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.

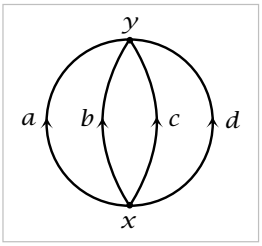

Having abelianized, let us switch to additive notation, so cycles become linear combinations of edges with integer coefficients, such as $a - b + c - d$. Let us call these linear combinations *chains* of edges. Some chains can be decomposed into cycles in several different ways, for example $(a - c) + (b - d) = (a - d) + (b - c)$, 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.

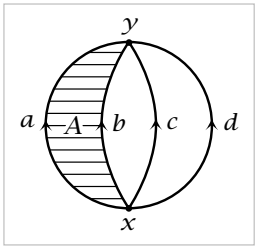

Algebraically, we can define now a pair of homomorphisms $C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0$ where C_2 is the infinite cyclic group generated by A and $\partial_2(A) = a - b$. The map ∂_1 is the boundary homomorphism in the previous example. The quotient group we are interested 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 $c - d$.

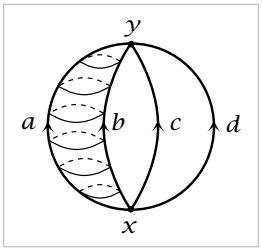

Let us continue one more step and construct a complex X_4 from X_3 by attaching a 3-cell C along the 2-sphere formed by A and B. This creates a chain group C_3

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 $H_2(X_4) = \operatorname{Ker}\partial_2 / \operatorname{Im}\partial_3$ has become trivial. Also $H_3(X_4) = \operatorname{Ker}\partial_3 = 0$. 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 $H_n(X) = \operatorname{Ker}\partial_n / \operatorname{Im}\partial_{n+1}$. 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 $H_n(X) = \operatorname{Ker}\partial_n / \operatorname{Im}\partial_{n+1}$, 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.