Introduction to Spectral Sequences

Kimball Strong

1 Introduction

This document serves as an introduction to the theory of spectral sequences, suited to a student who has completed a course in algebraic topology covering both homology and homotopy. In this document, we motivate and outline the general machinery of spectral sequences, using a slew of nontraditional examples to capture some of the basic phenomena.

There are two companion notes, and together these three documents comprise the notes from the author's A Exam: one covers the construction of the bar spectral sequence, which relates the homology of an H-space X to the homology of its delooping BX, and uses it to prove S^7 cannot be endowed with an associative (that is, A_∞) multiplication. The final note is an example of a different flavor, the EHP spectral sequence. This spectral sequence contains information about the (stable and unstable) homotopy groups of spheres.

Unless otherwise stated, by $H(-)$ we shall mean ordinary singular (or simplicial) homology with coefficients in $\mathbb{Z}$; or, when applied to a chain complex, the homology of that chain complex. All spaces will be topological, and pointed. By $*$ we mean the space with a single point. We denote weak homotopy equivalences with "$\simeq$". We will be guilty of using "B" to mean at least 3 distinct functors, and "$|-|$" to mean at least 2 distinct functors.

Any reader should be warned that we take slightly different conventions than is typical. In particular we use n to denote homological degree and i to denote filtration degree. By contrast, Hatcher takes n to be the total degree (the sum of the homological and filtration degrees) and i to be the filtration degree. This results in a relative "sheering" of our spectral sequences relative to Hatcher (and several other sources). We have chosen to break this convention because it simplifies the exposition of the basic theory, even though the established convention is better for most examples.

2 The General Machinery

Homology is "computable," which is to say that given a suitable representation of a space X, there is an algorithm which will output the homology groups of X. For example, if X is given as a finite simplicial complex, then its homology groups can be calculated using simplicial homology: one takes the chain complex

$$\cdots \xrightarrow{\partial_{n+1}} \mathbb{Z}^{|X_n|} \xrightarrow{\partial_n} \mathbb{Z}^{|X_{n-1}|} \xrightarrow{\partial_{n-1}} \cdots$$

The maps ∂_n are maps between free $\mathbb{Z}$-modules, and therefore have representations as integer matrices, which we denote $M(\partial_n)$. In fact they have very simple representations: order the elements of X_n, as $(x_1, \ldots, x_{|X_n|})$, and likewise for X_{n-1}, whose elements we denote by $(y_1, \ldots, y_{|X_{n-1}|})$. Letting d_k denote the kth boundary operator, we have that

$$M(\partial_n)_{i,j} = \begin{cases} 1 & \text{if there exist an even } k \geq 0 \text{ s.t. } d_k(x_i) = y_j \\ -1 & \text{if there exist an odd } k \geq 0 \text{ s.t. } d_k(x_i) = y_j \\ 0 & \text{otherwise} \end{cases}$$

Consequently, the calculation of homology is reduced to calculating the quotient $\ker(\partial_n)/\mathrm{im}(\partial_{n+1})$ of kernels and images of integer matrices. This can be done with techniques of integer linear algebra; calculating

homology is just a matter of doing elementary matrix operations.

In practice, of course, we are generally not given a space as an explicit simplicial complex, except while doing problem sets for a first course in Algebraic Topology. Still, this example gives us hope that much in the way of homological computations can be done by simply "keeping track of all the details" and then carrying out some algebraic computations. But in practice, organizing the massive amounts of information present even in simple computations can be prohibitively time consuming. A *spectral sequence* is a general technique for simplifying this.

