Chapter 2

Homology

The fundamental group $\pi_1(X)$ is especially useful when studying spaces of low dimension, as one would expect from its definition which involves only maps from low-dimensional spaces into X, namely loops $I \to X$ and homotopies of loops, maps $I \times I \to X$. The definition in terms of objects that are at most 2-dimensional manifests itself for example in the fact that when X is a CW complex, $\pi_1(X)$ depends only on the 2-skeleton of X. In view of the low-dimensional nature of the fundamental group, we should not expect it to be a very refined tool for dealing with high-dimensional spaces. Thus it cannot distinguish between spheres S^n with $n \ge 2$. This limitation to low dimensions can be removed by considering the natural higher-dimensional analogs of $\pi_1(X)$, the homotopy groups $\pi_n(X)$, which are defined in terms of maps of the n-dimensional cube I^n into X and homotopies $I^n \times I \to X$ of such maps. Not surprisingly, when X is a CW complex, $\pi_n(X)$ depends only on the $(n+1)$-skeleton of X. And as one might hope, homotopy groups do indeed distinguish spheres of all dimensions since $\pi_i(S^n)$ is 0 for $i < n$ and $\mathbb{Z}$ for $i = n$.

However, the higher-dimensional homotopy groups have the serious drawback that they are extremely difficult to compute in general. Even for simple spaces like spheres, the calculation of $\pi_i(S^n)$ for $i > n$ turns out to be a huge problem. Fortunately there is a more computable alternative to homotopy groups: the homology groups $H_n(X)$. Like $\pi_n(X)$, the homology group $H_n(X)$ for a CW complex X depends only on the $(n+1)$-skeleton. For spheres, the homology groups $H_i(S^n)$ are isomorphic to the homotopy groups $\pi_i(S^n)$ in the range $1 \le i \le n$, but homology groups have the advantage that $H_i(S^n) = 0$ for $i > n$.

The computability of homology groups does not come for free, unfortunately. The definition of homology groups is decidedly less transparent than the definition of homotopy groups, and once one gets beyond the definition there is a certain amount of technical machinery to be set up before any real calculations and applications can be given. In the exposition below we approach the definition of $H_n(X)$ by two preliminary stages, first giving a few motivating examples nonrigorously, then constructing a restricted model of homology theory called simplicial homology, before plunging into the general theory, known as singular homology. After the definition of singular homology has been assimilated, the real work of establishing its basic properties begins. This takes close to 20 pages, and there is no getting around the fact that it is a substantial effort. This takes up most of the first section of the chapter, with small digressions only for two applications to classical theorems of Brouwer: the fixed point theorem and 'invariance of dimension.'

The second section of the chapter gives more applications, including the homology definition of Euler characteristic and Brouwer's notion of degree for maps $S^n \to S^n$. However, the main thrust of this section is toward developing techniques for calculating homology groups efficiently. The maximally efficient method is known as cellular homology, whose power comes perhaps from the fact that it is 'homology squared' — homology defined in terms of homology. Another quite useful tool is Mayer–Vietoris sequences, the analog for homology of van Kampen's theorem for the fundamental group.

An interesting feature of homology that begins to emerge after one has worked with it for a while is that it is the basic properties of homology that are used most often, and not the actual definition itself. This suggests that an axiomatic approach to homology might be possible. This is indeed the case, and in the third section of the chapter we list axioms which completely characterize homology groups for CW complexes. One could take the viewpoint that these rather algebraic axioms are all that really matters about homology groups, that the geometry involved in the definition of homology is secondary, needed only to show that the axiomatic theory is not vacuous. The extent to which one adopts this viewpoint is a matter of taste, and the route taken here of postponing the axioms until the theory is well-established is just one of several possible approaches.

The chapter then concludes with three optional sections of Additional Topics. The first is rather brief, relating $H_1(X)$ to $\pi_1(X)$, while the other two contain a selection of classical applications of homology. These include the n-dimensional version of the Jordan curve theorem and the 'invariance of domain' theorem, both due to Brouwer, along with the Lefschetz fixed point theorem.

The Idea of Homology

The difficulty with the higher homotopy groups π_n is that they are not directly computable from a cell structure as π_1 is. For example, the 2-sphere has no cells in dimensions greater than 2, yet its n-dimensional homotopy group $\pi_n(S^2)$ is nonzero for infinitely many values of n. Homology groups, by contrast, are quite directly related to cell structures, and may indeed be regarded as simply an algebraization of the first layer of geometry in cell structures: how cells of dimension n attach to cells of dimension $n-1$.

