The chromatic tower

Aaron Mazel-Gee

Abstract

Much of chromatic homotopy theory organizes around the chromatic tower, a tower of certain Bousfield localizations of a given spectrum; the chromatic convergence theorem asserts that the limit of the tower recovers the original spectrum in many cases. After reviewing this background, we'll discuss the proof of the chromatic convergence theorem and then examine the individual layers of the tower.

0 Notation

Everything is ∞-categorical, and $\mathcal{S}$ denotes the category of spectra.

1 Introduction

Recall that $\langle E(n)\rangle = \langle K(0) \vee \cdots \vee K(n)\rangle$; this implies that there are natural transformations $\mathrm{id} \to \cdots \to L_2 \to L_1 \to L_0$, and applying these to a space X to a space gives the *chromatic tower* for the space X. As the following theorem makes precise, this tower encodes a good deal of information about X.

Theorem 1 (Chromatic convergence). *If X is a finite p-local spectrum, then $X \xrightarrow{\sim} \lim L_n X$.*

Remark 2. The chromatic convergence theorem closely mirrors a corresponding phenomenon in algebraic geometry, namely that we can understand a sheaf on a variety via its behavior on a nested sequence of open subvarieties. In this analogy, the variety corresponds to the moduli stack $\mathcal{M}_{FG,p\text{-typ}}$ of p-typical formal group laws, and the open subvarieties correspond to the open substacks $\mathcal{M}^{\leq n}_{FG,p\text{-typ}}$ of formal groups with height less than n.

Remark 3. It's unambiguous to say "E-local finite spectrum" versus "finite E-local spectrum" (i.e. a complex built out of finitely many E-local spheres and disks) if and only if L_E is an arithmetic localization. Otherwise, $L_E S$ has new self-maps, which already give rise to "exotic Moore spectra".

Remark 4. The chromatic tower might be thought of as a "Postnikov-type tower", only with respect to a much subtler filtration than dimension, namely the *chromatic filtration*. In fact, for all $n < \infty$ the Bousfield class $\langle K(n)\rangle$ is minimal. This implies that the chromatic tower is unrefineable: there is no localization functor that sits between L_n and L_{n-1}.

It seems that we've all made our peace with the fact that stable homotopy theory is generally done "one prime at a time". If we're additionally willing to grant that rational homotopy theory is some sort of "0^{th} degree approximation" to stable homotopy theory (in the sense that $\pi_* S_{\mathbb{Q}}$ is concentrated in dimension 0), then the chromatic tower gives a maximal factorization of the rationalization map, and the convergence theorem may be interpreted as saying that this assignment is faithful on finite p-local spectra.

In this talk, we'll first discuss the proof of the chromatic convergence theorem and then explore in some depth the levels of the chromatic tower.

2 Proof of chromatic convergence

Reduction 1. *The class* `ChromConv` *of finite p-local spectra for which the chromatic convergence theorem holds is thick. Thus, to prove the chromatic convergence theorem it suffices to prove that $S_{(p)} \in$* `ChromConv`*.*

Proof. Clearly `ChromConv` is closed under weak equivalences.

To see that `ChromConv` is closed under taking cofibers, suppose $X, Y \in$ `ChromConv` and suppose we have some $X \xrightarrow{f} Y$. Then we have

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, "\wr"] & Y \arrow[r] \arrow[d, "\wr"] & \mathrm{cofib}(f) \arrow[d, dashed] \\ \lim L_n X \arrow[r, "\lim L_n f"'] & \lim L_n Y \arrow[r] & \mathrm{cofib}(\lim L_n f), \end{tikzcd}$$

and the dashed arrow is an equivalence by comparing long exact sequences in homotopy. Moreover,

$$\mathrm{cofib}(\lim L_n f) \simeq \Sigma\ \mathrm{fib}(\lim L_n f) \simeq \Sigma \lim \mathrm{fib}(L_n f) \simeq \lim \Sigma\ \mathrm{fib}(L_n f) \simeq \lim \mathrm{cofib}(L_n f) \simeq \lim L_n \mathrm{cofib}(f).$$

To see that `ChromConv` is closed under retracts, suppose that $A \hookrightarrow X$ is the inclusion of a retract and suppose that $X \in$ `ChromConv`. Since any functor preserves retracts, we have the diagram

$$\begin{tikzcd} A \arrow[r, hook] \arrow[d] & X \arrow[l, bend right] \arrow[d, "\wr"] \\ \lim L_n A \arrow[r, hook] & \lim L_n X \arrow[l, bend right] \end{tikzcd}$$

where both rows are retract diagrams. Choosing a homotopy inverse for the right vertical arrow and chasing the resulting diagram yields that $A \xrightarrow{\sim} \lim L_n A$. $\square$

Remark 5. Actually, we don't need to rely on the thick subcategory theorem to prove Reduction 1. The localizations L_n are smashing (although all localizations are smashing on for finite spectra anyways), and moreover smashing with finite spectrum commutes with limits. Of course, this is just the essence of the argument used above to show that `ChromConv` is closed under cofibers, which proves Reduction 1 by itself anyways.