Suppose we are give a space X with a filtration $\{\} \subset X_0 \subset X_1 \subset \cdots \subset X$. By this we mean that we have spaces X_i for $i \geq 0$, equipped with inclusion maps $X_{i-1} \to X_i$, such that X is the colimit of the X_i. It is convenient to take the perspective that X is built out of the pieces $X_i \setminus X_{i-1}$; for instance, in the case where the filtration is the simplicial filtration of a simplicial complex, this is saying that X is iteratively formed from its i-simplices. By understanding the individual pieces, we hope to understand the space X. Of course, we also need to understand the way in which they attach to each other - hence we want to understand the relative pairs (X_i, X_{i-1}). For instance, suppose that X is a CW complex and X_i a CW filtration. Then we know that $X_i \setminus X_{i-1}$ is a disjoint union of i-disks, and $H_i(X_i, X_{i-1})$ is a free abelian group generated by these disks. These assemble into a chain complex

$$\cdots \to H_i(X_i, X_{i-1}) \to H_{i-1}(X_{i-1}, X_{i-2}) \to \cdots$$

with differentials determined by the attaching cellular maps, and the homology of this complex is the cellular homology of X, which is naturally isomorphic to the ordinary singular homology of X.

This example was rather simple because the filtration played nicely with homological dimension: the pair (X_i, X_{i-1}) has homology concentrated in degree i. Now let us consider the more general situation: given some arbitrary filtration $\bigcup X_i = X$, for each pair (X_i, X_{i-1}) we have homology groups $H_n(X_i, X_{i-1})$. We hope to maybe use this information to deduce the homology groups of X, as we did with the cellular filtration. So we write them all down in a big grid:

(X_3, X_2)	$H_0(X_3, X_2)$	$H_1(X_3, X_2)$	$H_2(X_3, X_2)$	$H_3(X_3, X_2)$
(X_2, X_1)	$H_0(X_2, X_1)$	$H_1(X_2, X_1)$	$H_2(X_2, X_1)$	$H_3(X_2, X_1)$
(X_1, X_0)	$H_0(X_1, X_0)$	$H_1(X_1, X_0)$	$H_2(X_1, X_0)$	$H_3(X_1, X_0)$
(X_0, X_{-1})	$H_0(X_0)$	$H_1(X_0)$	$H_2(X_0)$	$H_3(X_0)$
	H_0	H_1	H_2	H_3

If we were working with the special case of a CW filtration, only the terms on the main diagonal would be nonzero. In the CW filtration case, there were also some differential maps, obtained by taking a cell to its boundary. What do we have in this more general, less obviously geometric situation? Well, for each pair (X_i, X_{i-1}) we have long exact sequences of homology groups, and we can glue them together like so:

$$\begin{array}{ccccccccccc}
\cdots & \xrightarrow{\iota} & H_{n+1}(X_{i+2}) & \xrightarrow{r} & H_{n+1}(X_{i+2}, X_{i+1}) & \xrightarrow{\delta} & H_n(X_{i+1}) & \xrightarrow{\iota} & H_n(X_{i+2}) & \xrightarrow{r} & \cdots \\
 & & & & & & \| & & & & \\
\cdots & \xrightarrow{r} & H_{n+1}(X_{i+1}, X_i) & \xrightarrow{\delta} & H_n(X_i) & \xrightarrow{\iota} & H_n(X_{i+1}) & \xrightarrow{r} & H_n(X_{i+1}, X_i) & \xrightarrow{\delta} & \cdots \\
 & & & & \| & & & & & & \\
\cdots & \xrightarrow{\delta} & H_n(X_{i-1}) & \xrightarrow{\iota} & H_n(X_i) & \xrightarrow{r} & H_n(X_i, X_{i-1}) & \xrightarrow{\delta} & H_{n-1}(X_{i-1}) & \xrightarrow{\iota} & \cdots
\end{array}$$

Where by the vertical double lines we just mean the identity map. Exactness of the horizontal sequences tells us that the "staircase" sequences

$$\cdots \xrightarrow{r\circ\delta} H_{n+1}(X_{i+1}, X_i) \xrightarrow{r\circ\delta} H_n(X_i, X_{i-1}) \xrightarrow{r\circ\delta} \cdots$$

Are chain complexes; in the special case of the CW filtration and the complex of terms were $n = i$, this differential $r \circ \delta$ actually coincides with the cellular differential. But what does this differential tell us about our space in general? Let's explore a few examples, and then we'll regroup to prove some theorems.

