Algebraic Topology

Homework 11: Due Wednesday, November 17

Problem 1. Compute the homology of the Lens space $L(p,q)$.

$L(p,q)$ is obtained from two solid tori, with the meridian of each torus identified with p times the longitude, plus or minus q or q' times the meridian, of the other.

Let U be a neighborhood of one solid torus, and let V be a neighborhood of the other. The only nonzero $\tilde H_k(U)$ is $\tilde H_1 = \mathbf{Z}$, and likewise for V. The intersection $U \cap V$ retracts to an ordinary torus, which has $H_2 = \mathbf{Z}$ and $H_1 = \mathbf{Z}^2$. Mayer-Vietoris then says $0 \to H_3(L(p,q)) \to \mathbf{Z} \to 0 \to H_2(L(p,q)) \to \mathbf{Z}^2 \to \mathbf{Z}^2 \to H_1(L(p,q)) \to 0 \to 0 \to \tilde H_0(L(p,q))$, so $H_3 = \mathbf{Z}$, $H_0 = \mathbf{Z}$, H_2 is the kernel of the map $\mathbf{Z}^2 \to \mathbf{Z}^2$ and H_1 is the cokernel.

Take as the generators of $H_1(U \cap V)$ the meridian and longitude of the boundary of the first solid torus. The first generator maps to $0 \in H_1(U)$ and to p in $H_1(V)$. The second generator maps to 1 in $H_1(U)$ and to something (doesn't matter what – call it n) in $H_1(V)$. This makes the map $\begin{pmatrix} 0 & 1 \\ p & n \end{pmatrix}$, whose kernel is trivial and whose cokernel is $\mathbf{Z}_p$, so $H_1(L(p,q)) = \mathbf{Z}_p$ and $H_2(L(p,q)) = 0$.

Note that H_1 of the Lens space is the Abelianization of π_1 (which was already Abelian).

Problem 2. Let X be a 2-holed torus with two disks glued in, exactly as in the midterm. Compute the homology of X.

Let T be the 2-holed torus, let U be be a neighborhood of T (the deformation retracts to T) and let V be a neighborhood of the two disks (that deformation retracts to two points). $U \cap V$ then retracts to two circles. Note that $U \cap V$ and V are both disconnected, and that $\tilde H_0(U \cap V)$ and $\tilde H_0(V)$ are both $\mathbf{Z}$, with the map from $\tilde H_0(U \cap V)$ to $\tilde H_0(V)$ induced the the inclusion $U \cap V \subset V$ being an isomorphism.

We already computed the (reduced) homology of T in class. We have $\tilde H_2(T) = \mathbf{Z}$, $\tilde H_1(T) = \mathbf{Z}^4$ (with generators a_1, b_1, a_2 and b_2 in the usual picture of T as an octagon with edges identified), and $\tilde H_0(T) = 0$. The Mayer-Vietoris sequence then reads:

$$0 \to \mathbf{Z} \to \tilde H_2(X) \to \mathbf{Z}^2 \xrightarrow{i_1} \mathbf{Z}^4 \to \tilde H_1(X) \to \mathbf{Z}^2 \xrightarrow{i_0} \mathbf{Z}^2 \to \tilde H_0(X) \to 0.$$

The key is understanding the maps i_1 and i_0 induced by inclusion of $U \cap V$ into U and V. We have already seen that i_0 is an isomorphism, so $\tilde{H}_0(X) = 0$, and the map $H_1(X) \to \mathbf{Z}^2$ is the trivial map. That is, our sequence shortens to

$$0 \to \mathbf{Z} \to \tilde{H}_2(X) \to \mathbf{Z}^2 \xrightarrow{i_1} \mathbf{Z}^4 \to \tilde{H}_1(X) \to 0.$$

Now, the generators of $H_1(U \cap V)$ are a_1 and a_2, so $i_1 : \mathbf{Z}^2 \to \mathbf{Z}^4$ is just $i_1(x, y)^T = (x, y, 0, 0)^T$. Since the kernel of this is trivial, the image of the map $\partial_* : \tilde{H}_2(X) \to \mathbf{Z}^2$ is trivial, so the kernel of ∂_* is all of $\tilde{H}_2(X)$, so $\tilde{H}_2(X)$ is isomorphic to $\mathbf{Z}$. Meanwhile, $\tilde{H}_1(X)$ is the cokernel of i_1, namely $\mathbf{Z}^2$.