Now, we will actually rephrase the chromatic convergence theorem slightly. Let us define the functor $C_n = \mathrm{fib}(\mathrm{id} \to L_n)$, so that for any X we have a commutative diagram of fiber sequences

$$\begin{tikzcd} \lim C_n X \arrow[r] \arrow[d] & \cdots \arrow[r] & C_2 X \arrow[r] \arrow[d] & C_1 X \arrow[r] \arrow[d] & C_0 X \arrow[d] \\ X \arrow[r, equal] \arrow[d] & \cdots \arrow[r, equal] & X \arrow[r, equal] \arrow[d] & X \arrow[r, equal] \arrow[d] & X \arrow[d] \\ \lim L_n X \arrow[r] & \cdots \arrow[r] & L_2 X \arrow[r] & L_1 X \arrow[r] & L_0 X \end{tikzcd}$$

Then we have the following reduction.

Reduction 2. *To prove that $S_{(p)} \in$ `ChromConv`, it suffices to prove that for all $m \geq 0$, $\{\pi_m C_n S_{(p)}\}_{n\geq 0} \to \{0\}_{n \geq 0}$ is a pro-isomorphism of towers of abelian groups.*

Proof. This pro-isomorphism implies that $\lim C_n S_{(p)}$ is weakly contractible, which implies that $S_{(p)} \xrightarrow{\sim} \lim L_n S_{(p)}$. $\square$

Remark 6. Note that a tower may be pro-trivial even if it is nontrivial at all levels. For instance, to prove that the tower of groups $\cdots \to A_2 \to A_1 \to A_0$ is pro-trivial it suffices to prove that for any n, the map $A_{n+s} \to A_n$ is trivial for all sufficiently large s. This is what we will show to hold for the tower $\{\pi_m C_n S_{(p)}\}_{n\geq 0}$.

We now make a further reduction. We take as a black box the following theorem, which marks the entrance of MU into the story.

Theorem 7. *The natural map $C_n S_{(p)} \to C_{n-1} S_{(p)}$ induces the zero map $MU_* C_n S_{(p)} \to MU_* C_{n-1} S_{(p)}$.*

Proof sketch. Writing L for the Lazard ring, inductively define the L-modules M_n by $M_{-1} = L_{(p)}$ and $M_n = v_n^{-1}M_{n-1}/M_{n-1}$. One computes directly that there is a canonical isomorphism $MU_*C_nS_{(p)} \cong M_n$, and that the induced map $MU_*C_nS_{(p)} \to MU_*C_{n-1}S_{(p)}$ is indeed zero. $\square$

Remark 8. In fact, this computation is roughly the starting point for the chromatic spectral sequence. We'll discuss this later.

To show why this implies the pro-isomorphism statement of Reduction 2, we introduce the following terminology.

Definition 1. Write $\overline{MU} = \mathrm{fib}(S \xrightarrow{\eta} MU)$. The natural map $\overline{MU} \xrightarrow{\varepsilon} S$ induces a sequence

$$\cdots \xrightarrow{\varepsilon\wedge \mathrm{id}_{\overline{MU}^{\wedge 3}}} \overline{MU}^{\wedge 3} \xrightarrow{\varepsilon\wedge \mathrm{id}_{\overline{MU}^{\wedge 2}}} \overline{MU}^{\wedge 2} \xrightarrow{\varepsilon\wedge \mathrm{id}_{\overline{MU}}} \overline{MU} \xrightarrow{\varepsilon} S.$$

Then, we define the *Adams–Novikov filtration* on π_*X, denoted

$$\cdots \subseteq \mathcal{F}_3\pi_*X \subseteq \mathcal{F}_2\pi_*X \subseteq \mathcal{F}_1\pi_*X \subseteq \mathcal{F}_0\pi_*X = \pi_*X,$$

by defining $\mathcal{F}_s\pi_*X$ to be the image of $\pi_*\overline{MU}^{\wedge s}\wedge X \xrightarrow{\pi_*\varepsilon^{\circ s}} \pi_*X$. (This is the filtration coming from the normalized MU-Adams resolution of X.)