Let us look at some examples to see what the idea is. Consider the graph X_1 shown in the figure, consisting of two vertices joined by four edges. When studying the fundamental group of X_1 we consider loops formed by sequences of edges, starting and ending at a fixed basepoint. For example, at the basepoint x, the loop ab^{-1} travels forward along the edge a, then backward along b, as indicated by the exponent -1. A more complicated loop would be $ac^{-1}bd^{-1}ca^{-1}$. A salient feature of the fundamental group is that it is generally nonabelian, which both enriches and complicates the theory. Suppose we simplify matters by abelianizing. Thus for example the two loops ab^{-1} and $b^{-1}a$ are to be regarded as equal if we make a commute with b^{-1}. These two loops ab^{-1} and $b^{-1}a$ are really the same circle, just with a different choice of starting and ending point: x for ab^{-1} and y for $b^{-1}a$. The same thing happens for all loops: Rechoosing the basepoint in a loop just permutes its letters cyclically, so a byproduct of abelianizing is that we no longer have to pin all our loops down to a fixed basepoint. Thus loops become *cycles*, without a chosen basepoint.

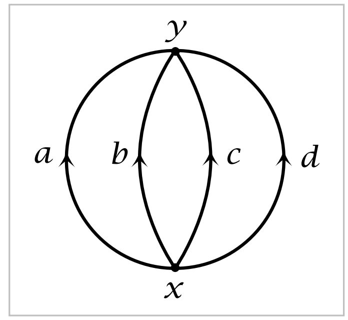

Having abelianized, let us switch to additive notation, so cycles become linear combinations of edges with integer coefficients, such as $a - b + c - d$. Let us call these linear combinations *chains* of edges. Some chains can be decomposed into cycles in several different ways, for example $(a - c) + (b - d) = (a - d) + (b - c)$, and if we adopt an algebraic viewpoint then we do not want to distinguish between these different decompositions. Thus we broaden the meaning of the term 'cycle' to be simply any linear combination of edges for which at least one decomposition into cycles in the previous more geometric sense exists.

What is the condition for a chain to be a cycle in this more algebraic sense? A geometric cycle, thought of as a path traversed in time, is distinguished by the property that it enters each vertex the same number of times that it leaves the vertex. For an arbitrary chain $ka + \ell b + mc + nd$, the net number of times this chain enters y is $k + \ell + m + n$ since each of a, b, c, and d enters y once. Similarly, each of the four edges leaves x once, so the net number of times the chain $ka + \ell b + mc + nd$ enters x is $-k - \ell - m - n$. Thus the condition for $ka + \ell b + mc + nd$ to be a cycle is simply $k + \ell + m + n = 0$.

To describe this result in a way that would generalize to all graphs, let C_1 be the free abelian group with basis the edges a, b, c, d and let C_0 be the free abelian group with basis the vertices x, y. Elements of C_1 are chains of edges, or 1-dimensional chains, and elements of C_0 are linear combinations of vertices, or 0-dimensional chains. Define a homomorphism $\partial : C_1 \to C_0$ by sending each basis element a, b, c, d to $y - x$, the vertex at the head of the edge minus the vertex at the tail. Thus we have $\partial(ka + \ell b + mc + nd) = (k + \ell + m + n)y - (k + \ell + m + n)x$, and the cycles are precisely the kernel of ∂. It is a simple calculation to verify that $a - b$, $b - c$, and $c - d$

form a basis for this kernel. Thus every cycle in X_1 is a unique linear combination of these three most obvious cycles. By means of these three basic cycles we convey the geometric information that the graph X_1 has three visible 'holes,' the empty spaces between the four edges.

Let us now enlarge the preceding graph X_1 by attaching a 2-cell A along the cycle $a - b$, producing a 2-dimensional cell complex X_2. If we think of the 2-cell A as being oriented clockwise, then we can regard its boundary as the cycle $a - b$. This cycle is now homotopically trivial since we can contract it to a point by sliding over A. In other words, it no longer encloses a hole in X_2. This suggests that we form a quotient of the group of cycles in the preceding example by factoring out the subgroup generated by $a - b$. In this quotient the cycles $a - c$ and $b - c$, for example, become equivalent, consistent with the fact that they are homotopic in X_2.

