Enchiridion: Mathematical User's Guides Volume 1 (October 2015)
mathusersguides.com

A user's guide: Coassembly and the K-theory of finite groups

Cary Malkiewich

2. Metaphors and imagery

2.1. Spectra. It will be easiest if I begin with how to picture spectra. A spectrum is a sequence of based spaces X_0, X_1, X_2, … with structure maps $\Sigma X_{n-1} \to X_n$. We'll assume that the spaces are cell complexes, and the maps are closed inclusions.

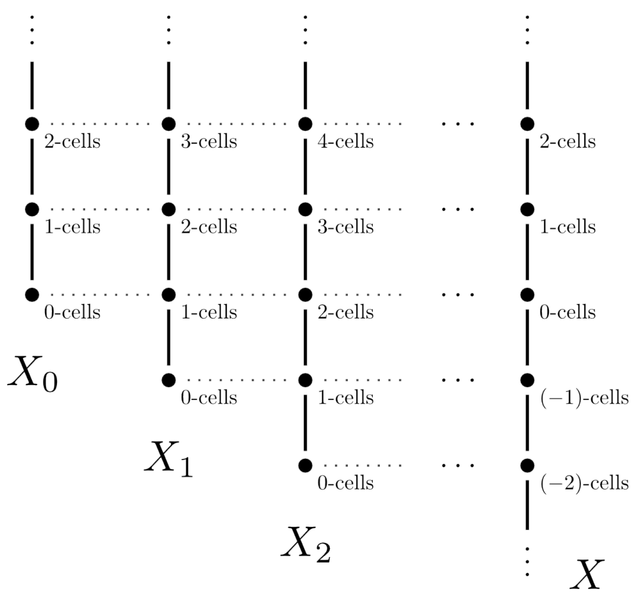

We arrange the spaces in a line as shown, and we think of X as their colimit. Their heights are staggered because each X_n is suspended before it is included into the next X_{n+1}.

We could say that X is made up of "elements," just like a set or an abelian group. An "element" of X is a point in the 0th space X_0, or a based loop

$S^1 \to X_1$, or a based sphere $S^n \to X_n$ for any n. We can always increment n, but we only care about the behavior as $n \to \infty$. So an element of X is a collection of maps $S^n \to X_n$, for all sufficiently large n, that agree along the structure maps of X. Picture a sequence of spheres S^n inside the spaces X_n, each one casting a shadow S^{n+1} in X_{n+1} that lines up with the next sphere.

These elements can be added together. Thinking of them as maps out of spheres, we add them by pinching the sphere at the equator. The choice of pinch $S^n \to S^n \vee S^n$ is not unique, but it is more or less equivalent to choosing two points in $\mathbb{R}^n$, where n can grow as large as we like. So we get a contractible space of such addition maps. Therefore the addition is commutative up to homotopy in a very strong sense – to be more precise, the elements form an E_∞ space.

Of course, you can have elements of other degrees too. An element of degree $k \in \mathbb{Z}$ is a collection of maps $S^{n+k} \to X_n$ for varying n; we may imagine these elements living in the kth row of the diagram above. The kth homotopy group $\pi_k(X)$ can be elegantly described as the homotopy classes of degree k elements. A map of spectra is an equivalence when it induces isomorphisms on these homotopy groups.

With this language of "elements," the sphere spectrum $\mathbb{S}$ is just a free spectrum on one element of degree 0. Similarly, the n-sphere $\Sigma^n\mathbb{S} = \mathbb{S}^n$ is freely generated by a single element of degree n.

Since spectra are similar to abelian groups, it makes sense to talk about ring spectra and module spectra. A ring spectrum R has a multiplication, which takes two elements of degrees k and ℓ as input and returns an element of degree $k + \ell$ as output. It also has a unit, which is just a degree 0 element of R. Similarly, if M is a module over R, that means that an element of degree k in R and an element of degree ℓ in M give an element of degree $k + \ell$ in M, and the unit of R acts as the identity. (Of course, multiplication should be associative and distribute over addition, and indeed this happens up to a contractible set of choices.) Every ordinary ring R becomes a ring spectrum HR, which only has interesting elements in degree zero, given by the actual elements of R itself.

