomposition of a space is to turn it into a collection of sewn-together cells, which are unit disks in $\mathbb{R}^n$. From a space X so decomposed, we can build a new space ΣX which contains all the same cells and attaching data as that of X, but with all of the cell dimensions shifted up by 1. Then, the invariant assigned to ΣX is exactly that assigned to X, shifted up by one. Many interesting invariants fit into this framework, and the study of stable homotopy theory is the study of this collection of invariants and how they relate to one another.

The homotopy groups of spheres (i.e., the homotopy classes of maps $S^n \to S^m$) are a hugely intractable object. While it's likely we'll never have full information about them, the size of the understood sector is a good yardstick for progress made in the field. Over the past few decades, all of the major advances in this problem have come from a single place: chromatic homotopy theory. Our goal in this talk is to visit the major components of chromatic homotopy, driven by a particular application, and hopefully to spark some interest in the subject. The topology involved will be standard fare for these talks, but the algebraic geometry requirements will be much steeper than usual.

2. Morava's portrait

The starting point of this subject is with MU, the homology theory of complex bordism. While complex bordism has a geometric construction and interpretation in terms of bordism classes of stably almost-complex manifolds, we will be interested in an entirely different interpretation, uncovered in the '60s:

Theorem 1 (Quillen)**.** *The stackification of the groupoid-valued functor corepresented by the Hopf algebroid (MU_*, MU_*MU) is $\mathcal{M}_{FG}$, the moduli of commutative, one-dimensional formal Lie groups*[1]*. Moreover, the output MU_*X of the homology theory MU can be reinterpreted as a quasicoherent sheaf $\mathcal{MU}(X)$ on $\mathcal{M}_{FG}$.*

This theorem is remarkable on its own, in that it's surprising that something geometrically defined has a connection to such a strange algebraic object. What is truly inspiring is how good this connection is:

Theorem 2 (Adams, Novikov)**.** *There is a conditionally convergent spectral sequence $H^*(\mathcal{M}_{FG}; \mathcal{MU}(X)) \Rightarrow \pi_*X$. As $\mathcal{MU}(\mathbb{S})$ is the structure sheaf of $\mathcal{M}_{FG}$, there is a convergent spectral sequence $H^*\mathcal{M}_{FG} \Rightarrow \pi_*\mathbb{S}$.*

[1]A great many calculations in homotopy theory end up being organized by power series, and so it's important to have some geometric receptacle for them. This role is filled by formal affine n-space: $\hat{\mathbb{A}}^n := \operatorname{colim} \operatorname{Spec} R[x_1, \ldots, x_n]/\langle x_1^{m_1}, \ldots, x_n^{m_n}\rangle$, which has the property that $\operatorname{Hom}(\hat{\mathbb{A}}^n, \hat{\mathbb{A}}^m) = \{(f_1, \ldots, f_m) \mid f_i \in R[\![x_1, \ldots, x_n]\!]\}$. These appear naturally in algebraic geometry when completing the ring of functions of a variety at a smooth point. An abelian n-dimensional formal Lie group is an abelian group object structure on $\hat{\mathbb{A}}^n$; such things appear when completing an algebraic group at the identity, and the algebraic group's multiplication and inversion laws Taylor expand to appropriate power series satisfying the group axioms.

So, given this spectral sequence, one thing we might do to inform our understanding of the stable category is to study the structure of $\mathcal{M}_{FG}$. Any facts, qualitative or quantitative, about its cohomology will translate rather directly into statements about stable homotopy groups of spaces. Digging through the literature of arithmetic geometry, we find several results, provided we first localize at an odd prime p; from now on, everything will be assumed so local.

Theorem 3 (Cartier)**.** *There is*[2] *a unique closed substack $S_n \subseteq \mathcal{M}_{FG}$ for each nonnegative codimension n. (While we're at it, we notate $U_n = \mathcal{M}_{FG} \setminus S_n$ and $i_n : U_n \to \mathcal{M}_{FG}$ the inclusion.)*

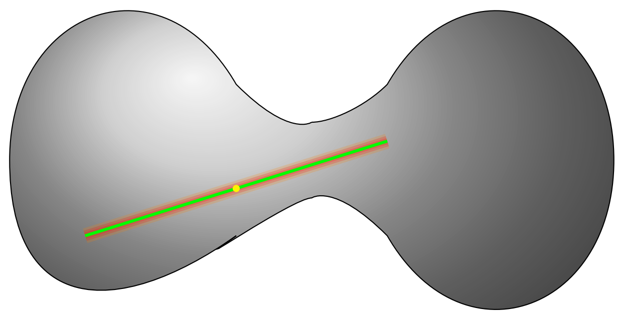

FIGURE 1. A diagram of $\mathcal{M}_{FG}$: the gray is S_0, the green is S_1, the yellow is S_2, and the red is the completion of U_2 along $U_2 \cap S_1$.