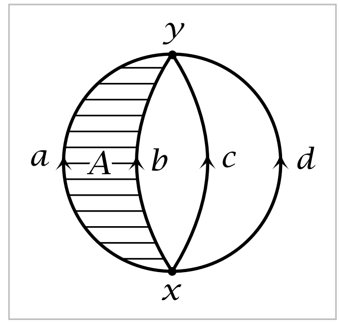

Algebraically, we can define now a pair of homomorphisms $C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0$ where C_2 is the infinite cyclic group generated by A and $\partial_2(A) = a - b$. The map ∂_1 is the boundary homomorphism in the previous example. The quotient group we are interested in is $\operatorname{Ker}\partial_1 / \operatorname{Im}\partial_2$, the kernel of ∂_1 modulo the image of ∂_2, or in other words, the 1-dimensional cycles modulo those that are boundaries, the multiples of $a - b$. This quotient group is the *homology group* $H_1(X_2)$. The previous example can be fit into this scheme too by taking C_2 to be zero since there are no 2-cells in X_1, so in this case $H_1(X_1) = \operatorname{Ker}\partial_1 / \operatorname{Im}\partial_2 = \operatorname{Ker}\partial_1$, which as we saw was free abelian on three generators. In the present example, $H_1(X_2)$ is free abelian on two generators, $b - c$ and $c - d$, expressing the geometric fact that by filling in the 2-cell A we have reduced the number of 'holes' in our space from three to two.

Suppose we enlarge X_2 to a space X_3 by attaching a second 2-cell B along the same cycle $a - b$. This gives a 2-dimensional chain group C_2 consisting of linear combinations of A and B, and the boundary homomorphism $\partial_2 : C_2 \to C_1$ sends both A and B to $a - b$. The homology group $H_1(X_3) = \operatorname{Ker}\partial_1 / \operatorname{Im}\partial_2$ is the same as for X_2, but now ∂_2 has a nontrivial kernel, the infinite cyclic group generated by $A - B$. We view $A - B$ as a 2-dimensional cycle, generating the homology group $H_2(X_3) = \operatorname{Ker}\partial_2 \approx \mathbb{Z}$. Topologically, the cycle $A - B$ is the sphere formed by the cells A and B together with their common boundary circle. This spherical cycle detects the presence of a 'hole' in X_3, the missing interior of the sphere. However, since this hole is enclosed by a sphere rather than a circle, it is of a different sort from the holes detected by $H_1(X_3) \approx \mathbb{Z} \times \mathbb{Z}$, which are detected by the cycles $b - c$ and $c - d$.

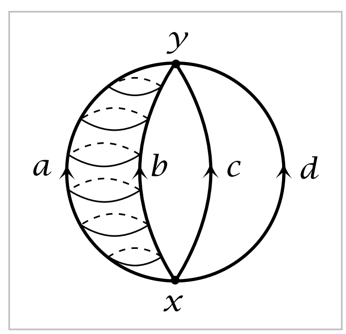

Let us continue one more step and construct a complex X_4 from X_3 by attaching a 3-cell C along the 2-sphere formed by A and B. This creates a chain group C_3

generated by this 3-cell C, and we define a boundary homomorphism $\partial_3 : C_3 \to C_2$ sending C to $A - B$ since the cycle $A - B$ should be viewed as the boundary of C in the same way that the 1-dimensional cycle $a - b$ is the boundary of A. Now we have a sequence of three boundary homomorphisms $C_3 \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0$ and the quotient $H_2(X_4) = \mathrm{Ker}\,\partial_2 / \mathrm{Im}\,\partial_3$ has become trivial. Also $H_3(X_4) = \mathrm{Ker}\,\partial_3 = 0$. The group $H_1(X_4)$ is the same as $H_1(X_3)$, namely $\mathbb{Z}\times\mathbb{Z}$, so this is the only nontrivial homology group of X_4.

It is clear what the general pattern of the examples is. For a cell complex X one has chain groups $C_n(X)$ which are free abelian groups with basis the n-cells of X, and there are boundary homomorphisms $\partial_n : C_n(X) \to C_{n-1}(X)$, in terms of which one defines the homology group $H_n(X) = \mathrm{Ker}\,\partial_n / \mathrm{Im}\,\partial_{n+1}$. The major difficulty is how to define ∂_n in general. For $n = 1$ this is easy: The boundary of an oriented edge is the vertex at its head minus the vertex at its tail. The next case $n = 2$ is also not hard, at least for cells attached along cycles that are simply loops of edges, for then the boundary of the cell is this cycle of edges, with the appropriate signs taking orientations into account. But for larger n, matters become more complicated. Even if one restricts attention to cell complexes formed from polyhedral cells with nice attaching maps, there is still the matter of orientations to sort out.

