FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 32

RAVI VAKIL

CONTENTS

Last day: Hilbert polynomials and Hilbert functions. Higher direct image sheaves.

Today: Applications of higher pushforwards; crash course in spectral sequences; towards the Leray spectral sequence.

1. A USEFUL ALGEBRAIC FACT

I'd like to start with an algebra exercise that is very useful.

1.1. *Exercise (Important algebra exercise).* Suppose $M_1 \xrightarrow{\alpha} M_2 \xrightarrow{\beta} M_3$ is a complex of A-modules (i.e. $\beta \circ \alpha = 0$), and N is an A-module. (a) Describe a natural homomorphism of the cohomology of the complex, tensored with N, with the cohomology of the complex you get when you tensor with N, $H(M_*) \otimes_A B \to H(M_* \otimes_A N)$, i.e.

$$\left(\frac{\ker \beta}{\operatorname{im} \alpha}\right) \otimes_A N \to \frac{\ker(\beta \otimes N)}{\operatorname{im}(\alpha \otimes N)}.$$

I always forget which way this map is supposed to go.
(b) If N is *flat*, i.e. $\otimes N$ is an exact functor, show that the morphism defined above is an isomorphism. (Hint: This is actually a categorical question: if M_* is an exact sequence in an abelian category, and F is a right-exact functor, then (a) there is a natural morphism $FH(M_*) \to H(FM_*)$, and (b) if F is an exact functor, this morphism is an isomorphism.)

Example: localization is exact, so $S^{-1}A$ is a *flat* A-algebra for all multiplicative sets S. In particular, A_f is a flat A-algebra. We used (b) implicitly last day, when I said that given a quasicompact, separated morphism $\pi : X \to Y$, and an affine open subset $\operatorname{Spec} A$ of Y, and a distinguished affine open $\operatorname{Spec} A_f$ of that, the cohomology of any Cech complex computing the cohomology $\pi^{-1}(\operatorname{Spec} A)$, tensored with A_f, would be naturally isomorphic to the cohomology of the complex you get when you tensor with A_f.

Date: Thursday, February 16, 2006. Updated June 26.

Here is another example.

1.2. *Exercise (Higher pushforwards and base change).* (a) Suppose $f : Z \to Y$ is any morphism, and $\pi : X \to Y$ as usual is quasicompact and separated. Suppose $\mathcal{F}$ is a quasicoherent sheaf on X. Let

$$\begin{tikzcd} W \arrow[r, "f'"] \arrow[d, "\pi'"] & X \arrow[d, "\pi"] \\ Z \arrow[r, "f"] & Y \end{tikzcd}$$

is a fiber diagram. Describe a natural morphism $f^*(R^i\pi_*\mathcal{F}) \to R^i\pi'_*(f')^*\mathcal{F}$.
(b) If $f : Z \to Y$ is an affine morphism, and for a cover $\operatorname{Spec} A_i$ of Y, where $f^{-1}(\operatorname{Spec} A_i) = \operatorname{Spec} B_i$, B_i is a *flat* A-algebra, show that the natural morphism of (a) is an isomorphism. (You can likely generalize this immediately, but this will lead us into the concept of flat morphisms, and we'll hold off discussing this notion for a while.)

A useful special case if the following. If f is a closed immersion of a closed point in Y, the right side is the cohomology of the fiber, and the left side is the fiber of the cohomology. In other words, the fiber of the higher pushforward maps naturally to the cohomology of the fiber. We'll later see that in good situations this is an isomorphism, and thus the higher direct image sheaf indeed "patches together" the cohomology on fibers.

Here is one more consequence of our algebraic fact.

1.3. *Exercise (projection formula).* Suppose $\pi : X \to Y$ is quasicompact and separated, and $\mathcal{E}, \mathcal{F}$ are quasicoherent sheaves on X and Y respectively. (a) Describe a natural morphism

$$(R^i\pi_*\mathcal{E}) \otimes \mathcal{F} \to R^i\pi_*(\mathcal{E} \otimes \pi^*\mathcal{F}).$$

(b) If $\mathcal{F}$ is locally free, show that this natural morphism is an isomorphism.

Here is another consequence, that I stated in class 33. (It is still also in the class 33 notes.)