Moreover, this structure is reflected in the existence of several interesting cohomology theories:

Theorem 4 (Johnson, Morava, Ravenel, Wilson, etc.)**.** *There exist cohomology theories $P(n)$, $E(n)$, and $K(n)$ modelling*[3] *S_n, U_n, and the unique geometric point of $S_n \setminus S_{n+1}$ respectively.*

To understand the sheaf cohomology $H^*(\mathcal{M}_{FG}; \mathcal{MU}(X))$, then, we might employ this descending filtration of the base to divide the problem into easier bits. These spectra realize the various morphisms of (sub)stacks in play, but to make use of them we still need to understand sheaves supported on S_n and how to push and pull sheaves around. The first piece of this is fairly easy to describe:

Theorem 5 (Hopkins, Smith)**.** *Let C_n denote*[4] *the category of finite spectra X for which $K(n-1)_* X = 0$. This is closed under cofibrations, weak equivalences, and retracts. There is a tower*

$$\cdots \subseteq C_n \subseteq C_{n-1} \subseteq \cdots \subseteq C_1 \subseteq C_0.$$

In Morava's picture, $X \in C_n$ corresponds to a sheaf supported on S_n. Moreover, we have the following useful detection property: the ground ring of $K(n)$-homology is $K(n)_* = \mathbb{F}_p[v_n^{\pm}]$. A finite spectrum X is in C_n if and only if it has a map $f_n : X \to \Sigma^{|v_n|N} X$ such that $K(n)_*(f_n) = v_n^N \cdot -$ for some $N \in \mathbb{N}$.

Then, in the 70s, p-primary homotopy theory was coming into its own with giants like Serre who very successfully computed the homotopy of various spaces once all the q-torsion, $q \neq p$, was stripped out. A lot of effort then went in to actually building spaces modelling the resulting homotopy groups. This was originally accomplished by Sullivan, and a broad generalization was subsequently suggested by Adams: fixing

[2] Actually, this theorem and all the consequences we draw from it belong to the moduli of p-typical formal group laws. In topology, this is reflected by a multiplicative splitting $L_p MU \simeq \bigvee \Sigma^* BP$, where BP, short for Brown-Peterson, is yet another cohomology theory. The p-localized moduli of formal groups also splits into a sum of isomorphic summands, and the unshifted such summand is the moduli of p-typical formal groups. It is not worth paying attention to that distinction here.

[3] It's orthogonal to the purpose of this talk, but topologists may be comforted by seeing the coefficient rings of these spectra. The largest one is BP, with $\pi_* BP = \mathbb{Z}_{(p)}[v_1, \ldots, v_n, \ldots]$ with $|v_i| = 2(p^i - 1)$. For $P(n)$, we have $\pi_* P(n) = \mathbb{Z}_{(p)}[v_{n+1}, \ldots]$. The complement is $\pi_* E(n) = \mathbb{Z}_{(p)}[v_1, \ldots, v_{n-1}, v_n^{\pm}]$. Finally, the quotient field of the index-shifted intersection is $\pi_* K(n) = \mathbb{F}_p[v_n^{\pm}]$.

[4] This is also commonly written $\langle K(n-1) \rangle$, using the notation of Bousfield classes.

a homology theory E, to each space X there should exist a fiber sequence $G_EX \to X \to L_EX$, where G_EX is E-acyclic and $[Y, L_EX] = 0$ for any E-acyclic Y. Adams' original statement of this idea had serious set-theoretic difficulties, which were successfully sorted out by Bousfield. Bousfield localization interacts very well with Morava's picture:

Theorem 6. *The localization functor $L_{E(n)}$ corresponds to the functor $Ri_{n*} \circ i_n^*$ on sheaves. The localization functor $L_{K(n)}$ corresponds to the completion of U_n along $U_n \cap S_{n-1}$.*

