07-13-2015

Contents

As we've been saying, this course is a bit experimental. Some of it is being worked out on the fly! – and some things we're used to from classical algebraic topology just don't work.

At Inna's request, we will be saving the calculation of the homology groups of a torus for a later date, when it can be presented in a more polished form. Finite models for spaces have a tendency to get large very quickly; consequently, calculating homology by hand and from the definition quickly becomes unwieldy.

In this set of notes, we're going to go after some more low-hanging fruit. First, we'll consider a notion of "dimension" related to homology, and then we'll develop a useful piece of machinery called the **Mayer-Vietoris Sequence**.

1. Dimension

Having a reasonable notion of dimension can be quite useful. For example, the rank-nullity theorem from linear algebra allows us to draw conclusions about linear maps with hardly any work. We're going to introduce two measures of dimension for a poset (equivalently, an A-space).

Suppose $\mathscr{C}$ is a poset. Given $x, y \in \mathscr{C}$, we say that x and y are **comparable** if either $x \leqq y$ or $y \leqq x$. We say they are **incomparable** otherwise. Since the order $\leqq$ on $\mathscr{C}$ is partial, there might be incomparable elements, but a subset $A \subseteqq \mathscr{C}$ is called a **chain** if any pair of elements $a, b \in A$ *are* comparable.

Exercise: Consider the poset of subsets of $\{0, 1, 2\}$ ordered under inclusion:

Date: July 15, 2015.

$$\begin{tikzcd}
 & \{0,1,2\} & \\
\{0,1\} \arrow[ur] & \{0,2\} \arrow[u] & \{1,2\} \arrow[ul] \\
\{0\} \arrow[u] \arrow[ur] & \{1\} \arrow[ul] \arrow[ur] & \{2\} \arrow[ul] \arrow[u] \\
 & \varnothing \arrow[ul] \arrow[u] \arrow[ur] &
\end{tikzcd}$$

Draw some chains (this is where the name comes from).

The **height** of a poset $\mathscr{C}$ to be the supremum (least upper bound) of the cardinalities of chains in $\mathscr{C}$:

$$\mathrm{height}(\mathscr{C}) := \sup_{\substack{A \subseteqq \mathscr{C} \\ A \text{ a chain}}} |A|.$$

In the case of a finite poset (which is what we're most interested in), this is just the largest size of a chain in $\mathscr{C}$. This gives one notion of how big/tall a poset is.

Exercise: Suppose $\mathscr{C}$ and $\mathscr{D}$ are finite posets and that $f : \mathscr{C} \to \mathscr{D}$ is a surjective order-preserving map such that $f(x) \leqq f(y)$ implies $x \leqq y$. Prove that $\mathrm{height}(\mathscr{C}) \geqq \mathrm{height}(\mathscr{D})$. Compare to a surjective linear map between vector spaces.

Exercise: Suppose $\mathscr{C}$ and $\mathscr{D}$ are finite posets and that $r : \mathscr{C} \rightleftarrows \mathscr{D} : s$ are order preserving maps such that $r \circ s = \mathrm{id}$. Prove that r is surjective, s is injective, and $\mathrm{height}(\mathscr{C}) \geqq \mathrm{height}(\mathscr{D})$.

Height is related to the geometry and homology of the poset. If A is a finite chain, then it can be written

$$A = \{a_0 < a_1 < \cdots < a_n\}.$$

Thus, a chain of size $(n+1) < \infty$ is the same thing as a nondegenerate n-simplex. We can therefore regard the height of a poset as a measure of geometric dimension:

$$\dim_{geom}(\mathscr{C}) := \mathrm{height}(\mathscr{C}) - 1,$$

where the geometric dimension $\dim_{geom}(\mathscr{C})$ is the dimension of the largest nondegenerate simplex in $\mathscr{C}$. It follows that if $\mathrm{height}(\mathscr{C}) < n+1$, then $\mathscr{C}$ has no nongenerate n-simplices. Recalling that the homology of $\mathscr{C}$ could be computed using only the nondegenerate simplices in $\mathscr{C}$, it follows that if $\mathrm{height}(\mathscr{C}) < n+1$, then $H_n(\mathscr{C}) = 0$. We define the "topological dimension"[1] of $\mathscr{C}$ to be the largest n such that $H_n(\mathscr{C}) \neq 0$:

$$\dim_{top}(\mathscr{C}) := \max\left\{ n \in \mathbb{N} \,\middle|\, H_n(\mathscr{C}) \neq 0 \right\}.$$

[1]This is nonstandard terminology.