Exercise. Suppose that X is a quasicompact separated k-scheme, where k is a field. Suppose $\mathcal{F}$ is a quasicoherent sheaf on X. Let $X_{\overline{k}} = X \times_{\operatorname{Spec} k} \operatorname{Spec} \overline{k}$, and $f : X_{\overline{k}} \to X$ the projection. Describe a natural isomorphism $H^i(X, \mathcal{F}) \otimes_k \overline{k} \to H^i(X_{\overline{k}}, f^*\mathcal{F})$. Recall that a k-scheme X is *geometrically integral* if $X_{\overline{k}}$ is integral. Show that if X is geometrically integral and projective, then $H^0(X, \mathcal{O}_X) \cong k$. (This is a clue that $\mathbb{P}^1_{\mathbb{C}}$ is not a geometrically integral $\mathbb{R}$-scheme.)

2. Fun applications of the higher pushforward

Last day we proved that if $\pi : X \to Y$ is a projective morphism, and $\mathcal{F}$ is a coherent sheaf on X, then $\pi_*\mathcal{F}$ is coherent (under a technical assumption: if either Y and hence X are Noetherian; or more generally if $\mathcal{O}_Y$ is a coherent sheaf).

As a nice immediate consequence is the following. Finite morphisms are affine (from the definition) and projective (an earlier exercise); the converse also holds.

2.1. *Corollary. — If $\pi : X \to Y$ is projective and affine and $\mathcal{O}_Y$ is coherent, then π is finite.*

In fact, more generally, if π is universally closed and affine, then π is finite. We won't use this, so I won't explain why, but you can read about it in Atiyah-Macdonald, Exercise 5.35.

Proof. By the theorem from last day, $\pi_* \mathcal{O}_X$ is coherent and hence finitely generated. $\square$

Here is another handy theorem.

2.2. *Theorem (relative dimensional vanishing). — If $f : X \to Y$ is a projective morphism and $\mathcal{O}_Y$ is coherent, then the higher pushforwards vanish in degree higher than the maximum dimension of the fibers.*

This is false without the projective hypothesis. Here is an example of why.

Exercise. Consider the open immersion $\pi : \mathbb{A}^n - 0 \to \mathbb{A}^n$. By direct calculation, show that $R^{n-1} f_* \mathcal{O}_{\mathbb{A}^n - 0} \neq 0$.

Proof. Let m be the maximum dimension of all the fibers.

The question is local on Y, so we'll show that the result holds near a point p of Y. We may assume that Y is affine, and hence that $X \hookrightarrow \mathbb{P}^n_Y$.

Let k be the residue field at p. Then $f^{-1}(p)$ is a projective k-scheme of dimension at most m. Thus we can find affine open sets $D(f_1), \dots, D(f_{m+1})$ that cover $f^{-1}(p)$. In other words, the intersection of $V(f_i)$ does not intersect $f^{-1}(p)$.

If $Y = \operatorname{Spec} A$ and $p = [\mathfrak{p}]$ (so $k = A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$), then arbitrarily lift each f_i from an element of $k[x_0, \dots, x_n]$ to an element f_i' of $A_{\mathfrak{p}}[x_0, \dots, x_n]$. Let F be the product of the denominators of the f_i'; note that $F \notin \mathfrak{p}$, i.e. $p = [\mathfrak{p}] \in D(F)$. Then $f_i' \in A_F[x_0, \dots, x_n]$. The intersection of their zero loci $\cap V(f_i') \subset \mathbb{P}^n_{A_F}$ is a closed subscheme of $\mathbb{P}^n_{A_F}$. Intersect it with X to get another closed subscheme of $\mathbb{P}^n_{A_F}$. Take its image under f; as projective morphisms are closed, we get a closed subset of $D(F) = \operatorname{Spec} A_F$. But this closed subset does not include p; hence we can find an affine neighborhood $\operatorname{Spec} B$ of p in Y missing the image. But if f_i'' are the restrictions of f_i' to $B[x_0, \dots, x_n]$, then $D(f_i'')$ cover $f^{-1}(\operatorname{Spec} B)$; in other words, over $f^{-1}(\operatorname{Spec} B)$ is covered by $m+1$ affine open sets, so by the affine-cover vanishing theorem, its cohomology vanishes in degree at least $m+1$. But the higher-direct image sheaf is computed using these cohomology groups, hence the higher direct image sheaf $R^i f_* \mathcal{F}$ vanishes on $\operatorname{Spec} B$ too. $\square$

2.3. *Important Exercise.* Use a similar argument to prove *semicontinuity of fiber dimension of projective morphisms*: suppose $\pi : X \to Y$ is a projective morphism where $\mathcal{O}_Y$ is coherent. Show that $\{y \in Y : \dim f^{-1}(y) > k\}$ is a Zariski-closed subset. In other words, the dimension of the fiber "jumps over Zariski-closed subsets". (You can interpret the case $k = -1$ as the fact that projective morphisms are closed.) This exercise is rather important for having a sense of how projective morphisms behave! Presumably the result is true more generally for proper morphisms.