The best solution to this problem seems to be to adopt an indirect approach. Arbitrary polyhedra can always be subdivided into special polyhedra called simplices (the triangle and the tetrahedron are the 2-dimensional and 3-dimensional instances) so there is no loss of generality, though initially there is some loss of efficiency, in restricting attention entirely to simplices. For simplices there is no difficulty in defining boundary maps or in handling orientations. So one obtains a homology theory, called simplicial homology, for cell complexes built from simplices. Still, this is a rather restricted class of spaces, and the theory itself has a certain rigidity that makes it awkward to work with.

The way around these obstacles is to step back from the geometry of spaces decomposed into simplices and to consider instead something which at first glance seems wildly more complicated, the collection of all possible continuous maps of simplices into a given space X. These maps generate tremendously large chain groups $C_n(X)$, but the quotients $H_n(X) = \mathrm{Ker}\,\partial_n / \mathrm{Im}\,\partial_{n+1}$, called singular homology groups, turn out to be much smaller, at least for reasonably nice spaces X. In particular, for spaces like those in the four examples above, the singular homology groups coincide with the homology groups we computed from the cellular chains. And as we shall see later in this chapter, singular homology allows one to define these nice cellular homology groups for all cell complexes, and in particular to solve the problem of defining the boundary maps for cellular chains.

2.1 Simplicial and Singular Homology

The most important homology theory in algebraic topology, and the one we shall be studying almost exclusively, is called singular homology. Since the technical apparatus of singular homology is somewhat complicated, we will first introduce a more primitive version called simplicial homology in order to see how some of the apparatus works in a simpler setting before beginning the general theory.

The natural domain of definition for simplicial homology is a class of spaces we call Δ-complexes, which are a mild generalization of the more classical notion of a simplicial complex. Historically, the modern definition of singular homology was first given in [Eilenberg 1944], and Δ-complexes were introduced soon thereafter in [Eilenberg-Zilber 1950] where they were called semisimplicial complexes. Within a few years this term came to be applied to what Eilenberg and Zilber called complete semisimplicial complexes, and later there was yet another shift in terminology as the latter objects came to be called simplicial sets. In theory this frees up the term semisimplicial complex to have its original meaning, but to avoid potential confusion it seems best to introduce a new name, and the term Δ-complex has at least the virtue of brevity.

Δ-Complexes

The torus, the projective plane, and the Klein bottle can each be obtained from a square by identifying opposite edges in the way indicated by the arrows in the following figures:

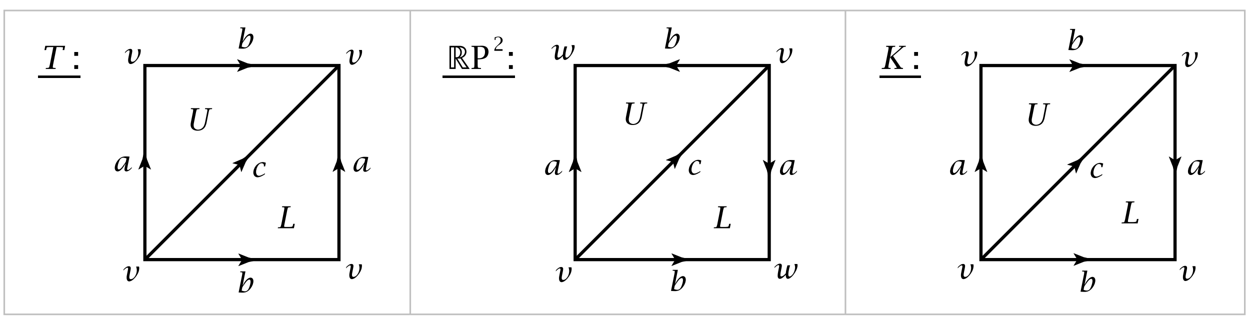

Cutting a square along a diagonal produces two triangles, so each of these surfaces can also be built from two triangles by identifying their edges in pairs. In similar fashion a polygon with any number of sides can be cut along diagonals into triangles, so in fact all closed surfaces can be constructed from triangles by identifying edges. Thus we have a single building block, the triangle, from which all surfaces can be constructed. Using only triangles we could also construct a large class of 2-dimensional spaces that are not surfaces in the strict sense, by allowing more than two edges to be identified together at a time.