Lemma 1. *If $X \xrightarrow{f} Y$ induces the zero map $MU_*X \xrightarrow{MU_*f} MU_*Y$, then $\pi_*X \xrightarrow{\pi_*f} \pi_*Y$ raises Adams–Novikov filtration degree.*

Proof. Suppose we are given the solid arrows in the diagram

$$\begin{tikzcd}
S^m \arrow[d, "x"'] \arrow[rr, "\overline{x}"] \arrow[ddr, dashed] & & & \\
X \arrow[dd, "f"'] & & \overline{MU}^{\wedge s}\wedge X \arrow[ll] \arrow[r] \arrow[dd] & MU\wedge\overline{MU}^{\wedge s}\wedge X \arrow[dd, "\mathrm{id}_{MU\wedge\overline{MU}^{\wedge s}}\wedge f"] \\
& \overline{MU}^{\wedge(s+1)}\wedge Y \arrow[dr] & & \\
Y & & \overline{MU}^{\wedge s}\wedge Y \arrow[ll] \arrow[r] & MU\wedge\overline{MU}^{\wedge s}\wedge Y.
\end{tikzcd}$$

We claim that $\pi_*(\mathrm{id}_{MU\wedge\overline{MU}^{\wedge s}}\wedge f)$ is the zero map, from which it follows that the dashed arrow exists. This is equivalent to the statement that the map $\overline{MU}^{\wedge s}\wedge X \xrightarrow{\mathrm{id}_{\overline{MU}^{\wedge s}}\wedge f} \overline{MU}^{\wedge s}\wedge Y$ is MU_*-null. From the MU_*-long exact sequence for $\overline{MU} \to S \to MU$, we see that $MU_*\overline{MU}$ is a free MU_*-module, and so we get Künneth isomorphisms $MU_*(\overline{MU}^{\wedge s}\wedge X) \cong (MU_*\overline{MU})^{\otimes s}\otimes_{MU_*} MU_*X$ and $MU_*(\overline{MU}^{\wedge s}\wedge Y) \cong (MU_*\overline{MU})^{\otimes s}\otimes_{MU_*} MU_*Y$. Since we assumed that $MU_*f = 0$, the claim follows. $\square$

This allows us the following further reduction.

Reduction 3. *To prove that $\{\pi_mC_nS_{(p)}\}_{n\geq 0} \to \{0\}_{n\geq 0}$ is a pro-isomorphism, it suffices to prove that the Adams–Novikov filtration on each $\pi_mC_nS_{(p)}$ is finite (i.e. there is some $s = s(m,n)$ such that $\mathcal{F}_s\pi_mC_nS_{(p)} = 0$).*

Proof. By Theorem 7, the image of $\pi_*C_{n+s}S_{(p)} \to \pi_*C_nS_{(p)}$ lands in $\mathcal{F}_s\pi_*C_nS_{(p)}$. Hence, the only endomorphism of the pro-object $\{\pi_*C_nS_{(p)}\}_{n\geq 0}$ is the zero map, which implies the claim. $\square$

To reduce this even further, we introduce a bit more terminology.

Definition 2. A map $X \xrightarrow{f} Y$ is called *phantom below dimension n*, or *n-phantom*, if for any finite spectrum F with $\dim F \leq n$, $[F,X] \xrightarrow{f\circ -} [F,Y]$ is zero. (So by definition, a phantom map is a map which is n-phantom for all n.)

Definition 3. X is called *MU-convergent* if for every n there is some $s = s(n)$ such that $\overline{MU}^{\wedge s} \wedge X \to X$ is n-phantom. (This implies that $\{\pi_* \overline{MU}^{\wedge s} \wedge X\}_{s \geq 0} \to \{\pi_* X\}_{s \geq 0}$ is a pro-isomorphism.) We will write MU-$\mathtt{Conv}$ for the the class of MU-convergent spectra.

We can now give our final reduction.

Reduction 4. *To prove that the Adams–Novikov filtration on each $\pi_n C_m S_{(p)}$ is finite, it suffices to prove that if X is connected (i.e. $\pi_i X = 0$ for $i \leq 0$), then $C_m X \in MU$-$\mathtt{Conv}$ for all m.*

Proof. Of course $S_{(p)}$ is connected, and so if $C_m S_{(p)} \in MU$-$\mathtt{Conv}$ and we choose $s = s(n)$ as in the definition of MU-convergence then we obtain that $\mathcal{F}_s \pi_n C_m S_{(p)} = 0$. □

Finally, we can prove the hypothesis of Reduction 4. It is an immediate consequence of the following three facts.

1. If X is connective (i.e. $\pi_i X = 0$ for $i < 0$), then $X \in MU$-$\mathtt{Conv}$.
2. MU-$\mathtt{Conv}$ is thick.
3. For any X, $L_n X \in MU$-$\mathtt{Conv}$.