From the definitions, we see that

$$\dim_{top}(\mathscr{C}) \leqq \dim_{geom}(\mathscr{C}) = \text{height}(\mathscr{C}) - 1.$$

The moral is that homology detects dimension in some ways. However, because homology is a homotopy invariant, $\dim_{top}(\mathscr{C})$ can be strictly smaller than $\dim_{geom}(\mathscr{C})$.

Exercise: Give an example of a poset $\mathscr{C}$ such that $\dim_{top}(\mathscr{C}) < \dim_{geom}(\mathscr{C})$. (Hint: what happens if $\mathscr{C}$ has a maximum element?)

Exercise: Suppose that $\mathscr{C}$ and $\mathscr{D}$ are posets and that $r : \mathscr{C} \rightleftarrows \mathscr{D} : s$ are order preserving maps such that $r \circ s = \text{id}$. Prove that r is surjective, s is injective, and $\dim_{top}(\mathscr{C}) \geqq \dim_{top}(\mathscr{D})$. Compare to linear maps between vector spaces.

2. The Mayer-Vietoris Sequence

The Mayer-Vietoris sequence is a tool that allows us to relate the homology of a space X to the homology of its subspaces. Informally, suppose that $X = A \cup B$ and that A and B are "sufficiently nice". In this case, the Mayer-Vietoris sequence will give us information about how the homology of X, A, B, and $A \cap B$ interact. Thus, if we choose A and B wisely, we can use this sequence to leverage (known) information about $H_*(A)$, $H_*(B)$ and $H_*(A \cap B)$ up to information about $H_*(X)$. More precisely, the Mayer-Vietoris sequence is a long exact sequence involving $H_*(X)$, the **direct sum**[2] $H_*(A) \oplus H_*(B)$, and $H_*(A \cap B)$.

Example: To illustrate how this might be useful, consider the torus $T = S^1 \times S^1$, where S^1 is regarded as the set of unit length complex numbers. Let

$$A = T \setminus (\{1\} \times S^1)$$
$$B = T \setminus (\{-1\} \times S^1)$$

be the subsets obtained by deleting two opposite bounding circles from T (draw the picture). A and B are both homeomorphic to the tube $(-1, 1) \times S^1$, which is homotopy equivalent to S^1 via the contraction onto $\{0\} \times S^1$. On the other hand, $A \cap B$ is homeomorphic to two copies of this tube, and hence homotopy equivalent to $S^1 \sqcup S^1$. Since homology is homotopy invariant, the Mayer-Vietoris sequence would allow us to understand $H_*(T)$ in relation to $H_*(S^1) \oplus H_*(S^1)$ and $H_*(S^1 \sqcup S^1)$. Thus, given sufficient information about $H_*(S^1)$, we might hope to recover $H_*(T)$.

The Mayer-Vietoris sequence also is useful for understanding the homology of certain inductively defined spaces; later on, we're going to show how to use it to compute the homology of all S^n inductively from $H_*(S^1)$. Note, by S^n, we mean a *finite (poset) model* for the space

$$\{x \in \mathbb{R}^{n+1} \mid ||x|| = 1\}$$

since we have not really considered the homology of spaces in general. To see how this will go, let us first recall how suspension works.

[2]We will define $\oplus$ momentarily.

2.1. **Suspension and Spheres.** Let us begin with the geometric situation. For a space X, define the **cone on** X to be

$$X \times [0,1]/(X \times \{1\}),$$

i.e. the space obtained by taking the cylinder $X \times [0,1]$ on X and then collapsing the top to a point. The (unreduced) **suspension** ΣX is obtained by gluing two copies of CX together along their base to form a double-cone over X:

$$\Sigma X = CX \cup_X CX.$$

It is sometimes convenient to regard the vertex of one cone as a "north pole" and the vertex of the other as a "south pole".

Exercise: Let $S^0 := \{0,1\}$. Show inductively that the Euclidean sphere $S^n \subset \mathbb{R}^{n+1}$ is homeomorphic to the n-fold suspension $\Sigma^n S^0$.

Poset suspension is completely analogous; the suspension $\Sigma\mathscr{C}$ is the double-cone on $\mathscr{C}$. Formally, the **suspension** $\Sigma\mathscr{C}$ of a poset $\mathscr{C}$ is obtained by adjoining two new incomparable points $\{N_1, N_2\}$ and declaring that $N_i > x$ for $i = 1, 2$ and every $x \in \mathscr{C}$. Adjoining a single new point $N > x$ corresponds to forming a cone over $\mathscr{C}$, so adjoining N_1 and N_2 forms a double-cone. Alternatively, we could set $N_i < x$ for $i = 1, 2$ and all $x \in \mathscr{C}$, but we *cannot* have $N_1 < x < N_2$. Orientation matters.