2.1 First Examples of Spectral Sequences

All of the following will be toy examples of spectral sequences. We will observe informal things about them, as we have not yet properly defined them - for now, the reader may think of a spectral sequence as being the

collection of relative homology groups associated to each adjacent pair in a filtration $\{\} \subset X_0 \subset X_1 \subset \cdots \subset X$. The reader is invited to peruse this section, and then later read through it again after reading the subsequent, more formal section on the general theory.

Example Let X be the space $D^2 \vee S^1$. Consider the filtration $X_{-1} \subset X_0 \subset X_2 \subset X_3 = X$ where $X_{-1} = \{\}$, $X_0 = * \vee *$, $X_1 = S^1 \vee *$, $X^2 = S^1 \vee S^1$, $X_3 = D^2 \vee S^1$. Another way of phrasing this is that X has a cellular structure consisting of 4 cells: the point $*$, the left circle e^1_ℓ, the right circle e^1_r, and the disk attaching to the left circle, e^2_ℓ. This filtration is

$$\{\} \subset \{*\} \subset \{*, e^1_\ell\} \subset \{*, e^1_\ell, e^1_r\} \subset \{*, e^1_\ell, e^1_r, e^2_\ell\} = X$$

Then we can organize the relative homology of adjacent pairs in the filtration in the following spreadsheet:

(X_3, X_2)	0	0	$\mathbb{Z}[e^2]$	0
(X_2, X_1)	0	$\mathbb{Z}[e^1_r]$	0	0
(X_1, X_0)	0	$\mathbb{Z}[e^1_\ell]$	0	0
(X_0, X_{-1})	$\mathbb{Z}[*]$	0	0	0
	H_0	H_1	H_2	H_3

The terms are all rank 1 free groups, generated by the cells of X. We consider our differential $d : H_n(X_i, X_{i-1}) \to H_{n-1}(X_{i-1}, X_{i-2})$ with the formula $d = r \circ \delta$. In this case, $d = 0$: to check this we really only need to consider two cases for degree reasons, as $r \circ \delta$ will always go to the left and down by one unit each. We consider first $d : \mathbb{Z}[e^1_\ell] \to Z[*]$, which is 0 as e^1_ℓ is a cycle. The second is $d : \mathbb{Z}[e^2_\ell] \to \mathbb{Z}[e^1_r]$, which is 0 as it sends the generator e^2_ℓ to its boundary e^1_ℓ which is 0 in the group $H_1(X_2, X_1)$.

Of course, we know there *ought* to be a nontrivial differential here, an isomorphism $\mathbb{Z}[e^2_\ell] \to \mathbb{Z}[e^1_\ell]$ expressing that the boundary of e^2_ℓ is e^1_ℓ. We can actually get this from our long exact sequences: define $d^2 = r \circ \iota^{-1} \circ \delta$. Then this is well defined (again, there are only really two cases to check, which we leave to the reader), and moreover, the component $d^2 : \mathbb{Z}[e^2_\ell] \to \mathbb{Z}[e^1_\ell]$ is an isomorphism. If we take the homology with respect to this d^2, then we obtain

