Chapter 0

Some Underlying Geometric Notions

The aim of this short preliminary chapter is to introduce a few of the most common geometric concepts and constructions in algebraic topology. The exposition is somewhat informal, with no theorems or proofs until the last couple pages, and it should be read in this informal spirit, skipping bits here and there. In fact, this whole chapter could be skipped now, to be referred back to later for basic definitions.

To avoid overusing the word 'continuous' we adopt the convention that maps between spaces are always assumed to be continuous unless otherwise stated.

Homotopy and Homotopy Type

One of the main ideas of algebraic topology is to consider two spaces to be equivalent if they have 'the same shape' in a sense that is much broader than homeomorphism. To take an everyday example, the letters of the alphabet can be written either as unions of finitely many straight and curved line segments, or in thickened forms that are compact regions in the plane bounded by one or more simple closed curves. In each case the thin letter is a subspace of the thick letter, and we can continuously shrink the thick letter to the thin one. A nice way to do this is to decompose a thick letter, call it $\mathbf{X}$, into line segments connecting each point on the outer boundary of $\mathbf{X}$ to a unique point of the thin subletter X, as indicated in the figure. Then we can shrink $\mathbf{X}$ to X by sliding each point of $\mathbf{X} - X$ into X along the line segment that contains it. Points that are already in X do not move.

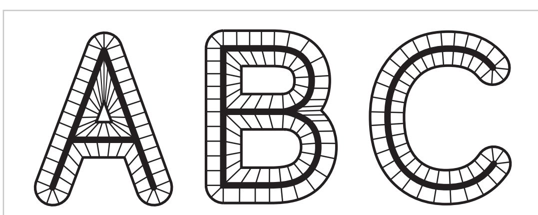

We can think of this shrinking process as taking place during a time interval $0 \le t \le 1$, and then it defines a family of functions $f_t : \mathbf{X} \to \mathbf{X}$ parametrized by $t \in I = [0,1]$, where $f_t(x)$ is the point to which a given point $x \in \mathbf{X}$ has moved at time t. Naturally we would like $f_t(x)$ to depend continuously on both t and x, and this will

Cambridge University Press
978-0-521-79540-1 - Algebraic Topology
Allen Hatcher
Excerpt
More information

be true if we have each $x \in$ **X** − X move along its line segment at constant speed so as to reach its image point in X at time $t = 1$, while points $x \in$ X are stationary, as remarked earlier.

Examples of this sort lead to the following general definition. A **deformation retraction** of a space X onto a subspace A is a family of maps $f_t : X \to X$, $t \in I$, such that $f_0 = \mathbb{1}$ (the identity map), $f_1(X) = A$, and $f_t \,|\, A = \mathbb{1}$ for all t. The family f_t should be continuous in the sense that the associated map $X \times I \to X$, $(x,t) \mapsto f_t(x)$, is continuous.

It is easy to produce many more examples similar to the letter examples, with the deformation retraction f_t obtained by sliding along line segments. The figure on the left below shows such a deformation retraction of a Möbius band onto its core circle.

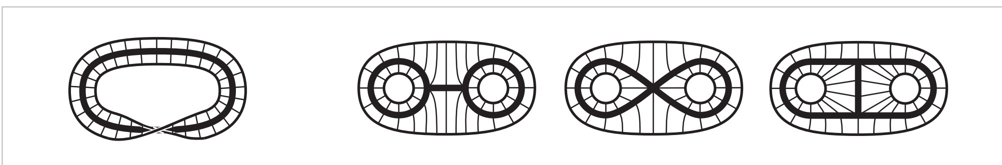

The three figures on the right show deformation retractions in which a disk with two smaller open subdisks removed shrinks to three different subspaces.

In all these examples the structure that gives rise to the deformation retraction can be described by means of the following definition. For a map $f : X \to Y$, the **mapping cylinder** M_f is the quotient space of the disjoint union $(X \times I) \amalg Y$ obtained by identifying each $(x,1) \in X \times I$ with $f(x) \in Y$. In the letter examples, the space X is the outer boundary of the thick letter, Y is the thin letter, and $f : X \to Y$ sends the outer endpoint of each line segment to its inner endpoint. A similar description applies to the other examples. Then it is a general fact that a mapping cylinder M_f deformation retracts to the subspace Y by sliding each point (x,t) along the segment $\{x\} \times I \subset M_f$ to the endpoint $f(x) \in Y$. Continuity of this deformation retraction is evident in the specific examples above, and for a general $f : X \to Y$ it can be verified using Proposition A.17 in the Appendix concerning the interplay between quotient spaces and product spaces.

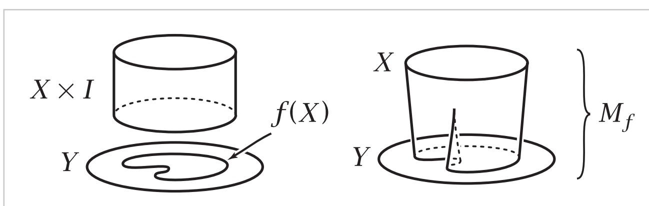

Not all deformation retractions arise in this simple way from mapping cylinders. For example, the thick **X** deformation retracts to the thin X, which in turn deformation retracts to the point of intersection of its two crossbars. The net result is a deformation retraction of **X** onto a point, during which certain pairs of points follow paths that merge before reaching their final destination. Later in this section we will describe a considerably more complicated example, the so-called 'house with two rooms.'

A deformation retraction $f_t : X \to X$ is a special case of the general notion of a **homotopy**, which is simply any family of maps $f_t : X \to Y$, $t \in I$, such that the associated map $F : X \times I \to Y$ given by $F(x,t) = f_t(x)$ is continuous. One says that two maps $f_0, f_1 : X \to Y$ are **homotopic** if there exists a homotopy f_t connecting them, and one writes $f_0 \simeq f_1$.