Here is another handy theorem, that is proved by a similar argument. We know that finite morphisms are projective, and have finite fibers. Here is the converse.

2.4. *Theorem (projective + finite fibers = finite). — Suppose $\pi : X \to Y$ is such that $\mathcal{O}_Y$ is coherent. Then π is projective and finite fibers if and only if it is finite. Equivalently, π is projective and quasifinite if and only it is finite.*

(Recall that quasifinite = finite fibers + finite type. But projective includes finite type.)

It is true more generally that proper + quasifinite = finite. (We may see that later.)

Proof. We show it is finite near a point $y \in Y$. Fix an affine open neighborhood $\operatorname{Spec} A$ of y in Y. Pick a hypersurface H in $\mathbb{P}^n_A$ missing the preimage of y, so $H \cap X$ is closed. (You can take this as a hint for Exercise 2.3!) Let $H' = \pi_*(H \cap X)$, which is closed, and doesn't contain y. Let $U = \operatorname{Spec} R - H'$, which is an open set containing y. Then above U, π is projective and affine, so we are done by the previous Corollary 2.1. $\square$

Here is one last potentially useful fact. (To be honest, I'm not sure if we'll use it in this course.)

2.5. *Exercise.* Suppose $f : X \to Y$ is a projective morphism, with $\mathcal{O}(1)$ on X. Suppose Y is quasicompact and $\mathcal{O}_Y$ is coherent. Let $\mathcal{F}$ be coherent on X. Show that

(a) $f^*f_*\mathcal{F}(n) \to \mathcal{F}(n)$ is surjective for $n \gg 0$. (First show that there is a natural map for any n! Hint: by adjointness of f_* with f_*.) Translation: for $n \gg 0$, $\mathcal{F}(n)$ is relatively generated by global sections.
(b) For $i > 0$ and $n \gg 0$, $R^i f_* \mathcal{F}(n) = 0$.

3. Toward the Leray spectral sequence: Crash course in spectral sequences

My goal now is to tell you enough about spectral sequences that you'll have a good handle on how to use them in practice, and why you shouldn't be frightened when they come up in a seminar. There will be some key points that I will not prove; it would be good, once in your life, to see a proof of these facts, or even better, to prove it yourself. Then in good conscience you'll know how the machine works, and you can close the hood once and for all and just happily drive the powerful machine.

My philosophy will be to tell you just about a stripped down version of spectral sequences, which frankly is what is used most of the time. You can always gussy it up later on. But it will be enough to give a quick proof of the Leray spectral sequence.

A good reference as always is Weibel. I learned it from Lang's *Algebra*. I don't necessarily endorse that, but at least his exposition is just a few pages long.

Let's get down to business.

For me, a *double complex* (in an abelian category) will be a bunch of objects $A^{p,q}$ ($p, q \in \mathbb{Z}$), which are zero unless $p, q \geq 0$, and morphisms $d^{p,q} : A^{p,q} \to A^{p+1,q}$ and $\delta^{p,q} : A^{p,q} \to A^{p,q+1}$ (we will always write these as d and δ and ignore the subscripts) satisfying $d^2 = 0$ and $\delta^2 = 0$, and one more condition: either $d\delta = \delta d$ ("all the squares commute") or $d\delta + \delta d = 0$ (they all "anticommute"). Both come up, and you can switch from one to the other by replacing $\delta^{p,q}$ with $(-1)^p \delta^{p,q}$. So I'll hereafter presume that all the squares anticommute, but that you know how to turn the commuting case into this one.

Also, there are variations on this definition, where for example the vertical arrows go downwards, or some different subset of the $A^{p,q}$ are required to be zero, but I'll leave these straightforward variations to you.

From the double complex (with the anticommuting convention), we construct a corresponding (single) complex A^* with $A^k = \oplus_i A^{i,k-i}$, with $D = d + \delta$. Note that $D^2 = (d+\delta)^2 = d^2 + (d\delta + \delta d) + \delta^2 = 0$, so A^* is indeed a complex. (Be sure you see how to interpret this in $A^{*,*}$!)