We prove these facts in turn.

Fact 1. *If X is connective, then $X \in MU$-$\mathtt{Conv}$.*

Proof. Suppose X is connective, and choose any integer n. Since $\overline{MU}$ is connected, it follows that $\overline{MU}^{\wedge(n+1)} \wedge X$ is n-connected. So if F is finite and $\dim F \leq n$, then $[F, \overline{MU}^{\wedge(n+1)} \wedge X] = 0$, so certainly $[F, \overline{MU}^{\wedge(n+1)} \wedge X] \to [F, X]$ is zero. Thus we can take $s = n + 1$. □

To prove Fact 2, we need the following lemma.

Lemma 2. *If $X \to Y$ is n-phantom and W is connective, then $X \wedge W \to Y \wedge W$ is also n-phantom.*

Proof. Suppose that F is finite with $\dim F \leq n$ and that we are given a map $F \to X \wedge W$. If we write W as a filtered colimit of finite connective spectra W_α, then there exists a factorization

$$\begin{tikzcd} F \arrow[r] \arrow[dr, dashed] & X \wedge W \arrow[r] & Y \wedge W \\ & X \wedge W_\alpha \arrow[r] \arrow[u] & Y \wedge W_\alpha . \arrow[u] \end{tikzcd}$$

Thus, it suffices to show that $X \wedge W_\alpha \to Y \wedge W_\alpha$ is n-phantom; that is, we may additionally assume that W is finite. Now, since W is connective then $\dim DW \leq 0$, so $\dim DW \wedge F \leq n$. Thus in the commutative diagram

$$\begin{tikzcd} {[F, X \wedge W]} \arrow[r, phantom, "\cong"] \arrow[d] & {[DW \wedge F, X]} \arrow[d] \\ {[F, Y \wedge W]} \arrow[r, phantom, "\cong"] & {[DW \wedge F, Y]}, \end{tikzcd}$$

the right map is zero since $X \to Y$ is n-phantom, so the left map is also zero. □

Fact 2. *MU-$\mathtt{Conv}$ is thick.*

Proof. Clearly MU-$\mathtt{Conv}$ is closed under weak equivalences.

To see that MU-$\mathtt{Conv}$ is closed under cofibers, clearly it is closed under suspension, so it suffices to show that if $X \to Y \to Z$ is a cofiber sequence with $X, Z \in MU$-$\mathtt{Conv}$ then $Y \in MU$-$\mathtt{Conv}$. So, fix any integer n, and choose s such that $\overline{MU}^{\wedge s} \wedge X \to X$ and $\overline{MU}^{\wedge s} \wedge Z \to Z$ are both n-phantom. We will show that $\overline{MU}^{\wedge 2s} \wedge Y \to Y$ is

n-phantom. First, by Lemma 2, since $\overline{MU}$ is connective then $\overline{MU}^{\wedge 2s} \wedge Z \to \overline{MU}^{\wedge s} \wedge Z$ is also n-phantom. Now for any finite F with $\dim F \leq n$, applying $[F, -]$ to the diagram

$$\begin{tikzcd}
 & \overline{MU}^{\wedge 2s} \wedge Y \arrow[r] \arrow[d] & \overline{MU}^{\wedge 2s} \wedge Z \arrow[d] \\
\overline{MU}^{\wedge s} \wedge X \arrow[r] \arrow[d] & \overline{MU}^{\wedge s} \wedge Y \arrow[r] \arrow[d] & \overline{MU}^{\wedge s} \wedge Z \\
X \arrow[r] & Y &
\end{tikzcd}$$

yields the diagram

$$\begin{tikzcd}
 & {[F, \overline{MU}^{\wedge 2s} \wedge Y]} \arrow[r] \arrow[d] & {[F, \overline{MU}^{\wedge 2s} \wedge Z]} \arrow[d] \\
{[F, \overline{MU}^{\wedge s} \wedge X]} \arrow[r] \arrow[d] & {[F, \overline{MU}^{\wedge s} \wedge Y]} \arrow[r] \arrow[d] & {[F, \overline{MU}^{\wedge s} \wedge Z]} \\
{[F, X]} \arrow[r] & {[F, Y].} &
\end{tikzcd}$$

The middle row is exact, and both the top right vertical arrow and the bottom left vertical arrow are zero by assumption; hence, the claim that the composition of the middle two vertical arrows is indeed zero follows from an easy diagram chase. So $Y \in MU\text{-}\texttt{Conv}$.