Stable homotopy theory	Algebraic geometry
MU (or really BP)	$\mathcal{M}_{FG}$
L_pX	quasicoherent sheaf / complex of sheaves
L_pF for F finite	finite complex of q.c. sheaves
$L_p\mathbb{S}$	$\mathcal{O}_{\mathcal{M}_{FG}}$
homotopy groups	hypercohomology
$P(n)$	substratum $S_n \subseteq \mathcal{M}_{FG}$ of codimension n
$E(n)$	$U_n = \mathcal{M}_{FG} \setminus S_n$
$K(n)$	total quotient field of $U_{n+1} \cap S_n$
category C_n	subcategory of sheaves supported on S_n
functor $L_{E(n)}$	functor $Ri_{n*} \circ i_n^*$
functor $L_{K(n)}$	completion of U_{n+1} along $U_{n+1} \cap S_n$
$\vdots$	$\vdots$
smash product	tensor product
function spectra	sheaf RHom(A^*, B^*)
chromatic tower	the Cousin complex
$\vdots$	$\vdots$

FIGURE 2. The Morava-Hopkins dictionary.

3. Calculations of two Picard groups

The stable homotopy category is symmetric monoidal with respect to the smash product $\wedge$ and with unit $\mathbb{S}$, the sphere spectrum. One thing we like to do in unital monoidal categories is to compute their Picard groups, which for us is

$$\mathrm{Pic} = \{X \in \mathsf{Spectra} \mid \exists X^{-1} \in \mathsf{Spectra}, X \wedge X^{-1} \simeq \mathbb{S}\}.$$

The answer, it turns out, is familiar:

Theorem 7. $\mathrm{Pic} = \{\mathbb{S}^n \mid n \in \mathbb{Z}\} \cong \mathbb{Z}$.

Proof. Applying the Künneth formula to $X \wedge X^{-1} \simeq \mathbb{S}$, we know that $H_*(X) = \mathbb{Z}$ in some dimension $k \in \mathbb{Z}$, giving a homology isomorphism $\mathbb{S}^k \to X$ in the Postnikov tower of X. Similarly, there is a map $\mathbb{S}^{-k} \to X^{-1}$, and smashing through with $\Sigma^k X$ gives a second map $X \to \mathbb{S}^k$, and together the two split X as $X = \mathbb{S}^k \vee A$. Similarly, $X^{-1} = \mathbb{S}^{-k} \vee B$. We compute $\mathbb{S} = X \wedge X^{-1} = \mathbb{S} \vee \Sigma^{-k} A \vee \Sigma^k B \vee (A \wedge B)$, and so, as the stable homotopy groups of spheres are finitely generated, we must have $\pi_* A = \pi_* B = 0$, and hence $A = B = \mathrm{pt}$. □

We saw in the previous section that the filtration on $\mathcal{M}_{FG}$ gives rise to a number of interesting localizations of the whole stable category, and so we turn our attention to the categories of $K(n)$-local spectra. These categories also carry a symmetric monoidal structure: while given $K(n)$-local X and Y it is not guaranteed $X \wedge Y$ will be again $K(n)$-local, we can make it so by defining $X \wedge_{K(n)} Y := L_{K(n)}(X \wedge Y)$. We can then ask a similar question to before: what is Pic_n, the Picard group of the $K(n)$-localized stable category? First, we have the following general characterization:

Theorem 8. *For X a $K(n)$-local spectrum, X is invertible if and only if $\dim K(n)_* X = 1$.*

Proof sketch. The rightward implication is handled by $K(n)$-homology Künneth isomorphisms. For the leftward implication, we guess[5] the inverse $Y = F(X, L_{K(n)}\mathbb{S}^0)$ with its evaluation map $X \wedge Y \to \mathbb{S}^0$. One shows that this is an isomorphism on $K(n)$-homology, and hence an equivalence after $K(n)$-localization; this is argued by replacing $\mathbb{S}^0$ with an arbitrary spectrum Z and reasoning about the class of such Z for which such an isomorphism exists. □

Now, we need some spectra to work with:

Lemma 9. *Consider the cofiber sequence*

$$\mathbb{S}^0 \xrightarrow{p^j} \mathbb{S}^0 \to M^0(p^j)$$

defining the "Moore spectrum" $M^0(p^j)$. There are maps $M^0(p^j) \xrightarrow{i_{j,k}} M^0(p^k)$ for $j < k$ which limit to give a $K(n)$-local equivalence $\operatorname{colim}_j M^0(p^j) = M^0(p^\infty) \simeq \mathbb{S}^1$.