Problem 3. Page 172, problem 5.2.

(a) If $H_n(X, A) = 0$ for all n, then we have $H_{n+1}(X, A) = 0 \to H_n(A) \xrightarrow{i_*} H_n(X) \to 0$, so i_* is an isomorphism. Conversely, if i_* is an isomorphism for all n, then the kernel of $j_* : H_n(X) \to H_n(X, A)$ is all of $H_n(X)$, so j_* is the zero map, and the image of j_* is trivial. However, the image of $\partial_* : H_n(X, A) \to H_{n-1}(A)$ is the kernel of $i_* : H_{n-1}(A) \to H_{n-1}(X)$, which is trivial, so the kernel of ∂_* is all of $H_n(X, A)$. Since the kernel of ∂_* equals the image of j_*, $H_n(X, A)$ must be trivial.

(b) This is the exact same argument shifted over by one. Explicitly:

If $H_n(A) = 0$ for all n, then we have $H_n(A) = 0 \to H_n(X) \xrightarrow{j_*} H_n(X, A) \to 0$, so j_* is an isomorphism. Conversely, if j_* is an isomorphism for all n, then the kernel of $\partial_* : H_n(X, A) \to H_{n-1}(A)$ is all of $H_n(X, A)$, so ∂_* is the zero map, and the image of ∂_* is trivial. However, the image of $i_* : H_n(A) \to H_n(X)$ is the kernel of $j_* : H_n(X) \to H_n(X, A)$, which is trivial, so the kernel of i_* is all of $H_n(A)$. Since the kernel of i_* equals the image of ∂_*, $H_n(A)$ must be trivial.

(c) If $H_n(X, A) = 0$ for all $n \leq q$, then we have $H_{n+1}(X, A) = 0 \to H_n(A) \xrightarrow{i_*} H_n(X) \to 0$ for all $n < q$, so i_* is an isomorphism for $n < q$. Furthermore, we have $H_q(A) \to H_q(X) \to 0$, so $i_* : H_q(A) \to H_q(X)$ is onto.

Conversely, if i_* is an isomorphism for all $n < q$ and is onto for $n = q$, then the kernel of $j_* : H_n(X) \to H_n(X, A)$ is all of $H_n(X)$ for $n \leq q$, so j_* is the zero map, and the image of j_* is trivial. However, the image of $\partial_* : H_n(X, A) \to H_{n-1}(A)$ is the kernel of $i_* : H_{n-1}(A) \to H_{n-1}(X)$, which is trivial for all $n \leq q$, so the kernel of ∂_* is all of $H_n(X, A)$. Since the kernel of ∂_* equals the image of j_*, $H_n(X, A)$ must be trivial for all $n \leq q$.

Problem 4. Page 172, problem 5.4.

By definition $Z_n(X, A)$ is the set of classes in $C_k(X, A)$ whose boundary is zero in $C_{n-1}(X, A)$. That is, the set of classes $[\alpha]$ (with $\alpha \in C_n(X)$ such that $0 = [\partial\alpha]$, i.e. for which $\partial\alpha \in C_{n-1}(A)$. But the set of classes (mod $C_n(A)$) of chains in X whose boundaries are in $C_{n-1}(A)$ is exactly $Z_n(X \bmod A)/C_n(A)$.

$B_n(X, A)$ is the image of $\partial : C_{n+1}(X, A) \to C_n(X, A)$. That is, it is the set of all classes $[\partial\alpha]$, where $\alpha \in C_{n+1}(X)$, which is the same as the set of all classes $[\partial\alpha + \beta]$, where $\beta \in C_n(A)$. But that's exactly $(B_n(X) + C_n(A))/C_n(A)$.

Since every class can be represented by an element of $B_n(X)$, we need only check when two such classes $[\partial\alpha]$ and $[\partial\alpha']$ are the same. That would be if $\partial\alpha - \partial\alpha' \in C_n(A)$. But $\partial\alpha - \partial\alpha'$ is automatically in $B_n(X)$, so the difference has to be in $B_n(X) \cap C_n(A)$. That is, we have $B_n(X)/(B_n(X) \cap C_n(A))$.