To see that $MU\text{-}\texttt{Conv}$ is closed under retracts, suppose that $A \hookrightarrow X$ is the inclusion of a retract and suppose that $X \in MU\text{-}\texttt{Conv}$. Given any n, choose s such that $\overline{MU}^{\wedge s} \wedge X \to X$ is n-phantom. Then so is $\overline{MU}^{\wedge s} \wedge A \to A$, since if F is finite with $\dim F \leq n$ then we have the commutative diagram

$$\begin{tikzcd}
{[F, \overline{MU}^{\wedge s} \wedge A]} \arrow[r, hook] \arrow[d] & {[F, \overline{MU}^{\wedge s} \wedge X]} \arrow[d] \arrow[l, bend right] \\
{[F, A]} \arrow[r, hook] & {[F, X],} \arrow[l, bend right]
\end{tikzcd}$$

and a retract of a zero map is zero. So $A \in MU\text{-}\texttt{Conv}$. □

Fact 3. *For any X, $L_n X \in MU\text{-}\texttt{Conv}$.*

Proof. By the smash product theorem, the natural map $\{L_n X\}_{m \geq 0} \to \{\mathrm{Tot}^m E(n)^{\wedge(\bullet+1)} \wedge X\}_{m \geq 0}$ is an equivalence of towers. So there exists some m such that $L_n X \hookrightarrow \mathrm{Tot}^m E(n)^{\wedge(\bullet+1)} \wedge X$ is the inclusion of a retract. Since $MU\text{-}\texttt{Conv}$ is thick, it suffices to show that $\mathrm{Tot}^m E(n)^{\wedge(\bullet+1)} \wedge X \in MU\text{-}\texttt{Conv}$. This is defined as a finite limit of $E(n)$-modules, hence of MU-modules, so again since $MU\text{-}\texttt{Conv}$ is thick it suffices to show that if M is an MU-module, then $M \in MU\text{-}\texttt{Conv}$. In fact, we can see that $\overline{MU} \wedge M \to M$ is null via the diagram

$$\begin{tikzcd}
\overline{MU} \wedge M \arrow[r] & M \arrow[r] \arrow[dr, equal] & MU \wedge M \arrow[d] \\
 & & M,
\end{tikzcd}$$

so in the definition of MU-convergent we can take $s(n) = 1$ (since null implies n-phantom for all n). □

3 The levels of the tower

There are a number of interrelated ways to try to understand the levels of the chromatic tower.

3.1 The chromatic fracture square

First, the map $L_n X \to L_{n-1} X$ fits into a natural pullback square

$$\begin{tikzcd}
L_n X \arrow[r] \arrow[d] & L_{K(n)} X \arrow[d] \\
L_{n-1} X \arrow[r] & L_{n-1} L_{K(n)} X,
\end{tikzcd}$$

called the *chromatic fracture square*. This also mirrors a corresponding phenomenon in algebraic geometry, namely that there is a Mayer–Vietoris principle to reconstruct a sheaf from its restrictions to a formal neighborhood of a closed subvariety and its open complement along with the appropriate gluing data.

Remark 9. As a converse to the chromatic fracture square, it's not hard to see that if Y is $E(n-1)$-local, Z is $K(n)$-local, and we are given a map $Y \to L_{n-1} Z$, then $X = \lim(Y \to L_{n-1} Z \leftarrow Z)$ is $E(n)$-local (and the pullback square is indeed the chromatic fracture square for $X = L_n X$). This implies that we can recover the category $L_n \mathcal{S}$ via the pullback diagram

$$\begin{tikzcd}
L_n \mathcal{S} \arrow[r, "L_{K(n)}"] \arrow[d, "L_{n-1}(\eta_{L_{K(n)}})"'] & L_{K(n)} \mathcal{S} \arrow[d, "L_{n-1}"] \\
\mathtt{Arr}(L_{n-1}\mathcal{S}) \arrow[r, "\text{target}"'] & L_{n-1}\mathcal{S},
\end{tikzcd}$$

where $\mathtt{Arr}$ denotes the arrow category. (We're keeping track of the map $Y \to L_{n-1} Z$, and the choice of localization map $Z \to L_{n-1} Z$ is swept into the definition of a homotopy pullback.) One might say that this gives a "semi-orthogonal decomposition" $L_n \mathcal{S} \simeq L_{n-1}\mathcal{S} \rtimes L_{K(n)}\mathcal{S}$. All together, this gives us the diagram