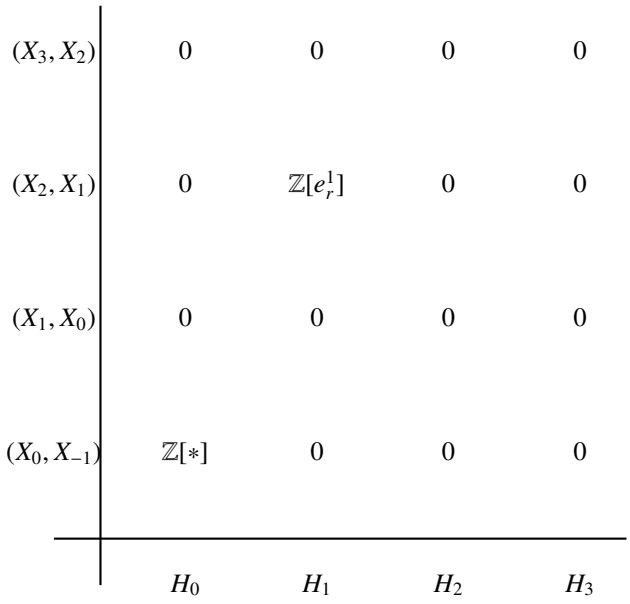

And we note that we can correctly read off the homology of our space X from this grid: it is $\mathbb{Z}[*]$ in dimension 0 and $\mathbb{Z}[e_r^1]$ in dimension 1, and 0 in other dimensions.

Example Let our space X be the torus $S^1 \times S^1$, with the filtration $* \subset S^1 \subset S^1 \vee S^1 \subset S^1 \times S^1$. If we write down our grid, we get

	H_0	H_1	H_2	H_3
(X_3, X_2)	0	0	$\mathbb{Z}$	0
(X_2, X_1)	0	$\mathbb{Z}$	0	0
(X_1, X_0)	0	$\mathbb{Z}$	0	0
(X_0, X_{-1})	$\mathbb{Z}$	0	0	0

And we can check that d is zero (this follows for the same reasons that the cellular differential for the cellular chain complex of the torus is 0). Furthermore, if we try to define d^2 as in the previous example, then this is also 0. This fits with the following observation: this grid contains precisely the homology of X! That is, if we sum up along the columns, we have a $\mathbb{Z}$ in the H_0-column, a $\mathbb{Z} \oplus \mathbb{Z}$ in the H_1-column, and $\mathbb{Z}$ in the H_2-column. So we appear to have again computed the homology from a non-cellular filtration. This case demonstrates a phenomenon we have not yet seen: there the H_1 column has two separate terms which we sum together to get the overall homology. The next example will further explore this.

Example Let X be the mapping cylinder of a degree two map $S^1 \to S^1$. Filter it by $X_0 = *$, $X_1 =$ the source circle, and $X_2 = X$. Then the grid we get looks as follows:

(X_2, X_1)	0	$\mathbb{Z}/2\mathbb{Z}$	0
(X_1, X_0)	0	$\mathbb{Z}$	0
(X_0, X_{-1})	$\mathbb{Z}$	0	0
	H_0	H_1	H_2

The first row is the homology of a point, the second the homology of a circle relative to a point, and the third the homology of the mapping cylinder relative to the source space, which is equivalently the reduced homology of the mapping cone - in this case, the mapping cone is $\mathbb{R}P^2$, hence the $\mathbb{Z}/2\mathbb{Z}$. We can again check that all the differentials are 0. So what are we to make of this? Well, certainly we can't just read off the homology again - our space X is homotopy equivalent to a circle, so it should have H_0 and H_1 both $\mathbb{Z}$, but here we have some strange copy of $\mathbb{Z}/2\mathbb{Z}$ floating around.

Well, we attempt to explain it as follows: if we consider the chain of inclusions $X_0 \subset X_1 \subset X_2 = X$ and apply H_1, we have a sequence of groups $0 \to \mathbb{Z} \to \mathbb{Z}$, where the last map is multiplication by 2. If we take quotients of adjacent terms, we therefore get $\mathbb{Z}$ and $\mathbb{Z}/2\mathbb{Z}$. So maybe that $\mathbb{Z}/2\mathbb{Z}$ has something to do with the fact that the inclusion of $H_1(X_1, X_0)$ into $H_1(X)$ isn't an isomorphism, but has $\mathbb{Z}/2\mathbb{Z}$ as its cokernel.