The third statement follows from the first two. A class in $H_n(X, A)$ is represented by an element of $Z_n(X, A)$, and two elements are the same if they differ by an element of $B_n(X, A)$. This means that every class in $H_n(X, A)$ is represented by an element of $Z_n(X \bmod A)$, and two such representatives are equivalent if they differ by an element of $C_n(A)$ (which doesn't change the class in $Z_n(X, A)$), or by an element of $B_n(X)$ (representing a class in $B_n(X, A)$), or by a linear combination of the two.

Problem 5. Page 175, problem 6.2.

In the exact sequence $0 \to A \xrightarrow{\phi} G \xrightarrow{\psi} B \to 0$, we have $A = \mathbf{Z}$ and $B = \mathbf{Z}_n$. The question is what G and ψ are. (Once we have G and ψ, then, up to isomorphism, A is then the kernel of ψ and ϕ is the inclusion map.)

If $n = ab$, with a and b positive integers, and if k is relatively prime to b, then we can take $G = \mathbf{Z} \oplus \mathbf{Z}_a$ and $\psi(x, y) = kx + by$. The kernel of ψ is all multiples of $(-b, k)$, so we can take $\phi(1) = (-b, k)$. I claim that every solution is of this form.

Note that G is a finitely generated Abelian group, and so is of the form $\mathbf{Z}^k \oplus T$, where T is the torsion subgroup of G. Since ϕ is injective and A is free, no element of T can be in the image of ϕ. This implies that ψ restricted to T is injective. In particular, T is a subgroup of $\mathbf{Z}_n$, so $T = \mathbf{Z}_a$, where a divides n. Let $b = n/a$. Then, up to a change in the choice of generator of $\mathbf{Z}_a$, we have $\psi(0, 1) = b$. Since ψ is onto, $\psi(1, 0)$ must be relatively prime to

b. Defining $k = \psi(1,0)$, we have the solution indicated above.

The only remaining question is when two solutions are isomorphic. Clearly different values of a give different groups. But do different values of k give the same map, up to isomorphism? I believe the answer is "yes", but haven't yet managed to prove it.

Problem 6. Suppose that $A \subset B \subset X$. Show that there is a long exact sequence $\to H_k(B,A) \to H_k(X,A) \to H_k(X,B) \to H_{k-1}(B,A) \to \cdots$. What are the maps in this sequence? [This sequence is called the long exact sequence of the triple (X,B,A). When A is the empty set, it becomes the usual long exact sequence of the pair (X,B).]

We must construct a short exact sequence of chain complexes

$$0 \to C_k(B,A) \xrightarrow{i} C_k(X,A) \xrightarrow{j} C_k(X,B) \to 0,$$

from which the snake lemma will give us the long exact sequence in homology.

There is an inclusion $C_k(B) \to C_k(X)$. Taking the quotient of both sides by $C_k(A)$ gives the inclusion i. Likewise, there is a projection j from $C_k(X,A)$ to $C_k(X,B)$, where the class of a chain in X, mod chains in A, is mapped to its class mod chains in B. This is clearly surjective, since if $[\alpha] \in C_k(X,B)$, then $\alpha \in C_k(X)$, and the class of α in $C_k(X,B)$ is just the image of the class of α in $C_k(X,A)$.

What remains is to show exactness at $C_k(X,A)$. If $\alpha \in C_k(X)$ and $j[\alpha] = 0$, then α must be in $C_k(B)$, so the class of α in $C_k(X,A)$ is i of the class of α in $C_k(B,A)$. Conversely, if $[\beta]$ is a class in $C_k(B,A)$, then β is a chain in B, and $j(i([\beta]))$ is zero.

The maps i_*, j_* and ∂_* in the long exact sequence are as follows: i_* is just induced by inclusion of B in X, j_* is induced by inclusion of A in B. As for ∂_*, suppose that $\gamma \in Z_n(X \text{ mod } B)$. That is $\gamma \in C_k(X)$ and $\partial\gamma \in C_k(B)$. Pull this back to $C_k(X,A)$ by looking at the class of γ in $C_k(X,A)$. Then take the boundary of γ. Then think of this as coming from a chain on B. In other words, $\partial_*[\gamma]$ is just $[\partial\gamma]$, viewed as a class in $H_{n-1}(B,A)$.

The usual long exact sequence of the pair (X,B) can be viewed as the exact sequence of the triple $(X,B,\emptyset)$, since $C_k(X,\emptyset) = C_k(X)$ and $C_k(B,\emptyset) = C_k(B)$.