Proof. We compute $K(n)_* M^0(p^j)$ to have copies of $K(n)_*$ in degrees 0 and 1. Each stage in this system is an isomorphism on the top copy and the zero map on the bottom copy, and so in the limit we have $\dim K(n)_* M^0(p^\infty) = 1$, so certainly it is invertible. But, we can say more: by limiting over n through the defining cofiber sequences as n grows large, we produce a fiber sequence $\mathbb{S}^0 \to p^{-1}\mathbb{S}^0 \to M^0(p^\infty)$. The middle object is $K(n)$-acyclic, and hence $M^0(p^\infty) \to \mathbb{S}^1$ is a $K(n)$-local equivalence. □

Since we're coning off the v_0-self-map on $\mathbb{S}^0$ to build the Moore spectrum, they are members of C_1 and as such admit v_1-self-maps.

Construction 10. *Let $0 \leq \lambda_i < p$ be the digit sequence of some p-adic integer $a \in \mathbb{Z}_p^\wedge$ with truncations $a_n = \sum_{i=0}^n \lambda_i p^i$. Then, we form a sequence of spectra using the recipe:*

$$M^{-1}(p) \xrightarrow{v_1^{\lambda_0}} M^{-|v_1|a_0-1}(p) \xrightarrow{i_{1,2}} M^{-|v_1|a_0-1}(p^2) \xrightarrow{v_1^{p\lambda_1}} M^{-|v_1|a_1-1}(p^2) \xrightarrow{i_{2,3}} \cdots .$$

The homotopy colimit of this sequence is denoted $\mathbb{S}^{-|v_1|a}$.

Lemma 11. *The spectrum $\mathbb{S}^{-|v_1|a}$ is a $K(1)$-locally invertible spectrum. If $a \in \mathbb{Z}_p$ is an ordinary integer, then the limit of this system is $M^{-|v_1|a-1}(p^\infty) \simeq_{K(n)} \mathbb{S}^{-|v_1|a}$.*

Proof. Taking $K(1)$-homology interacts well with homotopy colimits, and so we can can actually calculate $\dim K(1)_* \mathbb{S}^{-|v_1|a} = 1$. Second, if a is an ordinary integer, then $\lambda_i = 0$ for large i, and the system degenerates to one we have already described. □

Theorem 12 (Hopkins, Mahowald, Sadofsky)**.** *This is essentially all of* Pic_1*; we also have to take into account shifts by the spheres $\mathbb{S}^1, \ldots, \mathbb{S}^{|v_n|}$, and so in total we have the computation* $\mathrm{Pic}_1 = \mathbb{Z}_p^\wedge \times \mathbb{Z}/(2p-2)$*. In general, a similar construction gives an embedding $\mathbb{Z}_p^\wedge \to \mathrm{Pic}_n$.*

Indication of proof. Define the Morava module[6] of X by $\mathcal{K}_{1,*}(X) := \lim_n (KU_p^\wedge)_*(X \wedge M^0(p^n))$. It turns out that $X \in \mathrm{Pic}_1$ if and only if $\mathcal{K}_{1,*}(X) \cong \mathcal{K}_{1,*}(\mathbb{S}^k)$ for some k. When $K(1)_*X$ is even-concentrated, we say X lies in the subgroup Pic_1^0, and we reduce to the case $k = 0$ when identifying its Morava module. For a topological generator $\gamma \in (\mathbb{Z}_p^\wedge)^\times$, we define $ev(X)$ to be the eigenvalue of the completed Adams operation ψ^γ on $\mathcal{K}_{1,0}(X)$, which fairly clearly determines an injective homomorphism $\mathrm{Pic}_1^0 \to (\mathbb{Z}_p^\wedge)^\times$. That ev is also surjective requires more work, but boils down to considering truncations of $ev(X)$ and building a "convergent sequence;" in the end this gives an isomorphism $\mathrm{Pic}_1^0 \cong (\mathbb{Z}_p^\wedge)^\times \cong \mathbb{Z}_p^\wedge \times \mathbb{Z}/(p-1)$. Lastly, the sequence $\mathrm{Pic}_1^0 \to \mathrm{Pic}_1 \to \mathbb{Z}/2$ is seen to not be split. □

[5] It's worth emphasizing that this really is a guess. There is not, to my knowledge, a general criterion ensuring that the dual of a dualizable spectrum is the inverse in the monoidal structure, but it is often true and so seems like a place to start.

[6] This is a cousin of Morava E-theory, if that helps.

