An Introduction to Spectral Sequences

Matt Booth

December 4, 2016

This is the second half of a joint talk with Tim Weelinck. Tim introduced the concept of spectral sequences, and did some informal computations, including a spectral sequence proof of the Snake Lemma. Now it's time to make some of the ideas more rigorous, especially the intuitive notion of 'convergence'. Hopefully these notes hold up reasonably as a standalone work. We'll follow Weibel in [2].

1 Technicalities

Before getting into some of the details, it's important to be aware that there are two different conventions for spectral sequences: we can either index them homologically, or index them cohomologically. Of course these two notions are the same up to some sign conventions, but sign conventions are the bane of every homological algebraist. Vakil uses cohomological indexing, but Weibel prefers homological indexing.

Definition 1.1. Let $\mathcal{A}$ be an abelian category. A **homologically indexed spectral sequence** E is a collection of objects E^r_{pq} of $\mathcal{A}$ for $r \geq m \in \mathbb{N}$ and $p, q \in \mathbb{Z}$ and maps $d^r_{pq} : E^r_{pq} \to E^r_{p-r,q+r-1}$. If we fix an r, the doubly $\mathbb{Z}$-graded object E^r is referred to as a **page** (or **sheet**). Usually $m \in \{0, 1, 2\}$. The objects and maps are required to satisfy:

i) The d^r are differentials. These turn each page E^r into a collection of chain complexes that have 'slope' $-\frac{(r+1)}{r}$.

ii) $E^{r+1}_{pq} \cong H_{pq}(E^r)$.

We see that we can recover the whole of the spectral sequence from just the E^m page, since we can recover the E^{i+1} page from the E^i page by taking homology. It's useful to draw out the first few pages here: the differentials d^0 go down and the d^1 go to the left. The d^2 differentials carry out a Knight's move - they go two steps left and one step up. The higher differentials act like generalised Knight's moves.

The **total degree** of the term E^r_{pq} is $p+q$, so that the differentials always decrease the total degree by 1. We may also define a **cohomologically indexed spectral sequence** E^{pq}_r to simply be a homological spectral sequence E^r_{pq} with the terms reindexed as $E^{pq}_r = E^r_{-p,-q}$. We'll always indicate whether a spectral sequence is homological or cohomological by the placement of the pq subscript.

A homological spectral sequence E is said to be **first quadrant** if E^r_{pq} is only nonzero whenever (p,q) is in the first quadrant (i.e. $p \geq 0$ and $q \geq 0$). Let E^r_{pq} be a first quadrant spectral sequence, fix p and q, and consider the sequence $\{E^r_{pq} : r \geq m\}$. We see that eventually this sequence must stabilise, since the differentials entering and leaving E^r_{pq} become the zero map. Call this stable value E^∞_{pq}. We say that the spectral sequence E^r_{pq} **abuts to** E^∞_{pq}.

Remark 1.2. If $\mathcal{A}$ satisfies Grothendieck's axioms AB4 ($\mathcal{A}$ is cocomplete, and the coproduct of monomorphisms is a monomorphism) and AB4* ($\mathcal{A}$ is complete, and the product of epimorphisms is a epimorphism) then we can define the E^∞ page for any spectral sequence. These conditions hold for example if $\mathcal{A}$ is **mod-**R for some ring R, but not necessarily if $\mathcal{A}$ is a category of sheaves. However, we'll be interested only in first quadrant spectral sequences.

Definition 1.3. A first quadrant homological spectral sequence E is said to **converge** to a graded object H_*, denoted by $E^r_{pq} \Rightarrow H_{p+q}$, if for each n we have a finite filtration

$$H_n = F_t H_n \supseteq F_{t-1} H_n \supseteq \cdots \supseteq F_s H_n = 0$$

such that the quotients $F_p H_{p+q} / F_{p-1} H_{p+q}$ are isomorphic to E^∞_{pq}.