The *cohomology* of the single complex is sometimes called the *hypercohomology* of the double complex.

Our motivating goal will be to find the hypercohomology of the double complex. (You'll see later that we'll have other real goals, and that this is a red herring.)

Then here is recipe for computing (information) about the cohomology. We create a countable sequence of tables as follows. Table 0, denoted $E_0^{p,q}$, is defined as follows: $E_0^{p,q} = A^{p,q}$.

We then look just at the vertical arrows (the δ-arrows).

$$\begin{tikzcd}
\bullet & \bullet & \bullet \\
\bullet \arrow[u] & \bullet \arrow[u] & \bullet \arrow[u] \\
\bullet \arrow[u] & \bullet \arrow[u] & \bullet \arrow[u]
\end{tikzcd}$$

The columns are complexes, so we take cohomology of these vertical complexes, resulting in a new table, $E_1^{p,q}$. Then there are natural morphisms from each entry of the new table to the entry on the right. (This needs to be checked!)

$$\begin{tikzcd} \bullet \arrow[r] & \bullet \arrow[r] & \bullet \\ \bullet \arrow[r] & \bullet \arrow[r] & \bullet \\ \bullet \arrow[r] & \bullet \arrow[r] & \bullet \end{tikzcd}$$

The composition of two of these morphisms is again zero, so again we have complexes. We take cohomology of these as well, resulting in a new table, $E_2^{p,q}$. It turns out that there are natural morphisms from each entry to the entry two to the right and one below, and that the composition of these two is 0.

$$\begin{tikzcd} \bullet \arrow[drr] & \bullet & \bullet \\ \bullet \arrow[drr] & \bullet & \bullet \\ \bullet & \bullet & \bullet \end{tikzcd}$$

This can go on until the cows come home. The order of the morphisms is shown pictorially below.

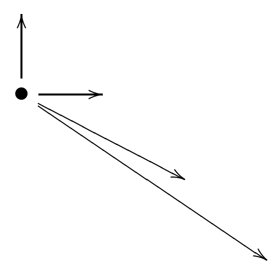

(Notice that the map always is "degree 1" in the grading of the single complex.)

Now if you follow any entry in our original table, eventually the arrow into it will come from outside of the first quadrant, and the arrow out of it will go to outside the first quadrant, so after a certain stage the complex will look like $0 \to E_?^{p,q} \to 0$. Then after that stage, the (p, q)-entry will never change. We define $E_\infty^{p,q}$ to be the table whose (p, q)th entry is this object. We say that $E_k^{p,q}$ *converges to* $E_\infty^{p,q}$.

Then it is a fact (or even a theorem) that there is a filtration of $H^k(A^*)$ by More precisely you can filter $H^k(A^*)$ with $k + 1$ objects whose successive quotients are $E_\infty^{i,k-i}$, where the sub-object is $E_\infty^{k,0}$, and the quotient $H^k(A^*)$ by the next biggest object is $E_\infty^{0,k}$. I hope that is clear; please let me know if I can say it better! The following may help:

$$\begin{array}{ccccccccc} & E_\infty^{0,k} & & E_\infty^{1,k-1} & & E_\infty^{k-1,1} & & E_\infty^{k,0} & \\ H^k(A^*) & \supset & ? & \supset & ? \supset \cdots \supset ? & \supset & ? & \supset & 0 \end{array}$$

(I always forget which way the quotients are supposed to go. One way of remembering it is by having some idea of how the result is proved. The picture here is that the double complex is filtered by subcomplexes $\oplus_{p\geq k, q\geq 0} A^{p,q}$, and the first term corresponding by taking the cohomology of the subquotients of this filtration. Then the "biggest quotient" corresponds to the left column, which remains true at the level of cohomology. If this doesn't help you, just ignore this parenthetical comment. If you have a better way of remembering this, even a mnemonic trick, please let me know!)

The sequence $E_k^{p,q}$ is called a *spectral sequence*, and we say that it *abuts* to $H^*(A^*)$. We often say that $E_2^{p,q}$ (or any other term) abuts to H.

Unfortunately, you only get partial information about $H^*(A^*)$. But there are some cases where you get more information: if all $E_\infty^{i,k-i}$ are zero, or if all but one of them are zero; or if we are in the category of vector spaces over a field k, and are interested only in the dimension of $H^*(A^*)$.

Also, in good circumstances, E_2 (or some other low term) already equals E_∞.