4. Continuity of the Picard group

There are many features of this answer that are remarkable, but first and foremost we see that Pic_1 comes with a topology induced by[7] picking off the eigenvalue of ψ^γ, or by the embedding of $\mathbb{Z}_p^\wedge$, which is not something expected from the definition. In general, where does the topology on Pic_n come from? At integers n larger than 1, the description of $K(1)$ and $\mathcal{K}_{1,*}$ in terms of complex K-theory breaks down, and so we lose the Adams operations with which we previously made our analysis. One idea is to use the collection of (generalized) Moore spectra, by which the presently known elements of Pic_n are constructed, to filter the group. Though we haven't discussed this point here, this idea already surfaces in the construction of the Morava spectral sequence $H^*(\mathbf{S}_n; \mathcal{K}_{n,*}F) \Rightarrow \pi_* L_{K(n)} F$, which, among other things, encodes an algebraic approximation to the Picard group as computed through continuous group cohomology.

Question 13. *How can we describe the topology on* Pic_n *in general? How does it interact with these two known topologies, coming from* $\mathbb{Z}_p^\wedge \to \mathrm{Pic}_n$ *and from the algebraic Picard group?*

Given that the Picard group of the whole stable category is exactly the sphere spectra, which we use to define homotopy groups of spaces, for a $K(n)$-local space we might define $\pi_\lambda X = [\mathbb{S}^\lambda, X]$, where $\mathbb{S}^\lambda$ is the invertible spectrum corresponding to a point $\lambda \in \mathrm{Pic}_n$. Putting these two thoughts together, then, the natural question is: is the topology on Pic_n reflected in some kind of continuity of the functors $\pi_\lambda(-)$ as λ varies? Again taking a cue from arithmetic geometry, one thing we might try is to construct a single module πX over a ring Λ, to be thought of as $\Lambda \approx \mathbb{W}_{\mathbb{F}_{p^n}}[\![\mathrm{Pic}_n^*]\!]$, the completed monoid-ring of continuous $\mathbb{W}_{\mathbb{F}_{p^n}}$-valued functions on the topologized group Pic_n. Then, each element $\lambda \in \mathrm{Pic}_n$ should have a corresponding Λ-module $\tilde{\lambda}$ selecting it, such that $\pi X \otimes_\Lambda \tilde{\lambda} \cong \pi_\lambda X$.

This is a bit complicated, so it is worth saying what this would mean in the untopologized case: writing $G = \mathbb{Z}$ for the indexing set and $R = \mathbb{Z}$ for the ground ring, we assemble the homotopy groups of X into the single module $M = \bigoplus_{g \in G} \pi_g X$. Then, we have an action of $R[G^*]$ on M by the formula

$$(rf) \cdot \sum_{g \in G} p_g = \sum_{g \in G} r f(g) \cdot p_g,$$

where $f \in G^*$ is a set theoretic function $G \to R$. For any $g \in G$ we can produce an $R[G^*]$-module $\tilde{g} \cong R$ picking out the copy of $R \subseteq R[G^*]$ corresponding to the characteristic function on $\{g\}$. This gives $M \otimes_{R[G^*]} \tilde{g} = \pi_g X$.

This setup becomes nontrivial when G carries a topology and we restrict to considering only continuous functions $f : G \to R$. In particular, there is not a naïve construction of $\tilde{g}$, since characteristic functions are not continuous — instead, a bump function peaking at g maybe also collect some information from the homotopy groups nearby. The assertion, then, that $\pi X \otimes_{R[G^*]} \tilde{g} = \pi_g X$ becomes highly nontrivial.

Question 14. *Does there exist a ring* Λ*,* Λ*-modules* $\tilde{\lambda}$*, and an* Λ*-module* πX *such that the above is true?*

One key thing to note about this line of thought is that Λ will be a fairly reasonable ring from the perspective of commutative algebra. Compare this with the setting of classical homotopy theory: we have two basic ring actions on the graded groups $\pi_* X$: that of $\mathbb{Z}$ and that of $\pi_*\mathbb{S}$. In the first instance, there is basically nothing to say, since a $\mathbb{Z}$-module is exactly an abelian group. In the second instance, the ring $\pi_*\mathbb{S}$ is so miserable to an algebraist, with its immense amounts of torsion and nilpotence, that almost nothing can be inferred about structure inherited by $\pi_* X$. On the other hand, in this new setting, we replace the $\mathbb{Z}$-action on $\pi_* X$ by a Λ-action on πX. Unlike $\mathbb{Z}$, the ring Λ is quite complicated, and unlike $\pi_*\mathbb{S}$ it is quite nice, hence one could expect its structure to reveal a lot about the structure of πX. As a testing ground, this can be manually built to work in the case of Pic_1, where on its own it very neatly captures important qualitative descriptions connecting $\mathrm{im}\, J$ and $\pi_* L_{K(1)} \mathbb{S}^1$.

[7] Of course, it's not obvious that this is independent of γ, and there does not naïvely appear to be a canonical choice of γ.