Example Let X be a CW-complex and X_i be the $2i$-skeleton of X. Then there are two types of nonzero terms, $H_{2i}(X_i, X_{i-1})$ and $H_{2i-1}(X_i, X_{i-1})$. We have that $H_{2i}(X_i, X_{i-1})$ is generated by the $2i$-cells of X, whereas $H_{2i-1}(X_i, X_{i-1})$ is isomorphic to $H_{2i-1}(X)$. Our spreadsheet looks like

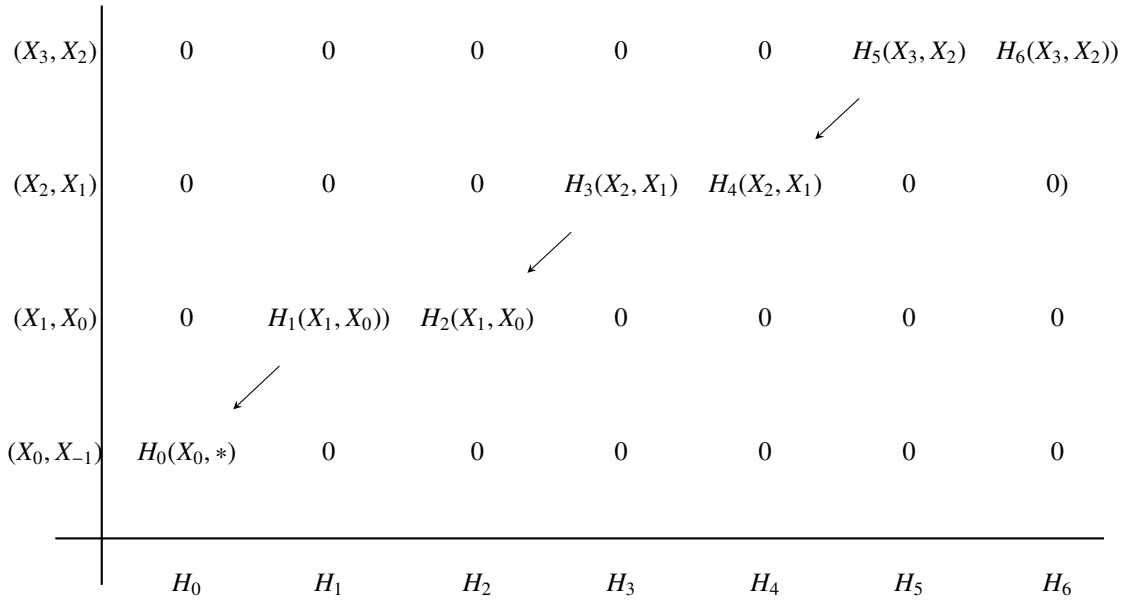

The differentials which may be nonzero are those which go $H_{2i-1}(X_{2i}, X_{2i-2}) \to H_{2i-1}(X_{2i-2}, X_{2i-4})$, and we have drawn them in. Some basic (though nontrivial) analysis of cellular homology tells us that if we take the homology of all differentials, we again end up with the homology groups of X, in the appropriate columns.

Example Let X be a CW-complex and $X_{2i} = X_{2i+1}$ be the i-skeleton of X. In this case, we end up with a spreadheet like:

(X_3, X_2)	0	0	0
(X_2, X_1)	0	$H_1(X_2, X_1)$	0
(X_1, X_0)	0	0	0
(X_0, X_{-1})	$H_0(X_0, *)$	0	0
	H_0	H_1	H_2

Here none of the differentials $r \circ \delta$ can be nonzero, simply for degree reasons. But again, we can define $d^2 = r \circ \iota^{-1} \circ \delta$, and we will get a stretched-out version of our first example, with the chain complex

$$\cdots \to H_2(X_4, X_3) \to H_1(X_2, X_1) \to H(0, X_0, \{\})$$

which is precisely the cellular chain complex of the space X, and so taking homology of this we will again get the homology of X.

These examples demonstrate various phenomena with spectral sequences that are absent from the special case of a cellular filtration: (1) that there may be nonzero groups $H_n(X_i, X_{i-1})$ for many values of n for a fixed i, and (2) that terms $H_n(X_i, X_{i-1})$ may interact with terms far away in the filtration, and (3) reconstructing the homology of X from the grid of groups may be nontrivial. We need some machinery that can organize all of this.