$$\begin{tikzcd}
\mathcal{S}_{(p)}^{fin} \arrow[d, "\text{fully faithful (and essentially surjective???)}"] & & & \\
\lim L_n \mathcal{S}^{fin} \arrow[d] & & & \\
\vdots \arrow[d] & & & \\
L_n \mathcal{S}^{fin} \arrow[rr, "L_{K(n)}"] \arrow[dr, "L_{n-1}(\eta_{L_{K(n)}})"] \arrow[dd, "L_{n-1}"'] & & L_{K(n)}\mathcal{S}^{fin} \arrow[dr, "L_{n-1}"] & \\
& \mathtt{Arr}(L_{n-1}\mathcal{S}^{fin}) \arrow[rr, "\text{target}"'] \arrow[dl, "\text{source}"] & & L_{n-1}\mathcal{S}^{fin} \\
L_{n-1}\mathcal{S}^{fin} \arrow[rr, "L_{K(n-1)}"] \arrow[dr, "L_{n-2}(\eta_{L_{K(n-1)}})"] \arrow[dd, "L_{n-2}"'] & & L_{K(n-1)}\mathcal{S}^{fin} \arrow[dr, "L_{n-2}"] & \\
& \mathtt{Arr}(L_{n-2}\mathcal{S}^{fin}) \arrow[rr, "\text{target}"'] \arrow[dl, "\text{source}"] & & L_{n-2}\mathcal{S}^{fin} \\
L_{n-2}\mathcal{S}^{fin} \arrow[d] & & & \\
\vdots & & &
\end{tikzcd}$$

in which the parallelograms are pullbacks. (It seems plausible that $\mathcal{S}^{fin}_{(p)} \to \lim L_n \mathcal{S}^{fin}$ might be an equivalence, but a priori there might be some infinite p-local spectrum whose $E(n)$-localizations are all equivalent to $E(n)$-localizations of various finite spectra.)

3.2 The monochromatic and $K(n)$-local categories

If we define $M_n = \mathrm{fib}(L_n \to L_{n-1})$, then $M_n X$ is called the n^{th} *chromatic layer* of X. Since $L_n L_{K(n)} \simeq L_{K(n)}$, the right vertical arrow in the chromatic fracture square is actually a map in the chromatic tower for $L_{K(n)}X$; the fact that it's a pullback square implies that $M_n L_{K(n)} X \simeq M_n X$.

On the other hand, since $L_{K(n)} L_{n-1} \simeq \mathrm{pt}$ and $L_{K(n)} L_n \simeq L_{K(n)}$, then the fiber sequence obtained by applying $L_{K(n)}$ to the fiber sequence $M_n X \to L_n X \to L_{n-1} X$ yields that $L_{K(n)} X \simeq L_{K(n)} M_n X$.

So, $M_n L_{K(n)} \simeq M_n$ and $L_{K(n)} M_n \simeq L_{K(n)}$. It is in this sense that monochromatic layers and $K(n)$-localizations determine each other; indeed, $L_{K(n)} : M_n \mathcal{S} \leftrightarrows L_{K(n)} \mathcal{S} : M_n$ defines an equivalence of categories. (Note that M_n – while not a localization – is a smashing functor, since it's defined as the fiber of two smashing functors. It follows that it is indeed an idempotent functor.) In fact, with a little fussing one can even check that this is an monoidal equivalence. **[the "internal" homotopy groups are equivalent, but the internal ones in $L_{K(n)}\mathcal{S}$ are the same as the external ones, while the internal ones in $M_n\mathcal{S}$ are distinct.]**

Remark 10. Recall that $\pi_* S$ contains periodic families of elements coming from v_n-self maps of finite spectra of type n. In fact, this same method generalizes very cleanly to any $\pi_* M_n X$.

Let us write $I = (i_0, \ldots, i_{n-1}) \in \mathbb{N}^n$. By the periodicity theorem, for some cofinal set of $I \in \mathbb{N}^n$ there exist "generalized Moore spectra" $M(I) = M(p^{i_0}, \ldots, v_{n-1}^{i_{n-1}})$ (with top cell in degree 0), defined inductively by the cofiber sequence

$$M(p^{i_0}, \ldots, v_{n-1}^{i_{n-1}}) \to M(p^{i_0}, \ldots, v_{n-2}^{i_{n-2}}) \xrightarrow{v_{n-1}^{i_{n-1}}} \Sigma^{-i_{n-1}\cdot|v_{n-1}|} M(p^{i_0}, \ldots, v_{n-2}^{i_{n-2}}).$$

Then we can present the n^{th} monochromatic layer of X as $M_n X \simeq \mathrm{colim}_{I \in \mathbb{N}^n} L_n X \wedge M(I)$.