Exercise: What goes wrong if we take $N_1 < x < N_2$? (Hint: maximum.)

Thus, for posets, it is **not** a good idea to think of the cone vertices as north and south poles, but rather as a doubled north (or south) pole.

In analogy to the case for Euclidean spheres, we start wth $S^0 := \{0,1\}$, regarded as a poset with two incomparable elements, and take $S^n := \Sigma^n S^0$. Observe S^n will have $2n+2$ points. In particular, S^1 is the familiar poset below.

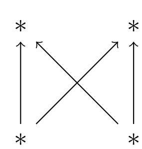

Exercise: Draw pictures of the posets S^n for some $n > 1$.

Now, the key observation for us is that $\Sigma\mathscr{C} = (\mathscr{C} \cup \{N_1\}) \cup (\mathscr{C} \cup \{N_2\})$ and $(\mathscr{C} \cup \{N_1\}) \cap (\mathscr{C} \cup \{N_2\}) = \mathscr{C}$. The cones $\mathscr{C} \cup \{N_\bullet\}$ are contractible because they have a maximum, hence $H_n(\mathscr{C} \cup \{N_\bullet\}) = 0$ for all $n > 0$ and $H_0(\mathscr{C} \cup \{N_\bullet\}) \cong \mathbb{Z}$. Thus, taking this decomposition as input into the Mayer-Vietoris Sequence will yield a direct comparison between $H_*(\mathscr{C})$ and $H_*(\Sigma\mathscr{C})$. This will let us compute $H_*(S^n)$ inductively.

2.2. **Direct Sums.** We shall now give a brief account of the direct sum. Suppose that G and H are abelian groups. The **direct sum of** G **and** H is the set of ordered pairs $(g,h) \in G \times H$ equipped with componentwise operations:

$$\begin{aligned} (g,h) + (g',h') &= (g +_G g', h +_H h') \\ 0 &= (0_G, 0_H) \\ -(g,h) &= (-g,-h). \end{aligned}$$

It can be thought of as analogous to the product of two spaces, or the Cartesian product of sets. For this reason, it is sometimes called the Cartesian product of abelian groups.

Interestingly, there is a second perspective on the direct sum, which better accounts for its name. We claim that $G \oplus H$ is the abelian group obtained by formally allowing ourselves to add elements of G and H together (with a commutative sum), while respecting the existing relations in G and H. For the time being, denote this new group $G \boxplus H$. The generic element of $G \boxplus H$ is a finite sum of the form

$$g_{1,1} + \cdots + g_{1,n_1} + h_{1,1} + \cdots + h_{1,m_1} + g_{2,1} + \cdots + g_{2,n_2} + h_{2,1} + \cdots + g_{N,n_N}/h_{N,m_N},$$

which can be regrouped simply into

$$\sum g_{i,j} + \sum h_{k,l} = g + h.$$

Note, the sum $\sum g_{i,j}$ takes place in G, while the sum $\sum h_{k,l}$ takes place in H. Since we assume the group operation $+ = +_{\boxplus}$ on $G \boxplus H$ is commutative, $+_{\boxplus}$ is actually "componentwise":

$$(g +_{\boxplus} h) +_{\boxplus} (g' +_{\boxplus} h') = g +_{\boxplus} g' +_{\boxplus} h +_{\boxplus} h' = (g +_G g') +_{\boxplus} (h +_H h').$$

Consequently, identities and inversion are also: $0_{G \boxplus H} = 0_G + 0_H$ and $-(g+h) = (-g) + (-h)$. Thus, we recover the same group structure as $G \oplus H$, provided we identify $g +_{\boxplus} h$ with the ordered pair (g,h).

Exercise: (Direct Sums are categorical products) Suppose G, H, K are abelian groups. Check that the projection maps

$$\begin{tikzcd} G & G \oplus H \arrow[l, "p"'] \arrow[r, "q"] & H \end{tikzcd}$$

defined by $p(g,h) = g$ and $q(g,h) = h$ are group homomorphisms. Show that for any pair of homomorphisms $f : K \to G$ and $g : K \to H$, there is a unique homomorphism $k : K \to G \oplus H$ such that the diagram below commutes.

$$\begin{tikzcd} G & G \oplus H \arrow[l, "p"'] \arrow[r, "q"] & H \\ & K \arrow[ul, "f"] \arrow[u, dashed, "k"] \arrow[ur, "g"'] & \end{tikzcd}$$