3.1. *Exercise.* Show that $H^0(A^*) = E_\infty^{0,0} = E_2^{0,0}$ and

$$0 \to E_2^{1,0} \to H^1(A^*) \to E_2^{0,1} \to E_2^{2,0} \to H^2(A^*).$$

3.2. *Exercise.* Suppose we are working in the category of vector spaces over a field k, and $\oplus_{p,q} E_2^{p,q}$ is a finite-dimensional vector space. Show that $\chi(H^*(A^*))$ is well-defined, and equals $\sum_{p,q}(-1)^{p+q} E_2^{p,q}$. (It will sometimes happen that $\oplus E_0^{p,q}$ will be an infinite-dimensional vector space, but that $E_2^{p,q}$ will be finite-dimensional!)

Eric pointed out that I was being a moron, and I could just as well have done everything in the opposite direction, i.e. reversing the roles of horizontal and vertical morphisms. Then the sequences of arrows giving the spectral sequence would look like this:

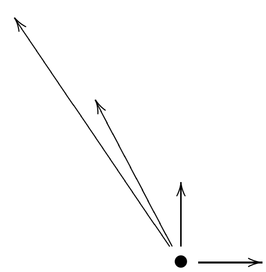

Then we would again get pieces of a filtration of $H^*(A^*)$ (where we have to be a bit careful with the order with which $E_\infty^{p,q}$ corresponds to the subquotients — it in the opposite order to the previous case).

I tried unsuccessfully to convince that Eric that I am *not* a moron, and that this was my secret plan all along. Both algorithms compute the same thing, and usually we don't care about the final answer — we often care about the answer we get in one way, and we get at it by doing the spectral sequence in the *other* way.

Now we're ready to try this out, and see how to use it in practice.

The moral of these examples is what follows: in the past, you've had prove various facts involving various sorts of diagrams, which involved chasing elements all around. Now, you'll just plug them into a spectral sequence, and let the spectral sequence machinery do your chasing for you.

Example: Proving the snake lemma. Consider the diagram

$$\begin{tikzcd} 0 \arrow[r] & D \arrow[r] & E \arrow[r] & F \arrow[r] & 0 \\ 0 \arrow[r] & A \arrow[r] \arrow[u, "\alpha"] & B \arrow[r] \arrow[u, "\beta"] & C \arrow[r] \arrow[u, "\gamma"] & 0 \end{tikzcd}$$

where the rows are exact and the squares commute. (Normally the snake lemma is described with the vertical arrows pointing downwards, but I want to fit this into my spectral sequence conventions.) We wish to show that there is an exact sequence

$$0 \to \ker \alpha \to \ker \beta \to \ker \gamma \to \operatorname{im} \alpha \to \operatorname{im} \beta \to \operatorname{im} \gamma \to 0. \tag{1}$$

We plug this into our spectral sequence machinery. We first compute the hypercohomology by taking the rightward morphisms first, i.e. using the order

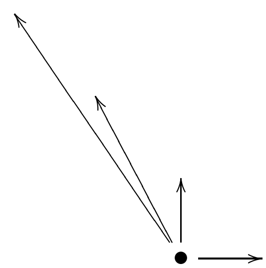

Then because the rows are exact, $E_1^{p,q} = 0$, so the spectral sequence has already converged: $E_\infty^{p,q} = 0$.

We next compute this "0" in another way, by computing the spectral sequence starting in the other direction.

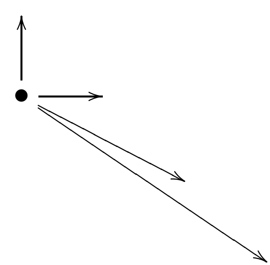

Then $E_1^{*,*}$ (with its arrows) is:

$$\begin{tikzcd} 0 \arrow[r] & \operatorname{im} \alpha \arrow[r] & \operatorname{im} \beta \arrow[r] & \operatorname{im} \gamma \arrow[r] & 0 \\ 0 \arrow[r] & \ker \alpha \arrow[r] & \ker \beta \arrow[r] & \ker \gamma \arrow[r] & 0. \end{tikzcd}$$

Then we compute $E_2^{*,*}$ and find:

$$\begin{tikzcd}
0 \arrow[drr] & ?? \arrow[drr] & ? \arrow[drr] & ? & 0 \\
0 & ? & ? & ?? & 0.
\end{tikzcd}$$

Then we see that after E_2, all the terms will stabilize except for the double question marks; and after E_3, even these two will stabilize. But in the end our complex must be the 0 complex. This means that in E_2, all the entries must be zero, except for the two double question marks; and these two must be the same. This means that $0 \to \ker\alpha \to \ker\beta \to \ker\gamma$ and $\operatorname{im}\alpha \to \operatorname{im}\beta \to \operatorname{im}\gamma \to 0$ are both exact (that comes from the vanishing of the single-question-marks), and

$$\operatorname{coker}(\ker\beta \to \ker\gamma) \cong \ker(\operatorname{im}\alpha \to \operatorname{im}\beta)$$

is an isomorphism (that comes from the equality of the double-question-marks). Taken together, we have proved the snake lemma (1)!