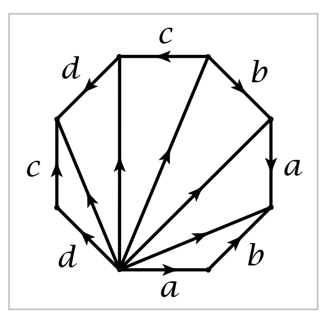

The idea of a Δ-complex is to generalize constructions like these to any number of dimensions. The n-dimensional analog of the triangle is the **n-simplex**. This is the smallest convex set in a Euclidean space $\mathbb{R}^m$ containing $n+1$ points $v_0, \cdots, v_n$ that do not lie in a hyperplane of dimension less than n, where by a hyperplane we mean the set of solutions of a system of linear equations. An equivalent condition would be that the difference vectors $v_1 - v_0, \cdots, v_n - v_0$ are linearly independent. The points v_i are the **vertices** of the simplex, and the simplex itself is denoted $[v_0, \cdots, v_n]$. For example, there is the standard n-simplex

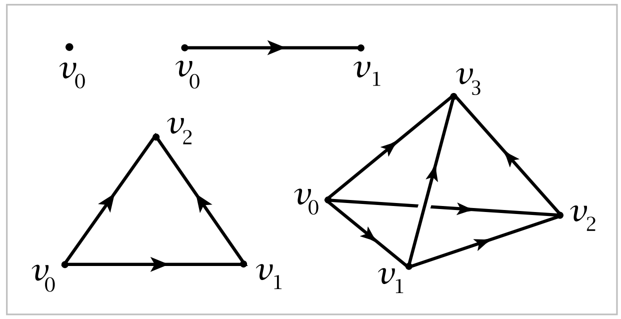

$$\Delta^n = \{\, (t_0, \cdots, t_n) \in \mathbb{R}^{n+1} \mid \textstyle\sum_i t_i = 1 \text{ and } t_i \geq 0 \text{ for all } i \,\}$$

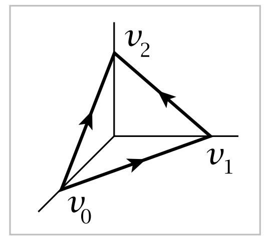

whose vertices are the unit vectors along the coordinate axes.

For purposes of homology it will be important to keep track of the order of the vertices of a simplex, so 'n-simplex' will really mean 'n-simplex with an ordering of its vertices.' A by-product of ordering the vertices of a simplex $[v_0, \cdots, v_n]$ is that this determines orientations of the edges $[v_i, v_j]$ according to increasing subscripts, as shown in the two preceding figures. Specifying the ordering of the vertices also determines a canonical linear homeomorphism from the standard n-simplex Δ^n onto any other n-simplex $[v_0, \cdots, v_n]$, preserving the order of vertices, namely, $(t_0, \cdots, t_n) \mapsto \sum_i t_i v_i$. The coefficients t_i are the **barycentric coordinates** of the point $\sum_i t_i v_i$ in $[v_0, \cdots, v_n]$.

If we delete one of the $n+1$ vertices of an n-simplex $[v_0, \cdots, v_n]$, then the remaining n vertices span an $(n-1)$-simplex, called a **face** of $[v_0, \cdots, v_n]$. We adopt the following convention:

The vertices of a face, or of any subsimplex spanned by a subset of the vertices, will always be ordered according to their order in the larger simplex.

The union of all the faces of Δ^n is the **boundary** of Δ^n, written $\partial\Delta^n$. The **open simplex** $\mathring{\Delta}^n$ is $\Delta^n - \partial\Delta^n$, the interior of Δ^n.

A **Δ-complex** structure on a space X is a collection of maps $\sigma_\alpha : \Delta^n \to X$, with n depending on the index α, such that:

(i) The restriction $\sigma_\alpha \,|\, \mathring{\Delta}^n$ is injective, and each point of X is in the image of exactly one such restriction $\sigma_\alpha \,|\, \mathring{\Delta}^n$.

(ii) Each restriction of σ_α to a face of Δ^n is one of the maps $\sigma_\beta : \Delta^{n-1} \to X$. Here we are identifying the face of Δ^n with Δ^{n-1} by the canonical linear homeomorphism between them that preserves the ordering of the vertices.

(iii) A set $A \subset X$ is open iff $\sigma_\alpha^{-1}(A)$ is open in Δ^n for each σ_α.

Among other things, this last condition rules out trivialities like regarding all the points of X as individual vertices. The earlier decompositions of the torus, projective plane, and Klein bottle into two triangles, three edges, and one or two vertices define Δ-complex structures with a total of six σ_α's for the torus and Klein bottle and seven for the projective plane. The orientations on the edges in the pictures are compatible with a unique ordering of the vertices of each simplex, and these orderings determine the maps σ_α.