If we fix a ring spectrum R then it is easy to build some modules over it. I can take a finite number of copies of R, multiply them by discs, and glue them together along the boundaries of these discs, to make something like a cell complex. Then, if I want, I could cut down to some smaller module spectrum sitting inside the bigger one as a retract. Every module built this way is "perfect." This is the analogue of being a module that is finitely-generated and projective. Each cell in my complex plays the role of a generator or a relation, or sometimes a little bit of both. If I allowed infinitely many cells, then I could capture every R-module this way, up to equivalence.

Finally, if R is a ring spectrum and G is a topological group or monoid, there is a group ring spectrum $R[G]$, generated freely by R and by a degree 0 element for every point of G. Formally, the nth level of $R[G]$ is just $R_n \wedge G_+$.

As we mentioned in the first part of the user's guide, these group rings are closely connected to the study of parametrized spectra, or families of spectra that vary continuously over a base space B. Just picture a fiber bundle over B, each fiber of which is one of these spectra. The main idea is that spaces over B are essentially the same as spaces with an action of ΩB, so the category of parametrized spectra is essentially the same as the category of module spectra over the ring spectrum $\mathbb{S}[\Omega B]$.

2.2. Algebraic K-theory. In [**Mal15**] we study the algebraic K-theory of a ring or ring spectrum R. This is a sequence of abelian groups $K_n(R)$, which are the homotopy groups of a spectrum $K(R)$. So we can think of $K_n(R)$ as the degree n "elements" of the spectrum $K(R)$, in the sense we discussed in the last section. Let's discuss what these elements look like.

We begin with the first space $K(R)_0$ in the spectrum $K(R)$. If R is a discrete ring, then $K(R)_0$ is the moduli space of all finitely-generated projective R-modules. A point in this space is an R-module, and a path between two points is an isomorphism of modules. Similarly, if R is a ring spectrum, then $K(R)_0$ is the moduli space of all perfect R-modules. A point is a perfect R-module, and a path between points is a weak equivalence of modules. Already, we can conclude that each perfect R-module M gives us some degree zero element $[M]$ of the spectrum $K(R)$.

When I move up to the space $K(R)_1$, each of these R-modules is now given by a based loop $S^1 \to K(R)_1$. Each isomorphism or weak equivalence of modules is now a homotopy of loops. But $K(R)_1$ is more than $\Sigma K(R)_0$, and the extra points do something neat. If I have a cofiber sequence of modules

$$A \longrightarrow X \longrightarrow X/A,$$

then each of the modules A, X, and X/A is a loop in $K(R)_1$. I add in a 2-simplex so that the composite of the loops for A and X/A is homotopic to the loop for X:

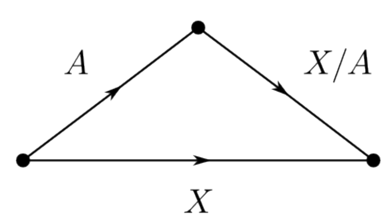

Then I add some higher-dimensional simplices in a similar way, to make the rest of $K(R)_1$. (If you want more detail, I'm thinking of Waldhausen's $\mathcal{S}_\cdot$-construction [**Wal85**].) As a result, in the spectrum $K(R)$, the elements $[A]$ and $[X/A]$ now sum to $[X]$, up to homotopy. This is *additivity*, the fundamental property of algebraic K-theory.

Additivity gives conceptual meaning to the addition of elements of $K(R)$: now $[A]+[B]$ is equivalent to $[A \vee B]$, and to any $[X]$ with a cofiber sequence $A \to X \to B$. This means that $K(R)$ is a certain kind of "group completion" of the space of R-modules.

I can repeat this process in a reasonable way to get $K(R)_2$ from $K(R)_1$, and so on. It turns out that after $K(R)_1$, nothing else interesting happens to our elements, so I can read off all the K-theory groups by taking the homotopy groups of the space $\Omega K(R)_1$.

There is one final important point. This $\mathcal{S}.$ construction on the category of perfect R-modules builds a tremendously large space. You could say this is the reason why K-theory groups are hard to compute. We don't have a nice, small model to calculate them, as we do in more basic problems in algebraic topology, like computing the homology of a torus.