Exercise: (Direct Sums are categorical sums) Suppose G, H, K are abelian groups. Check that the inclusion maps

$$\begin{tikzcd} G \arrow[r, "i"] & G \oplus H & H \arrow[l, "j"'] \end{tikzcd}$$

defined by $i(g) = (g, 0)$ and $j(h) = (0, h)$ are group homomorphisms. Show that for any pair of homomorphisms $f : G \to K$ and $g : H \to K$, there is a unique homomorphism $k : G \oplus H \to K$ such that the diagram below commutes.

$$\begin{tikzcd} G \arrow[r, "i"] \arrow[dr, "f"'] & G \oplus H \arrow[d, dashed, "k"] & H \arrow[l, "j"'] \arrow[dl, "g"] \\ & K & \end{tikzcd}$$

Exercise: Make sense of what a 2×2 matrix of homomorphisms $G \oplus H \to G' \oplus H'$ should mean. (Hint: regard $G \oplus H$ as a sum and $G' \oplus H'$ as a product.)

2.3. **Constuction of the Mayer-Vietoris Sequence.** We will need one more definition before presenting the Mayer-Vietoris sequence. Let $\mathscr{C}$ be a poset. A subset $\mathscr{A} \subseteqq \mathscr{C}$ is called a **sieve in** $\mathscr{C}$ if, for any $x, y \in \mathscr{C}$, if $x \leqq y$ and $y \in \mathscr{A}$, then $x \in \mathscr{A}$. Succinctly, a sieve in $\mathscr{C}$ is a downward-closed subset of $\mathscr{C}$.

Exercise: Suppose $\mathscr{C}$ is a poset and $\mathscr{A} \subseteqq \mathscr{C}$. Prove that $\mathscr{A}$ is a sieve if and only if it is open in the associated Alexandroff topology.

Thus, sieves in a poset $\mathscr{C}$ are "reasonable" subspaces by any standard.

Theorem (Mayer-Vietoris)**.** *Suppose $\mathscr{C}$ is a poset, $\mathscr{A}, \mathscr{B} \subseteqq \mathscr{C}$ are sieves, and $\mathscr{C} = \mathscr{A} \cup \mathscr{B}$. Then there is a Long Exact Sequence (LES)*

$$\cdots \to H_n(\mathscr{A} \cap \mathscr{B}) \xrightarrow{i_n} H_n(\mathscr{A}) \oplus H_n(\mathscr{B}) \xrightarrow{p_n} H_n(\mathscr{C}) \xrightarrow{d_n} H_{n-1}(\mathscr{A} \cap \mathscr{B}) \to \cdots .$$

This sequence extends infinitely to the left and right, but it is eventually 0 *as one goes out to the right.*

Proof. This is one of the few cases of a long exact sequence where everything is constructed concretely.[3] To define i_n, note that the inclusion maps $\mathscr{A} \cap \mathscr{B} \xrightarrow{i_A} \mathscr{A}$ and $\mathscr{A} \cap \mathscr{B} \xrightarrow{i_B} \mathscr{B}$ are (trivially) order-preserving. Since H_n is a functor, it induces group homomorphisms $H_n(i_A)$ and $H_n(i_B)$, which we take as the components of $i_n : H_n(\mathscr{A} \cap \mathscr{B}) \to H_n(\mathscr{A}) \oplus H_n(\mathscr{B})$. Explicitly,

$$i_n(\alpha) = (\alpha, \alpha).$$

Similarly, the inclusions $\mathscr{A} \xrightarrow{j_A} \mathscr{C}$ and $\mathscr{B} \xrightarrow{j_B} \mathscr{C}$ induce maps $H_n(j_A)$ and $H_n(j_B)$ on homology. Together, $H_n(j_A) : H_n(\mathscr{A}) \to H_n(\mathscr{C})$ and $-H_n(j_B) : H_n(\mathscr{B}) \to H_n(\mathscr{C})$ induce the homomorphism $p_n : H_n(\mathscr{A}) \oplus H_n(\mathscr{B}) \to H_n(\mathscr{C})$, which is given by the formula

$$p_n(\alpha, \beta) = \alpha - \beta.$$

Finally, for d_n we do the following. Suppose $\gamma \in H_n(\mathscr{C})$ is represented by the cycle $\sum_{i=1}^{k} a_i[x_0^i \leqq \ldots \leqq x_n^i] = \sum a_i x^i$. For each i, the vertex x_n^i is either in $\mathscr{A}$ or $\mathscr{B}$, because $\mathscr{C} = \mathscr{A} \cup \mathscr{B}$. Since $\mathscr{A}$ (resp. $\mathscr{B}$) is downward closed, it follows that if

[3]For most LESs, the "connecting homomorphism" $H_n \to H_{n-1}$ is constructed abstractly using the *Snake Lemma*.