What's the intuition behind this definition? Certainly we could simply say that E converges to $\bigoplus_{p+q=n} E^\infty_{pq}$ and be done with it. But this is too restrictive - really we should only be able to compute the H_n up to extension. For example, if we have a first quadrant spectral sequence with $E^\infty_{10} = \mathbb{Z}/2 = E^\infty_{01}$, this does not determine whether H_1 'should' be $\mathbb{Z}/4$ or the Klein 4-group. So the same spectral sequence can converge to two different things. However there are theorems (e.g. [2], 5.2.12) that can help establish some form of uniqueness.

Definition 1.4. A spectral sequence E^r_{pq} is said to **collapse** at M if there is at most one nonzero row or column in E^M_{pq}, and all differentials are zero (this is automatic if $M \geq 2$).

Many spectral sequences found 'in the wild' collapse at the E^1 or E^2 pages. If a spectral sequence collapses, then we can read off its limit: it's clear that $E^\infty_{pq} = E^M_{pq}$, so H_n is simply the unique E^M_{pq} with $p + q = n$. In this case the limit is unique. It's clear that collapse is not a necessary condition for us to be able to 'read off the unique limit': it's sufficient that every diagonal has at most one nonzero entry.

If we want to get off the ground with proving things, we'll need a couple of convergence results. Here's the most important:

Theorem 1.5 ([2], 5.6.1)**.** *Let C be a first-quadrant double complex. We can define a spectral sequence ${}^\vee E$ with ${}^\vee E^0_{pq} = C_{pq}$ with the obvious vertical differentials. Then ${}^\vee E$ converges to the homology $H_*(C)$ of the total complex of C.*

Observe that in the situation above we can also define another spectral sequence ${}^< E$ with ${}^< E^0_{pq} = C_{qp}$ (note the index flip!) and horizontal differentials.

Proposition 1.6 ([2], 5.6.2)**.** *Let C be a first-quadrant double complex. Then ${}^< E$ also converges to $H_*(C)$.*

Remark 1.7. The proof is exactly the same as the proof of Theorem 1.5.

So now we have two ways to compute the same thing: if C is a double complex then we may compute $H_*(C)$ by the two different spectral sequences ${}^{\vee}E$ and ${}^{<}E$. This often provides interesting results.

Remark 1.8. Note that a double complex is the same thing as a graded complex. We can play the same game above if the complex C is only filtered, not graded. Of course, things get more complicated!

2 Example: Balancing of Tor

If R is a ring and M, N are (right, left) R-modules, we can define two functors

$$\mathrm{Tor}_*^R(M,-) := \mathbf{L}_*(M \otimes_R -) : R\text{-}\mathbf{mod} \to \mathbf{Ab}$$
$$\mathrm{Tor}_*^R(-,N) := \mathbf{L}_*(- \otimes_R N) : \mathbf{mod}\text{-}R \to \mathbf{Ab}$$

To say that Tor is **balanced** is to say that $\mathrm{Tor}_*^R(M,-)(N) \cong \mathrm{Tor}_*^R(-,N)(M)$, naturally in M and N. There's an easy proof of this using spectral sequences:

Take projective resolutions $P_\bullet \to M$ and $Q_\bullet \to N$, and tensor them together to get a first-quadrant double complex $P_\bullet \otimes_R Q_\bullet$. From $P_\bullet \otimes_R Q_\bullet$ we get two associated spectral sequences, ${}^{\vee}E$ and ${}^{<}E$. Using the Universal Coefficient Theorem for homology, it's not hard to check that

$${}^{\vee}E^2_{pq} = \begin{cases} H_p(P_\bullet \otimes_R N) & q = 0 \\ 0 & \text{else} \end{cases} \quad \text{and} \quad {}^{<}E^2_{pq} = \begin{cases} H_p(M_\bullet \otimes_R Q) & q = 0 \\ 0 & \text{else} \end{cases}$$

So both spectral sequences collapse at the E^2 page, and we see that

$$\mathrm{Tor}_p^R(-,N)(M) = H_p(P_\bullet \otimes_R N) = H_p(P_\bullet \otimes_R Q_\bullet) = H_p(M_\bullet \otimes_R Q) = \mathrm{Tor}_p^R(M,-)(N)$$

3 Some Examples of Spectral Sequences

Many more examples can be found at the nLab's entry on spectral sequences, available at `https://ncatlab.org/nlab/show/spectral+sequence`. Lots of our examples will come from the world of sheaf theory.