2.2 Some actual definitions

Definition A (bigraded) **exact couple** is a diagram

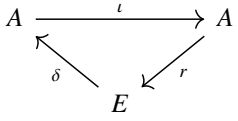

of (bigraded) modules which is exact at each term.

CONVENTION ALERT: we will omit the word "bigraded" from now on when describing exact couples, as all our exact couples will be of bigraded modules.

Example Let X be a space and (X_i) a filtration. Then there is an exact couple

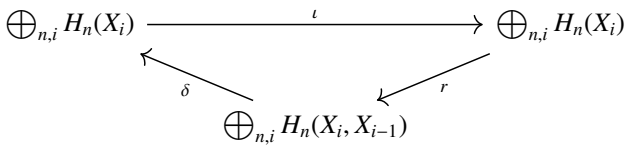

Note the degrees of each map: ι has bidegree $(0, 1)$, r has bidegree $(0, 0)$, and δ has bidegree $(-1, -1)$.

The point of this formalism is that it allows us to neatly bundle up what we did before:

Definition Given an exact couple

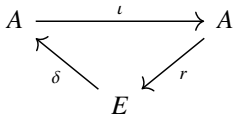

There is a differential $d : E \to E$ defined by $d = r \circ \delta$. The **derived couple** is the exact couple given by

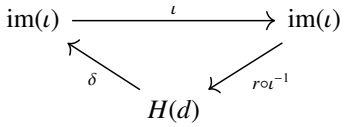

Where by $H(d)$ we mean the homology of d, by $\bar{\delta}$ we mean the map which sends a homology class represented by x to $\delta(x)$, and by $r \circ \iota^{-1}$ we mean the map which sends $\iota(x)$ to $r(x)$.

Theorem 2.1. *The above is a well defined exact couple.*

The proof of this statement is a not-too-difficult exercise in homological algebra, so we omit it.

In the second and fifth examples above, we used the differential $d = r \circ \iota^{-1} \circ \delta$, which is the differential associated to the derived exact couple of our starting exact couple. We can think of this as taking the derived couple, and then taking the homology of its differential to calculate the homology of our space X - or, equivalently, taking the derived couple of our derived couple. This iterated process is what leads to our main definition:

Definition Let

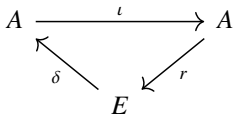

be an exact couple. Its associated **spectral sequence** is the sequence of differential bigraded modules (E^i, d^i) defined as follows: E^i is the bottom term of the $(i-1)$-fold derived couple, and d^i the associated differential. In particular, $E^1 = E$, $d^1 = r \circ \delta$, and $E^{i+1} = H(d^i)$. We refer to each E^i as the ith page of the spectral

sequence. If for some r, we have that $\forall i \geq r\ d^i = 0$, then it follows that $E^r \cong E^{r+1} \cong \cdots$ and we denote this as E^∞. In this case we say the spectral sequence **collapses** at the E^r-page. More weakly, if for any fixed i and n, $E^r_{n,i} \cong E^{r+1}_{n,i} \cong \cdots$ then we say that the spectral sequence *converges* and call the stable values $E^\infty_{i,n}$, the *E-infinity page*. Note that collapse implies convergence, but the reverse implication does not necessarily hold.

Definition Let $\bullet \subset X_0 \subset X_1 \subset \cdots$ be a filtration of a space X. The (homological) spectral sequence associated to this filtration is the spectral sequence associated to the exact couple

$$\begin{array}{ccc} \bigoplus_{n,i} H_n(X_i) & \xrightarrow{\quad\iota\quad} & \bigoplus_{n,i} H_n(X_i) \\ & {\scriptstyle\delta}\nwarrow \quad \swarrow{\scriptstyle r} & \\ & \bigoplus_{n,i} H_n(X_i, X_{i-1}) & \end{array}$$

In our previous examples, the spectral sequences collapsed at the E^2, E^2, and E^3 pages, and these pages contained the homology of X. More generally, we might wonder exactly what the E^∞ page of the spectral sequence contains. We now make a simplifying assumption about our exact couples that will hold in all cases of interest to us.