Now, if we are given some $\alpha \in \pi_m M_n X$, then by the small object argument there exists a factorization

$$\begin{tikzcd} S^m \arrow[rr, "\alpha"] \arrow[dr, dashed, "\widetilde{\alpha}"'] & & M_n X \\ & L_n X \wedge M(I) \arrow[ur, hook] & \end{tikzcd}$$

for some $I = (i_0, \ldots, i_{n-1})$. Hence for any v_n-self map $\Sigma^{i_n \cdot |v_n|} M(I) \xrightarrow{v_n^{i_n}} M(I)$ we can define a family of elements $\alpha_{s \cdot i_n} \in \pi_{m + s\cdot i_n \cdot |v_n|} M_n X$ (beginning with $\alpha_0 = \alpha$), where α_s is given as the composite

$$S^{m+s\cdot i_n \cdot |v_n|} \xrightarrow{\widetilde{\alpha}} \Sigma^{s \cdot i_n \cdot |v_n|} L_n X \wedge M(I) \xrightarrow{(v_n^{i_n})^{\circ s}} L_n X \wedge M(I) \hookrightarrow M_n X.$$

One can check that the maps $\Sigma^{s\cdot i_n \cdot |v_n|} L_n X \wedge M(I) \to L_n X \wedge M(I)$ are all equivalences (since $v_n^{i_n}$ is a $K(n)$-equivalence and the smash product of an $E(n)$-local spectrum with a finite spectrum of type n is $K(n)$-local); this doesn't imply that nontriviality of α guarantees nontriviality of the α_t, but at least it's some indication that this isn't too naive of an operation. Note also that by the asymptotic uniqueness and centrality of v_n-self maps, the α_t are asymptotically independent of the choices of I and $v_n^{i_n}$.

4 All the spectral sequences

There are a bunch of spectral sequences running around, all computing related things. In this section we attempt to organize them.

4.1 The chromatic spectral sequence

As we mentioned above, the computation that $MU_* C_n S_{(p)} \to MU_* C_n S_{(p)}$ is the zero map is roughly the starting point for the chromatic spectral sequence, aside from the small point that rather than apply MU_* to p-local spectra we simply apply BP_* instead.

To construct the chromatic spectral sequence, we will want to use the chromatic tower, but we must amend it slightly so that it gives a resolution of $S_{(p)}$ instead of $L_0 S_{(p)}$. Actually, so that we don't have to keep making exceptions at the low levels, we replace the bottom of the chromatic tower with

$$\begin{tikzcd}
 & \vdots \arrow[d] \\
M_2 S_{(p)} \arrow[r] & L_2 S_{(p)} \arrow[d] \\
M_1 S_{(p)} \arrow[r] & L_1 S_{(p)} \arrow[d] \\
 & L_0 S_{(p)} \arrow[d] \\
M_0 S_{(p)} \arrow[r] \arrow[ur, equal] & S_{(p)}/p^\infty \arrow[d] \\
 & \Sigma S_{(p)}
\end{tikzcd}$$

(Note that $M_0 S_{(p)} \simeq L_0 S_{(p)} \simeq S_{\mathbb{Q}} = p^{-1} S_{(p)}$.) For notational convenience, we denote the levels of this tower by $\mathcal{L}_n S_{(p)}$; that is, we set $\mathcal{L}_n S_{(p)} = L_n S_{(p)}$ for $n \geq 1$, and then we set $\mathcal{L}_0 S_{(p)} = S_{(p)}/p^\infty$ and $\mathcal{L}_{-1} S_{(p)} = S_{(p)}$. We retain the $M_n S_{(p)}$, although note that $M_1 S_{(p)} \to \mathcal{L}_1 S_{(p)} \to \mathcal{L}_0 S_{(p)}$ is *not* a fiber sequence.

We begin with the fact that the fiber sequence $S_{(p)} \to L_n S_{(p)} \to \Sigma C_n S_{(p)}$ induces a short exact sequence in BP_*, which is split for $n \geq 1$. Thus in that case, $BP_* L_n S_{(p)} \cong BP_* \oplus BP_* C_n S_{(p)}$, and in fact the aforementioned computation can be used to show that for $n \geq 2$, the map $BP_* L_n S_{(p)} \to BP_* L_{n-1} S_{(p)}$ is an isomorphism on the first summands and zero on the second summands. At the bottom, we have that $BP_* \mathcal{L}_0 S_{(p)} \cong BP_*/p^\infty$ and that $BP_* \mathcal{L}_{-1} S_{(p)} \cong \Sigma BP_*$, and that the maps $BP_* \mathcal{L}_1 S_{(p)} \to BP_* \mathcal{L}_0 S_{(p)}$ and $BP_* \mathcal{L}_0 S_{(p)} \to BP_* \mathcal{L}_{-1} S_{(p)}$ are both zero.