3.1 The Grothendieck Spectral Sequence

Introduced in the *Tôhoku* paper, this is one of the most useful spectral sequences, which computes the derived functor of a composition GF in terms of the derived functors of F and of G. To be precise, let $\mathcal{A} \xrightarrow{F} \mathcal{B} \xrightarrow{G} \mathcal{C}$ be a diagram of abelian categories and additive, left-exact functors. Suppose that the following conditions are satisfied:

i) $\mathcal{A}$ and $\mathcal{B}$ have enough injectives.

ii) F takes F-acyclic objects to G-acyclic objects.

Then for each $A \in \mathcal{A}$ there is a spectral sequence E (the **Grothendieck spectral sequence**), whose E_2 page is $E_2^{pq} = (\mathbf{R}^p G \circ \mathbf{R}^q F)(A)$, converging to $\mathbf{R}^{p+q}(GF)(A)$. So if we know the right derived functors of F, G we can compute the right derived functors of GF.

Many special cases of the Grothendieck spectral sequence are important enough to get their own name:

i) Let X and Y be topological spaces and $f : X \to Y$ be a continuous map. Consider the diagram $\mathbf{Ab}(X) \xrightarrow{f_*} \mathbf{Ab}(Y) \xrightarrow{\Gamma_Y} \mathbf{Ab}$. Note that the composition $\Gamma_Y \circ f_*$ is just Γ_X. The associated Grothendieck spectral sequence is the **Leray spectral sequence** of a sheaf $\mathcal{F}$ on X, whose E_2 page is $H^p(Y, \mathbf{R}^q f_* \mathcal{F}) \Rightarrow H^{p+q}(X, \mathcal{F})$.

ii) Now let G be a group, N a normal subgroup and A a G-module. Consider the **invariants functor** $A \mapsto A^G$. Then we may compute A^G in two steps as $A^G = (A^N)^{G/N}$. The Grothendieck spectral sequence of the functors $A \mapsto A^N$ and $A \mapsto A^{G/N}$ is called the **Hochschild–Serre spectral sequence**. Recalling that group cohomology is exactly the derived functor of the invariants functor, we see that the Hochschild–Serre spectral sequence is the spectral sequence whose E_2 page is

$$E_2^{pq} = H^p(G/N, H^q(N, A)) \Rightarrow H^{p+q}(G, A)$$

So if we know the group cohomology of some 'pieces' of G, we can put this together to get the group cohomology of G.

iii) Let X be a scheme, with structure sheaf $\mathcal{O}$. Recall that given two $\mathcal{O}$-modules $\mathcal{F}$, $\mathcal{G}$, we can define a sheaf $\mathscr{H}om_{\mathcal{O}}(\mathcal{F}, \mathcal{G})$ with $\mathscr{H}om_{\mathcal{O}}(\mathcal{F}, \mathcal{G})(U) = \mathrm{Hom}_{\mathcal{O}|_U}(\mathcal{F}|_U, \mathcal{G}|_U)$. Moreover, we may define another sheaf $\mathscr{E}xt$ as the right derived functor of $\mathscr{H}om$. Note that $\mathrm{Hom}_{\mathcal{O}}(\mathcal{F}, \mathcal{G})$ is the set of global sections of $\mathscr{H}om_{\mathcal{O}}(\mathcal{F}, \mathcal{G})$. Putting this together we obtain a spectral sequence whose E_2 page is

$$H^p(X, \mathscr{E}xt^q_{\mathcal{O}}(\mathcal{F}, \mathcal{G})) \Rightarrow \mathrm{Ext}^{p+q}_{\mathcal{O}}(\mathcal{F}, \mathcal{G})$$

the **local-to-global Ext spectral sequence**.