$x^i_n \in \mathscr{A}$ (resp. $\mathscr{B}$), then $[x^i_0 \leqq \cdots \leqq x^i_n] \in C_n(\mathscr{A})$ (resp. $C_n(\mathscr{B})$). Thus, for each i, x^i is in $C_n(\mathscr{A})$, or $C_n(\mathscr{B})$, or both. Choose one of them for each i, and write

$$\sum_{i=1}^{k} a_i[x^i_0 \leqq \ldots \leqq x^i_n] = \sum_{x_i \in C_n(\mathscr{A})} a_i x^i + \sum_{x_i \in C_n(\mathscr{B})} a_i x^i = \alpha + \beta.$$

Consider $\partial\alpha$. Since $\sum a_i x^i$ is a cycle, $0 = \partial(\sum a_i x^i) = \partial(\alpha + \beta) = \partial\alpha + \partial\beta$. Therefore $\partial\alpha = -\partial\beta$ as $(n-1)$-cycles in $\mathscr{C}$. Since $\partial\alpha$ is a cycle in $\mathscr{A}$ and $\partial\beta$ is a cycle in $\mathscr{B}$, it follows $\partial\alpha$ is a cycle in $\mathscr{A} \cap \mathscr{B}$. W define $d_n(\sum a_i x^i)$ to be the homology class in $H_{n-1}(\mathscr{A} \cap \mathscr{B})$ represented by $\partial\alpha$. $\square$

Exercise: Check the details. We made a lot of (noncanonical) choices in the construction of d_n. Make sure the end results are independent of these choices.

2.4. **A Sample Calculation.** Let's look at how to compute the homology of S^2 using the Mayer-Vietoris sequence and the homology of S^1.

From the work of the previous lecture, we know that

$$H_n(S^1) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, 1 \\ 0 & \text{else} \end{cases} .$$

We decompose S^2 into two cones $\mathscr{A}, \mathscr{B} \cong CS^1$ based over S^1, with $\mathscr{A} \cap \mathscr{B} = S^1$ (two hemispheres). See section 2.1. The cones are contractible, therefore

$$H_n(CS^1) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0 \\ 0 & \text{else} \end{cases} .$$

Write down the Mayer-Vietoris sequence for $n = 2, 1, 0$ and plug in what we know:

$$0 \oplus 0 \to H_2(S^2) \to \mathbb{Z} \to 0 \oplus 0 \to H_1(S^2) \to \mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z} \to \mathbb{Z} \to 0$$

$0 \oplus 0 \cong 0$, so there are lots of 0's showing up in this *exact* sequence. That's a good thing. Do the following exercise:

Exercise: If $0 \to A \to B$ is exact, then $A \to B$ is injective. If $A \to B \to 0$ is exact, then $A \to B$ is surjective.

So, the exactness of $0 \to H_2(S^2) \to \mathbb{Z} \to 0$ implies $H_2(S^2) \cong \mathbb{Z}$. The exactness of of $0 \to H_1(S^2) \to \mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}$ implies

$$H_1(S^2) \cong \operatorname{im}(H_1(S^2) \to \mathbb{Z}) = \ker(\mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}).$$

By the definition of the Mayer-Vietoris sequence, we know that $\mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}$ corresponds to the "diagonal map" $H_0(S^1) \to H_0(CS^1) \oplus H_0(CS^1)$ sending $\alpha \mapsto (\alpha, \alpha)$. This is injective, so $\ker(\mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}) = 0$. Therefore $H^1(S^2) \cong 0$. We know S^2 is connected, so $H_0(S^2) \cong \mathbb{Z}$. Finally, since S^2 has no nondegenerate n-simplices for $n > 2$, $H_n(S^2) \cong 0$ for all $n > 2$. To summarize:

$$H_n(S^2) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, 2 \\ 0 & \text{else} \end{cases} .$$

Exercise: When showing $\mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}$ above is injective, did we need to use the specific formula given by the Mayer-Vietoris sequence, or would it have been enough to know everything is exact?

We pause to emphasize that the best way to learn these methods is to do things yourself. It is generally accepted that diagram chasing is mostly (only?) beneficial to the person doing the chasing, and is best done in the privacy of one's own home. In any event, try the following exercise.

Exercise: Understand what just happened. Now prove inductively that

$$H_k(S^n) \cong \begin{cases} \mathbb{Z} & k = 0, n \\ 0 & \text{else} \end{cases} .$$