A consequence of (iii) is that X can be built as a quotient space of a collection of disjoint simplices Δ^n_α, one for each $\sigma_\alpha : \Delta^n \to X$, the quotient space obtained by identifying each face of a Δ^n_α with the Δ^{n-1}_β corresponding to the restriction σ_β of σ_α to the face in question, as in condition (ii). One can think of building the quotient space inductively, starting with a discrete set of vertices, then attaching edges to these to produce a graph, then attaching 2-simplices to the graph, and so on. From this viewpoint we see that the data specifying a Δ-complex can be described purely combinatorially as collections of n-simplices Δ^n_α for each n together with functions associating to each face of each n-simplex Δ^n_α an $(n-1)$-simplex Δ^{n-1}_β.

More generally, Δ-complexes can be built from collections of disjoint simplices by identifying various subsimplices spanned by subsets of the vertices, where the identifications are performed using the canonical linear homeomorphisms that preserve the orderings of the vertices. The earlier Δ-complex structures on a torus, projective plane, or Klein bottle can be obtained in this way, by identifying pairs of edges of two 2-simplices. If one starts with a single 2-simplex and identifies all three edges to a single edge, preserving the orientations given by the ordering of the vertices, this produces a Δ-complex known as the 'dunce hat.' By contrast, if the three edges of a 2-simplex are identified preserving a cyclic orientation of the three edges, as in the first figure at the right, this does not produce a Δ-complex structure, although if the 2-simplex is subdivided into three smaller 2-simplices about a central vertex, then one does obtain a Δ-complex structure on the quotient space.

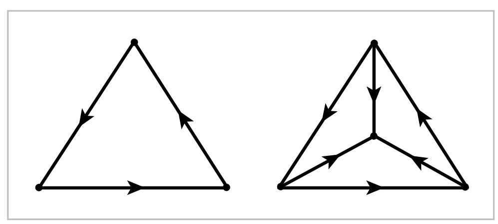

Thinking of a Δ-complex X as a quotient space of a collection of disjoint simplices, it is not hard to see that X must be a Hausdorff space. Condition (iii) then implies that each restriction $\sigma_\alpha | \mathring{\Delta}^n$ is a homeomorphism onto its image, which is thus an open simplex in X. It follows from Proposition A.2 in the Appendix that these open simplices $\sigma_\alpha(\mathring{\Delta}^n)$ are the cells e^n_α of a CW complex structure on X with the σ_α's as characteristic maps. We will not need this fact at present, however.

Simplicial Homology

Our goal now is to define the simplicial homology groups of a Δ-complex X. Let $\Delta_n(X)$ be the free abelian group with basis the open n-simplices e^n_α of X. Elements of $\Delta_n(X)$, called ***n*-chains**, can be written as finite formal sums $\sum_\alpha n_\alpha e_\alpha^n$ with coefficients $n_\alpha \in \mathbb{Z}$. Equivalently, we could write $\sum_\alpha n_\alpha \sigma_\alpha$ where $\sigma_\alpha : \Delta^n \to X$ is the characteristic map of e_α^n, with image the closure of e_α^n as described above. Such a sum $\sum_\alpha n_\alpha \sigma_\alpha$ can be thought of as a finite collection, or 'chain,' of n-simplices in X with integer multiplicities, the coefficients n_α.

As one can see in the next figure, the boundary of the n-simplex $[v_0, \cdots, v_n]$ consists of the various $(n-1)$-dimensional simplices $[v_0, \cdots, \hat{v}_i, \cdots, v_n]$, where the 'hat' symbol $\hat{}$ over v_i indicates that this vertex is deleted from the sequence $v_0, \cdots, v_n$. In terms of chains, we might then wish to say that the boundary of $[v_0, \cdots, v_n]$ is the $(n-1)$-chain formed by the sum of the faces $[v_0, \cdots, \hat{v}_i, \cdots, v_n]$. However, it turns out to be better to insert certain signs and instead let the boundary of $[v_0, \cdots, v_n]$ be $\sum_i (-1)^i [v_0, \cdots, \hat{v}_i, \cdots, v_n]$. Heuristically, the signs are inserted to take orientations into account, so that all the faces of a simplex are coherently oriented, as indicated in the following figure:

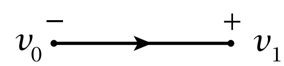

$$\partial[v_0, v_1] = [v_1] - [v_0]$$

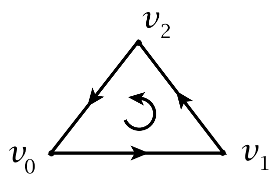

$$\partial[v_0, v_1, v_2] = [v_1, v_2] - [v_0, v_2] + [v_0, v_1]$$

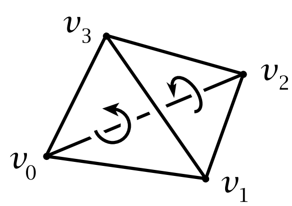

$$\begin{aligned}\partial[v_0, v_1, v_2, v_3] = [v_1, v_2, v_3] &- [v_0, v_2, v_3] \\ &+ [v_0, v_1, v_3] - [v_0, v_1, v_2]\end{aligned}$$

In the last case, the orientations of the two hidden faces are also counterclockwise when viewed from outside the 3-simplex.