3.2 The Leray-Serre Spectral Sequence

Let's say we have a Serre fibration $F \to E \to B$ of topological spaces (e.g. a fibre bundle). There's a long exact sequence linking the homotopy groups of the three spaces F, E, B. But is there an analogous object for cohomology? Assuming that B is simply-connected, Serre proved that there is a spectral sequence with E_2 term $E_2^{pq} = H^p(B, H^q(F)) \Rightarrow H^{p+q}(E)$. A similar statement holds if B is not simply-connected and we take into account the action of $\pi_1 B$ on the cohomology of the fibres.

The Leray-Serre spectral sequence has a generalisation to any Eilenberg-Steenrod cohomology theory, the **Atiyah-Hirzebruch spectral sequence**. If K is any generalised cohomology theory, the Atiyah-Hirzebruch sequence has E_2 page $H^p(B, K^q(F)) \Rightarrow K^{p+q}(E)$. If we take the fibration $\mathrm{pt} \to X \to X$ then we recover the K-cohomology of X in terms of the ordinary cohomology with coefficients in the graded ring $K^*(\mathrm{pt})$.

3.3 The Künneth Spectral Sequence

Lots of standard theorems from algebraic topology have 'spectral counterparts'. Recall that if X and Y are any two topological spaces, and k is a field, then there is an isomorphism

$$\bigoplus_{i+j=n} H_i(X, k) \otimes_k H_j(Y, k) \cong H_n(X \times Y, k)$$

More generally, the Künneth Theorem tells us that if we use homology with coefficients in a principal ideal domain R, we recover the following short exact sequence:

$$0 \to \bigoplus_{i+j=n} H_i(X,R) \otimes_R H_j(Y,R) \to H_n(X \times Y, R) \quad \to \bigoplus_{i+j=n-1} \mathrm{Tor}_1^R(H_i(X,R), H_j(Y,R)) \to 0$$

Even more generally, if R is any commutative ring, there is a **Künneth spectral sequence** with E^2 page

$$E^2_{pq} = \bigoplus_{i+j=q} \mathrm{Tor}_p^R(H_i(X,R), H_j(Y,R)) \;\Rightarrow\; H_{p+q}(X \times Y, R)$$

which, in the situations above, collapses to yield the given relation.

There are also 'spectral versions' of the Seifert-Van Kampen Theorem and the Hurewicz theorem, although these are more complicated.

3.4 The Cartan-Leray Spectral Sequence

Let X be a pointed connected topological space and G a group acting on X freely and properly, i.e.

i) For all $x \in X$, the stabiliser $\mathrm{Stab}(x) \subseteq G$ is trivial.

ii) Every $x \in X$ has a neighbourhood U with $gU \cap U = \emptyset$ for all $g \neq e$.

Then there is a spectral sequence, the **Cartan-Leray spectral sequence**, with E^2 page $E^2_{pq} = H_p(G, H_q(X)) \Rightarrow H_{p+q}(X/G)$. For example, if X satisfies some mild point-set conditions (connected and locally-simply-connected) then $G = \pi_1(X)$ acting on the universal cover $\tilde{X}$ satisfies the hypotheses above. Since $\tilde{X}/G = X$, the spectral sequence becomes $H_p(\pi_1(X), H_q(\tilde{X})) \Rightarrow H_{p+q}(X)$.

3.5 The Mayer-Vietoris Spectral Sequence

Let X be a topological space, $\mathcal{F}$ a sheaf on X, and $\mathcal{U}$ a cover of X. Say that $\mathcal{U}$ is $\mathcal{F}$-acyclic if $\mathcal{F}$ is acyclic on any finite intersection of elements from $\mathcal{U}$. Then **Leray's Theorem** says that if $\mathcal{U}$ is an $\mathcal{F}$-acyclic cover, then the Čech cohomology $H^n(\mathcal{U}, \mathcal{F})$ is the same as the sheaf cohomology $H^n(X, \mathcal{F})$.