1. For each n, the sequence $\cdots \to A_{n,i} \to A_{n,i+1} \to \cdots$ stabilizes. Equivalently, for a fixed n, only finitely many of the maps $\iota_{n,i}$ are not isomorphisms. Also equivalently, for a fixed n, only finitely many of the terms $E_{n,i}$ are nonzero.[1]

2. For each n, the sequence $\cdots \to A_{n,i} \to A_{n,i+1} \to \cdots$ is zero for sufficiently low i.

If condition (1) is satisfied, we denote the stable term as $A_{n,\infty}$. This is satisfied for the homological exact couple of any filtration such that X_i is obtained from X_{i-1} by gluing cells of dimension increasing with i. This includes all of our examples so far. Since we take our filtrations to begin with $X_{-1} = \{\bullet\}$, condition (2) is satisfied for the homological exact couple of any filtration.

With these assumptions, we can prove the theorem which lends spectral sequences their computational utility:

Theorem 2.2. *Let*

$$\begin{array}{ccc} A & \xrightarrow{\quad\iota\quad} & A \\ & {\scriptstyle\delta}\nwarrow \quad \swarrow{\scriptstyle r} & \\ & E & \end{array}$$

Be an exact couple satisfying (1) and (2). Then the associated spectral sequence converges. Furthermore, if we denote by $F_{n,i}$ the image of $A_{n,i}$ in $A_{n,\infty}$, then

$$E^\infty_{n,i} \cong F_{n,i}/F_{n,i-1}$$

Thus, modulo the extension problem, the E^∞ page of the spectral sequence tells us the stable values $A_{n,\infty}$.

[1] A stronger condition that is often satisfied is that the sequence $\iota, \iota^2, \iota^3, \ldots$ stabilizes up to isomorphism.

Proof. For each $r \geq 0$, the rth derived couple gives us an exact sequence

$$E^r_{n+1,i+r-1} \to A^r_{n,i+r-2} \to A^r_{n,i+r-1} \to E^r_{n,i} \to A^r_{n-1,i-1}$$

Condition (1) tells us that for sufficiently large r, the term $E^r_{n+1,i+r-1}$ is 0. Condition (2) tells us that for sufficiently large r, the term $A^r_{n-1,i-1}$ is 0. Consequently, for sufficiently large r, we have a short exact sequence

$$0 \to A^r_{n,i+r-2} \to A^r_{n,i+r-1} \to E^r_{n,i} \to 0$$

Which, referring to the definition of the derived exact couple, is

$$0 \to \iota^{r-1}(A_{n,i-1}) \to \iota^{r-1}(A_{n,i}) \to E^r_{n,i} \to 0$$

which proves the claim. □

It is typical to write the conclusion of this theorem as $E^r_{n,i} \Rightarrow A_{n,\infty}$. As an immediate corollary of this, we get the following:

Theorem 2.3. *Let X be a space with a filtration $\bullet \subset X_0 \subset \cdots \subset X$, such that the homological exact couple satisfies (1). Then there is a spectral sequence with $E^1_{n,i} = H_n(X_i, X_{i-1})$ and $E^r_{n,i} \Rightarrow H_n(X)$. The differentials d^r have degree $(-1, -1-r)$*

This may seem a daunting theorem to apply: there is a potentially immense amount of data necessary to write down the E^1 page, and furthermore computing further pages (up to the E^∞ page) may require very explicit, tedious computations of iterated exact couples. For this reason, the existence of a spectral sequence is not enough to perform efficient calculations. However, in many cases, purely formal arguments can yield powerful results, and the general framework tends to greatly illuminate the study of a particular space.

References

[1] Allen Hatcher. *Algebraic Topology*