2.3. Assembly and coassembly. Now that we know what $K(R)$ is, we'll try to understand how $K(R)$ is connected to $K(R[G])$ when G is a group. The relationship between them is the *assembly map*

$$BG_+ \wedge K(R) \longrightarrow K(R[G])$$

We picture the left-hand side as being much smaller than the right-hand side, though in truth it is still quite big.

There are two ways to understand the assembly map, and in the paper we explicitly prove that they give the same thing. The first and more classical perspective is this. Recall that a G-torsor is a left G-space that is isomorphic to G itself. We can interpret BG as the moduli space of all G-torsors, and $K(R)_0$ as the moduli space of R-modules. Given an R-module M and G-torsor $\widetilde{G}$, the tensor product

$$M \wedge \widetilde{G}_+ = \bigoplus^{\widetilde{G}} M$$

is a module over $R[G]$. This gives the assembly map $BG_+ \wedge K(R) \longrightarrow K(R[G])$ at spectrum level 0. It extends in a reasonable way to the higher levels of the spectrum as well. And if we want to think in terms of matrices, rather than modules and torsors, the assembly map sends a matrix $A \in GL_n(R)$ and an element $g \in G$ to a matrix in $GL_n(R[G])$ where every entry of A has been multiplied by g.

The second perspective is more homotopy-theoretic. Suppose F is a functor from unbased spaces to spectra. Each point of X gives a map $* \to X$. Apply F to get a map $F(*) \to F(X)$. When F is nice (either topological or a homotopy functor) then these maps $F(*) \to F(X)$ "assemble" together into a map

$$X_+ \wedge F(*) \xrightarrow{\alpha} F(X)$$

Why does this apply to $K(R[G])$? Remember that $R[G]$-modules are bundles of R-modules over BG. So I can think of $K(R[G])$ as a group-completed moduli space, the space of bundles of R-modules over BG, which came from perfect $R[G]$-modules. This extends to a functor on other spaces too: for each space X we take the bundles of R-modules over X. Given a map $X \to Y$, and a bundle E of modules over X, I can extend E to $\widetilde{E}$ over Y while keeping its total space weakly equivalent to that of E. This rule allows us to define an assembly map as above.

We can see why this gives the same assembly map. If M is an R-module, then for any point $b \in BG$, I can try to make a bundle whose total space is equivalent to just M sitting above the point b. The way to do this is to thicken the map $* \to BG$ to the bundle $EG \to BG$, then multiply by M to get the bundle $M \times EG$. If we let the point b vary, then EG itself does not change, but our choice of basepoint $* \in EG$ does vary. If b goes around a closed loop corresponding to $g \in G$, the basepoint of EG changes by multiplication by G. Imagine pulling $\bigoplus^G M$ around this loop in the bundle, and when we get back to the basepoint, we end up with a map

$$\bigoplus^G M \longrightarrow \bigoplus^G M$$

which multiplies on the left by G. So our bundle has fiber $\bigoplus^G M$ and monodromy given by left multiplication. It comes from the $R[G]$-module $\bigoplus^G M$, which is exactly the module we described in the classical version of the assembly map.

This entire discussion can be dualized, and the dual story is shorter. Any bundle over BG can be restricted to a smaller subspace, giving another bundle. This allows us to define a *coassembly map*

$$K(R[G]) \longrightarrow \mathrm{Map}_*(BG_+, K(R))$$

Each point in $K(R[G])$ is a bundle of modules E over BG. Given such a bundle, and a point $b \in BG$, we take the fiber E_b. That is the complete mental picture of the coassembly map.

Unfortunately we are lying a bit, because $K(R[G])$ isn't really a contravariant functor. The problem is with the finiteness conditions: the fiber E_b is certainly an R-module, but it may not be perfect. So this coassembly map isn't always defined.

If G happens to be a finite group, then E_b is finite, so coassembly is defined. More generally, coassembly is defined on those bundles of modules whose *fibers* are finite. So, we are studying $R[G]$-modules whose *underlying* R-module is finite. The K-theory of such things is the Swan theory $G^R(R[G])$. There is always a coassembly map

$$G^R(R[G]) \longrightarrow \mathrm{Map}_*(BG_+, K(R))$$

When G is finite, every finite $R[G]$-module is a finite R-module, so $K(R[G])$ includes into $G^R(R[G])$. This inclusion goes under the fancy name of the *Cartan map*.