With this geometry in mind we define for a general Δ-complex X a **boundary homomorphism** $\partial_n : \Delta_n(X) \to \Delta_{n-1}(X)$ by specifying its values on basis elements:

$$\partial_n(\sigma_\alpha) = \sum_i (-1)^i \sigma_\alpha \,|\, [v_0, \cdots, \hat{v}_i, \cdots, v_n]$$

Note that the right side of this equation does indeed lie in $\Delta_{n-1}(X)$ since each restriction $\sigma_\alpha \,|\, [v_0, \cdots, \hat{v}_i, \cdots, v_n]$ is the characteristic map of an $(n-1)$-simplex of X.

Lemma 2.1. *The composition* $\Delta_n(X) \xrightarrow{\partial_n} \Delta_{n-1}(X) \xrightarrow{\partial_{n-1}} \Delta_{n-2}(X)$ *is zero.*

Proof: We have $\partial_n(\sigma) = \sum_i (-1)^i \sigma \,|\, [v_0, \cdots, \hat{v}_i, \cdots, v_n]$, and hence

$$\begin{aligned}\partial_{n-1}\partial_n(\sigma) &= \sum_{j<i} (-1)^i (-1)^j \sigma | [v_0, \cdots, \hat{v}_j, \cdots, \hat{v}_i, \cdots, v_n] \\ &\quad + \sum_{j>i} (-1)^i (-1)^{j-1} \sigma | [v_0, \cdots, \hat{v}_i, \cdots, \hat{v}_j, \cdots, v_n]\end{aligned}$$

The latter two summations cancel since after switching i and j in the second sum, it becomes the negative of the first. $\square$

The algebraic situation we have now is a sequence of homomorphisms of abelian groups

$$\cdots \longrightarrow C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \longrightarrow \cdots \longrightarrow C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0$$

with $\partial_n\partial_{n+1} = 0$ for each n. Such a sequence is called a **chain complex**. Note that we have extended the sequence by a 0 at the right end, with $\partial_0 = 0$. The equation $\partial_n\partial_{n+1} = 0$ is equivalent to the inclusion $\operatorname{Im}\partial_{n+1} \subset \operatorname{Ker}\partial_n$, where Im and Ker denote image and kernel. So we can define the n^{th} **homology group** of the chain complex to be the quotient group $H_n = \operatorname{Ker}\partial_n / \operatorname{Im}\partial_{n+1}$. Elements of $\operatorname{Ker}\partial_n$ are called **cycles** and elements of $\operatorname{Im}\partial_{n+1}$ are called **boundaries**. Elements of H_n are cosets of $\operatorname{Im}\partial_{n+1}$, called **homology classes**. Two cycles representing the same homology class are said to be **homologous**. This means their difference is a boundary.

Returning to the case that $C_n = \Delta_n(X)$, the homology group $\operatorname{Ker}\partial_n / \operatorname{Im}\partial_{n+1}$ will be denoted $H_n^\Delta(X)$ and called the n^{th} **simplicial homology group** of X.

Example 2.2. $X = S^1$, with one vertex v and one edge e. Then $\Delta_0(S^1)$ and $\Delta_1(S^1)$ are both $\mathbb{Z}$ and the boundary map ∂_1 is zero since $\partial e = v - v$. The groups $\Delta_n(S^1)$ are 0 for $n \ge 2$ since there are no simplices in these dimensions. Hence

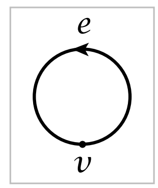

$$H_n^\Delta(S^1) \approx \begin{cases} \mathbb{Z} & \text{for } n = 0,1 \\ 0 & \text{for } n \ge 2 \end{cases}$$

This is an illustration of the general fact that if the boundary maps in a chain complex are all zero, then the homology groups of the complex are isomorphic to the chain groups themselves.