Now, returning to the modified chromatic tower, this all implies that for $n \geq 2$ the fiber sequence $\Sigma^{-1} L_{n-1} S_{(p)} \to M_n S_{(p)} \to L_n S_{(p)}$ induces a short exact sequence $0 \to \Sigma \widetilde{BP}_* \mathcal{L}_{n-1} S_{(p)} \to BP_* M_n S_{(p)} \to \widetilde{BP}_* \mathcal{L}_n S_{(p)} \to 0$ (where by an abuse of notation we write $\widetilde{BP}_* \mathcal{L}_n S_{(p)}$ to mean $BP_* C_n S_{(p)}$, despite the fact that the homology of a spectrum isn't ever really *un*reduced). However, letting $\widetilde{BP}_* \mathcal{L}_0 S_{(p)} = BP_* \mathcal{L}_0 S_{(p)}$ and $\widetilde{BP}_* \mathcal{L}_{-1} S_{(p)} = BP_* \mathcal{L}_{-1} S_{(p)}$ for notational convenience, we actually have this short exact sequence for all n.

Then, taking cohomology (i.e. Ext in BP_*BP-comodules) yields the exact couple

$$\begin{tikzcd}
\bigoplus_{i,j,n} H^{i,j}(\widetilde{BP}_* \mathcal{L}_n S_{(p)}) \arrow[rr, "{(i,j,n) \rightsquigarrow (i+1,j+1,n-1)}"] & & \bigoplus_{i,j,n} H^{i,j}(\widetilde{BP}_* \mathcal{L}_n S_{(p)}) \arrow[dl, "{(i,j+1,n-1) \rightsquigarrow (i,j,n)}"] \\
 & \bigoplus_{i,j,n} H^{i,j}(BP_* M_n S_{(p)}), \arrow[ul, "{(i,j,n) \rightsquigarrow (i,j,n)}"] &
\end{tikzcd}$$

and the associated spectral sequence – the *chromatic spectral sequence* – abuts to $H^{\star,\star}(\widetilde{BP}_* \mathcal{L}_{-1} S_{(p)})$, i.e. to the E_2-page of the Adams–Novikov spectral sequence for $\pi_* S_{(p)}$. (This exact couple gives the E_1-page of the chromatic spectral sequence, and so we have $d_r : E_r^{i,j,n} \to E_r^{i,j-r,n+r}$.)

4.2 Another route to $\pi_* S_{(p)}$

If we apply π_* to the chromatic tower, we get a spectral sequence $\pi_* M_* S_{(p)} \Rightarrow \pi_* S_{(p)}$. This also might reasonably be called the "chromatic spectral sequence" – though it isn't, so we'll refer to the usual chromatic spectral sequence as the *algebraic chromatic spectral sequence* and the present spectral sequence as the *topological spectral sequence* (even though these aren't exactly analogs of each other in the algebraic and topological worlds or anything like that).

To get at the topological chromatic spectral sequence, we have Adams–Novikov spectral sequences $H^{*,*}(BP_* M_n S_{(p)}) \Rightarrow \pi_* M_n S_{(p)}$. Note that the input to these is the same as the input to the chromatic spectral sequence.

Remark 11. Let E_n denote the Morava E-theory spectrum obtained from the height-n Honda formal group over $\overline{\mathbb{F}}_p$, and let $\mathbb{G}_n = \mathrm{Aut}(H_n) \rtimes \mathrm{Gal}(\overline{\mathbb{F}}_p/\mathbb{F}_p)$ denote the n^{th} extended Morava stabilizer group. By the Morava change-of-rings theorem, the Adams–Novikov spectral sequences for $M_n S$ and $L_{K(n)}S$ take the form

$$\begin{aligned} H^*_{cts}(\mathbb{G}_n; (E_n)_*/(p^\infty, \ldots, u_{n-1}^\infty)) &\Rightarrow \pi_* M_n S \\ H^*_{cts}(\mathbb{G}_n; (E_n)_*) &\Rightarrow \pi_* L_{K(n)} S. \end{aligned}$$

In fact, there is a topological refinement of the second statement: $\mathbb{G}_n$ can be made to act by $\mathcal{E}_\infty$-ring maps on E_n, and then the unit map $S \to E_n$ induces an equivalence $L_{K(n)}S \xrightarrow{\sim} E_n^{h\mathbb{G}_n}$. That is, the spectrum $L_{K(n)}S$ itself is the (hyper)cohomology of the group $\mathbb{G}_n$ with coefficients in the spectrum E_n. (Given this, the above spectral sequence for $\pi_* L_{K(n)}S$ can also be viewed as a descent spectral sequence.)

[mention Galois extensions?]