2.4. The transfer and the norm. The main theorem of [**Mal15**] identifies the composite of the assembly and coassembly maps as a norm, so we end by painting a picture of transfer and norm maps.

Suppose we have two spectra X and Y, and a map $X \to Y$ whose fibers are finite sets. We would like to define a map $Y \to X$ going the other way. It would be nice if we could continuously choose, for each element of Y, some preimage in X. But this is usually impossible; most covering spaces do not admit a section. So instead, we take each element of Y to the sum of all its preimages in X. This eliminates all choices, and it is possible because the elements can be added together in a somewhat commutative way.

More specifically, suppose G is a finite group, X is a spectrum with a free left G-action (that is, free away from the basepoint), and $Y = X_G$ is the orbits. I "define" the transfer $Y \to X$ by taking each element of Y to the sum of its G preimages in X. We have to be careful, because the sum of elements is only defined up to a choice of point in some contractible space. In particular, we have a rule for how to sum together $|G|$ different maps, each time I pick an embedding of G into $\mathbb{R}^\infty$. So for each element $y \in Y$, I should choose some embedding $p^{-1}(y) \to \mathbb{R}^\infty$ to define this sum, but I have to make these choices in a continuous way. In particular, if I pass around a loop based at y, the set of embedded points will move through $\mathbb{R}^\infty$and come back to itself, but each point x inside the set will travel to gx, for some fixed $g \in G$.

So my rule associates to each $y \in Y$ the G-fold sum of its preimages in X, but this rule has a twisting, or monodromy, as I rove around Y. It seems that I could just pick the rule once, and then change it by a G-action as I rove around Y. So let's pick a map of spectra

$$p : \mathbb{S} \longrightarrow \Sigma^\infty_+ G$$

It should be equivariant with respect to G acting on the right, and it should send the degree 0 generator on the left to the sum of the $|G|$ distinct degree 0 generators on the right. This isn't quite possible on the nose, but it becomes possible if we allow ourselves to replace $\Sigma^\infty_+ G$ up to equivalence. Anyway, this map p is called the *pretransfer*. Once I have the pretransfer, I use it on every point of Y, only changing it by multiplication by G as I rove around Y. In other words, I smash the pretransfer with the identity map of X, and divide out by the G-action:

$$Y \cong \mathbb{S} \wedge_G X \longrightarrow \Sigma^\infty_+ G \wedge_G X \cong X$$

The resulting map is the *equivariant transfer*.

If instead $X \to Y$ were a covering space with fiber $\underline{n} = \{1, \ldots, n\}$ (away from the basepoint), then I could express

$$X = \widetilde{X} \wedge_{\Sigma_n} \underline{n},$$

where $\widetilde{X}$ has a free Σ_n-action and $Y = \widetilde{X}_{\Sigma_n}$. Then I could do essentially the same thing as above, except that the pretransfer is a Σ_n-equivariant map $\mathbb{S} \to \Sigma^\infty_+ \underline{n}$. This recipe gives the classical transfer map.

Returning to the case where the fiber of $X \to Y$ is G, we observe that if we sum up all the points in a single orbit of X, that sum should really be fixed by the action of G. So the transfer map $Y \to X$ actually factors through the homotopy fixed points X^{hG}. The resulting map

$$X_{hG} \longrightarrow X^{hG}$$

is the *equivariant norm map*. It takes each orbit of X to the sum of the points in that orbit, regarded as a point of X which is fixed under the G-action.

Let's tie the pictures from the last four sections together, to see why our theorem should be true. If I apply the assembly and coassembly maps to $BG_+ \wedge K(R)$, the element $[M]$ is sent to $[\bigoplus^G M]$. But this agrees with the G-fold sum of the element $[M]$, by the additivity of K-theory. The transfer and norm maps are given by a similar sum, so we are led to guess that assembly and coassembly give some sort of transfer. This is more than just intuition; it is a rough outline of the proof!

References

[Mal15] C. Malkiewich, *Coassembly and the K-theory of finite groups*, arXiv preprint arXiv:1503.06504 (2015).

[Wal85] F. Waldhausen, *Algebraic K-theory of spaces*, Algebraic and geometric topology, Springer, 1985, pp. 318–419.

Dept of Mathematics, University of Illinois at Urbana-Champaign, Urbana, IL 61801

E-mail address: cmalkiew@illinois.edu