But what if $\mathcal{U}$ is not $\mathcal{F}$-acyclic? In general, there is a spectral sequence with E_2 page $E_2^{pq} = H^p(\mathcal{U}, \mathscr{H}^q(\mathcal{F})) \Rightarrow H^{p+q}(X, \mathcal{F})$, where $\mathscr{H}^q(\mathcal{F})$ is the presheaf on X that takes U to $H^q(U, \mathcal{F}|_U)$. If $\mathcal{U}$ was already $\mathcal{F}$-acyclic, then the spectral sequence collapses at the E_2 page to yield Leray's Theorem. If we take $\mathcal{U}$ to consist of just two open sets, then we recover the Mayer-Vietoris sequence for sheaf cohomology.

3.6 The Hodge-de Rham Spectral Sequence

Also called the **Frölicher spectral sequence**. Let X be a complex manifold, and let Ω^q be the sheaf of **holomorphic q-forms** on X (the complex q-forms whose coefficients are holomorphic functions). If X is compact Kähler (e.g. a smooth projective variety over $\mathbb{C}$), then we have the **Hodge decomposition**

$$\bigoplus_{p+q=l} H^p(X, \Omega^q) = H^l(X, \mathbb{C})$$

Can we recover a similar statement if X is not compact Kähler?

Pick (complex) local coordinates $z_1, \ldots, z_n$. Write $z_j = x_j + iy_j$ and define vector fields, the **Wirtinger derivatives**

$$\frac{\partial}{\partial z_j} := \frac{1}{2}(\frac{\partial}{\partial x_j} - i\frac{\partial}{\partial y_j})$$
$$\frac{\partial}{\partial \bar{z}_j} := \frac{1}{2}(\frac{\partial}{\partial x_j} + i\frac{\partial}{\partial y_j})$$

Note that a smooth function f on X is holomorphic precisely if all of the derivatives $\frac{\partial}{\partial \bar{z}_j}(f)$ vanish. We get associated complex differential forms

$$\mathrm{d}z_j := \mathrm{d}x_j + i\mathrm{d}y_j$$
$$\mathrm{d}\bar{z}_j := \mathrm{d}x_j - i\mathrm{d}y_j$$

Then any complex differential 1-form is a linear combination, over the ring of smooth functions on X, of such forms. For $p, q \geq 0$ let $\Omega^{p,q}$ (unfortunate but common notation!) be the sheaf of complex differential $(p+q)$-forms that are locally of the type

$$\alpha_1 \wedge \cdots \wedge \alpha_p \wedge \beta_1 \wedge \cdots \wedge \beta_q$$

where the α_k are linear combinations of the $\mathrm{d}z_j$ and the β_k are linear combinations of the $\mathrm{d}\bar{z}_j$. It's easy to see that the sheaf of complex differential n-forms is the sum $\oplus_{p+q=n}\Omega^{p,q}$.

There are boundary maps turning the collection $\Omega^{p,q}$ into a (cohomological) double complex. The cohomology in the q direction is the **Dolbeault cohomology** of X, written $H^{p,q}(X)$. It's not a homotopy invariant of X, since it depends on the complex structure. **Dolbeault's Theorem** is a complex version of de Rham's theorem that gives an isomorphism $H^{q,p}(X) \cong H^p(X, \Omega^q)$.

The total cohomology of $\Omega^{p,q}$ is the de Rham cohomology with complex coefficients. Hence we get a spectral sequence with E_1 page $H^p(X, \Omega^q) \Rightarrow H^{p+q}(X, \mathbb{C})$. If X is compact Kähler, this sequence collapses at the E_1 page and we recover the Hodge decomposition.

References

[1] Ravi Vakil, *Spectral Sequences: Friend or Foe?*
Available at `http://math.stanford.edu/~vakil/0708-216/216ss.pdf`

[2] Charles Weibel, *An Introduction to Homological Algebra*, CUP 1994.

[3] John McCleary, *A User's Guide to Spectral Sequences*, CUP 2001.