Example: the Five Lemma. Suppose

(2)
$$\begin{tikzcd}
F \arrow[r] & G \arrow[r] & H \arrow[r] & I \arrow[r] & J \\
A \arrow[r] \arrow[u, "\alpha"] & B \arrow[r] \arrow[u, "\beta"] & C \arrow[r] \arrow[u, "\gamma"] & D \arrow[r] \arrow[u, "\delta"] & E \arrow[u, "\epsilon"]
\end{tikzcd}$$

where the rows are exact and the squares commute.

Suppose α, β, δ, ϵ are isomorphisms. We'll show that γ is an isomorphism.

We first compute the cohomology of the total complex by starting with the rightward arrows:

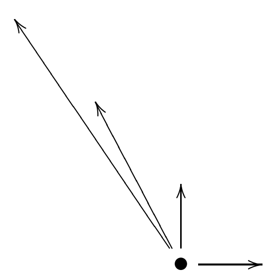

(I chose this because I see that we will get lots of zeros.) Then E_1 looks like this:

$$\begin{tikzcd}
? & 0 & 0 & 0 & ? \\
? \arrow[u] & 0 \arrow[u] & 0 \arrow[u] & 0 \arrow[u] & ? \arrow[u]
\end{tikzcd}$$

Then E_2 looks similar, and the sequence will converge by E_2 (as we'll never get any arrows between two non-zero entries in a table thereafter). We can't conclude that the cohomology of the total complex vanishes, but we can note that it vanishes in all but four degrees — and most important, in the two degrees corresponding to the entries C and H (the source and target of γ).

We next compute this in the other direction:

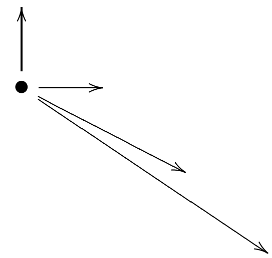

Then E_1 looks like this:

$$0 \longrightarrow 0 \longrightarrow ? \longrightarrow 0 \longrightarrow 0$$

$$0 \longrightarrow 0 \longrightarrow ? \longrightarrow 0 \longrightarrow 0$$

and the spectral sequence converges at this step. We wish to show that those two ?'s are zero. But they are precisely the cohomology groups of the total complex that we just showed *were* zero — so we're done!

Exercise. By looking at this proof, prove a subtler version of the five lemma, where one of the isomorphisms can instead just be required to be an injection, and another can instead just be required to be a surjection. (I'm deliberately not telling you which ones, so you can see how the spectral sequence is telling you how to improve the result.) I've heard this called the "subtle five lemma", but I like calling it the $4\frac{1}{2}$-lemma.

Exercise. If β and δ (in (2)) are injective, and α is surjective, show that γ is injective. State the dual statement. (The proof of the dual statement will be essentially the same.)

Exercise. Use spectral sequences to show that a short exact sequence of complexes gives a long exact sequence in cohomology.

3.3. *Exercise.* Suppose $\mu : A^* \to B^*$ is a morphism of complexes. Suppose C^* is the single complex associated to the double complex $A^* \to B^*$. (C^* is called the *mapping cone* of μ.) Show that there is a long exact sequence of complexes:

$$\cdots \to H^{i-1}(C^*) \to H^i(A^*) \to H^i(B^*) \to H^i(C^*) \to H^{i+1}(A^*) \to \cdots .$$

(There is a slight notational ambiguity here; depending on how you index your double complex, your long exact sequence might look slightly different.) In particular, people often use the fact μ induces an isomorphism on cohomology if and only if the mapping cone is exact.

(Does anyone else have some classical important fact that would be useful practice for people learning spectral sequences?)

Next day, I'll state and prove the Leray spectral sequence in algebraic geometry.

E-mail address: `vakil@math.stanford.edu`