Example 2.3. $X = T$, the torus with the Δ-complex structure pictured earlier, having one vertex, three edges a, b, and c, and two 2-simplices U and L. As in the previous example, $\partial_1 = 0$ so $H_0^\Delta(T) \approx \mathbb{Z}$. Since $\partial_2 U = a + b - c = \partial_2 L$ and $\{a, b, a+b-c\}$ is a basis for $\Delta_1(T)$, it follows that $H_1^\Delta(T) \approx \mathbb{Z}\oplus\mathbb{Z}$ with basis the homology classes $[a]$ and $[b]$. Since there are no 3-simplices, $H_2^\Delta(T)$ is equal to $\operatorname{Ker}\partial_2$, which is infinite cyclic generated by $U - L$ since $\partial(pU + qL) = (p+q)(a+b-c) = 0$ only if $p = -q$. Thus

$$H_n^\Delta(T) \approx \begin{cases} \mathbb{Z}\oplus\mathbb{Z} & \text{for } n = 1 \\ \mathbb{Z} & \text{for } n = 0,2 \\ 0 & \text{for } n \ge 3 \end{cases}$$

Example 2.4. $X = \mathbb{R}\mathrm{P}^2$, as pictured earlier, with two vertices v and w, three edges a, b, and c, and two 2-simplices U and L. Then $\operatorname{Im}\partial_1$ is generated by $w - v$, so $H_0^\Delta(X) \approx \mathbb{Z}$ with either vertex as a generator. Since $\partial_2 U = -a+b+c$ and $\partial_2 L = a-b+c$, we see that ∂_2 is injective, so $H_2^\Delta(X) = 0$. Further, $\operatorname{Ker}\partial_1 \approx \mathbb{Z}\oplus\mathbb{Z}$ with basis $a - b$ and c, and $\operatorname{Im}\partial_2$ is an index-two subgroup of $\operatorname{Ker}\partial_1$ since we can choose c and $a - b + c$

as a basis for $\operatorname{Ker}\partial_1$ and $a-b+c$ and $2c=(a-b+c)+(-a+b+c)$ as a basis for $\operatorname{Im}\partial_2$. Thus $H_1^\Delta(X)\approx\mathbb{Z}_2$.

Example 2.5. We can obtain a Δ-complex structure on S^n by taking two copies of Δ^n and identifying their boundaries via the identity map. Labeling these two n-simplices U and L, then it is obvious that $\operatorname{Ker}\partial_n$ is infinite cyclic generated by $U-L$. Thus $H_n^\Delta(S^n)\approx\mathbb{Z}$ for this Δ-complex structure on S^n. Computing the other homology groups would be more difficult.

Many similar examples could be worked out without much trouble, such as the other closed orientable and nonorientable surfaces. However, the calculations do tend to increase in complexity before long, particularly for higher-dimensional complexes.

Some obvious general questions arise: Are the groups $H_n^\Delta(X)$ independent of the choice of Δ-complex structure on X? In other words, if two Δ-complexes are homeomorphic, do they have isomorphic homology groups? More generally, do they have isomorphic homology groups if they are merely homotopy equivalent? To answer such questions and to develop a general theory it is best to leave the rather rigid simplicial realm and introduce the singular homology groups. These have the added advantage that they are defined for all spaces, not just Δ-complexes. At the end of this section, after some theory has been developed, we will show that simplicial and singular homology groups coincide for Δ-complexes.

Traditionally, simplicial homology is defined for **simplicial complexes**, which are the Δ-complexes whose simplices are uniquely determined by their vertices. This amounts to saying that each n-simplex has $n+1$ distinct vertices, and that no other n-simplex has this same set of vertices. Thus a simplicial complex can be described combinatorially as a set X_0 of vertices together with sets X_n of n-simplices, which are $(n+1)$-element subsets of X_0. The only requirement is that each $(k+1)$-element subset of the vertices of an n-simplex in X_n is a k-simplex, in X_k. From this combinatorial data a Δ-complex X can be constructed, once we choose a partial ordering of the vertices X_0 that restricts to a linear ordering on the vertices of each simplex in X_n. For example, we could just choose a linear ordering of all the vertices. This might perhaps involve invoking the Axiom of Choice for large vertex sets.

An exercise at the end of this section is to show that every Δ-complex can be subdivided to be a simplicial complex. In particular, every Δ-complex is then homeomorphic to a simplicial complex.

Compared with simplicial complexes, Δ-complexes have the advantage of simpler computations since fewer simplices are required. For example, to put a simplicial complex structure on the torus one needs at least 14 triangles, 21 edges, and 7 vertices, and for $\mathbb{R}P^2$ one needs at least 10 triangles, 15 edges, and 6 vertices. This would slow down calculations considerably!