In these terms, a deformation retraction of X onto a subspace A is a homotopy from the identity map of X to a **retraction** of X onto A, a map $r : X \to X$ such that $r(X) = A$ and $r \mid A = 1\!\!1$. One could equally well regard a retraction as a map $X \to A$ restricting to the identity on the subspace $A \subset X$. From a more formal viewpoint a retraction is a map $r : X \to X$ with $r^2 = r$, since this equation says exactly that r is the identity on its image. Retractions are the topological analogs of projection operators in other parts of mathematics.

Not all retractions come from deformation retractions. For example, a space X always retracts onto any point $x_0 \in X$ via the constant map sending all of X to x_0, but a space that deformation retracts onto a point must be path-connected since a deformation retraction of X to x_0 gives a path joining each $x \in X$ to x_0. It is less trivial to show that there are path-connected spaces that do not deformation retract onto a point. One would expect this to be the case for the letters 'with holes,' A, B, D, O, P, Q, R. In Chapter 1 we will develop techniques to prove this.

A homotopy $f_t : X \to X$ that gives a deformation retraction of X onto a subspace A has the property that $f_t \mid A = 1\!\!1$ for all t. In general, a homotopy $f_t : X \to Y$ whose restriction to a subspace $A \subset X$ is independent of t is called a **homotopy relative to** A, or more concisely, a homotopy rel A. Thus, a deformation retraction of X onto A is a homotopy rel A from the identity map of X to a retraction of X onto A.

If a space X deformation retracts onto a subspace A via $f_t : X \to X$, then if $r : X \to A$ denotes the resulting retraction and $i : A \to X$ the inclusion, we have $ri = 1\!\!1$ and $ir \simeq 1\!\!1$, the latter homotopy being given by f_t. Generalizing this situation, a map $f : X \to Y$ is called a **homotopy equivalence** if there is a map $g : Y \to X$ such that $fg \simeq 1\!\!1$ and $gf \simeq 1\!\!1$. The spaces X and Y are said to be **homotopy equivalent** or to have the same **homotopy type**. The notation is $X \simeq Y$. It is an easy exercise to check that this is an equivalence relation, in contrast with the nonsymmetric notion of deformation retraction. For example, the three graphs  are all homotopy equivalent since they are deformation retracts of the same space, as we saw earlier, but none of the three is a deformation retract of any other.

It is true in general that two spaces X and Y are homotopy equivalent if and only if there exists a third space Z containing both X and Y as deformation retracts. For the less trivial implication one can in fact take Z to be the mapping cylinder M_f of any homotopy equivalence $f : X \to Y$. We observed previously that M_f deformation retracts to Y, so what needs to be proved is that M_f also deformation retracts to its other end X if f is a homotopy equivalence. This is shown in Corollary 0.21.

A space having the homotopy type of a point is called **contractible**. This amounts to requiring that the identity map of the space be **nullhomotopic**, that is, homotopic to a constant map. In general, this is slightly weaker than saying the space deformation retracts to a point; see the exercises at the end of the chapter for an example distinguishing these two notions.

Let us describe now an example of a 2-dimensional subspace of $\mathbb{R}^3$, known as the *house with two rooms*, which is contractible but not in any obvious way. To build this

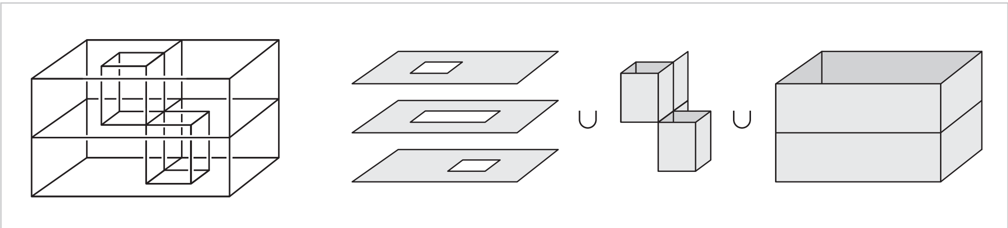

space, start with a box divided into two chambers by a horizontal rectangle, where by a 'rectangle' we mean not just the four edges of a rectangle but also its interior. Access to the two chambers from outside the box is provided by two vertical tunnels. The upper tunnel is made by punching out a square from the top of the box and another square directly below it from the middle horizontal rectangle, then inserting four vertical rectangles, the walls of the tunnel. This tunnel allows entry to the lower chamber from outside the box. The lower tunnel is formed in similar fashion, providing entry to the upper chamber. Finally, two vertical rectangles are inserted to form 'support walls' for the two tunnels. The resulting space X thus consists of three horizontal pieces homeomorphic to annuli plus all the vertical rectangles that form the walls of the two chambers.

To see that X is contractible, consider a closed ε-neighborhood $N(X)$ of X. This clearly deformation retracts onto X if ε is sufficiently small. In fact, $N(X)$ is the mapping cylinder of a map from the boundary surface of $N(X)$ to X. Less obvious is the fact that $N(X)$ is homeomorphic to D^3, the unit ball in $\mathbb{R}^3$. To see this, imagine forming $N(X)$ from a ball of clay by pushing a finger into the ball to create the upper tunnel, then gradually hollowing out the lower chamber, and similarly pushing a finger in to create the lower tunnel and hollowing out the upper chamber. Mathematically, this process gives a family of embeddings $h_t : D^3 \to \mathbb{R}^3$ starting with the usual inclusion $D^3 \hookrightarrow \mathbb{R}^3$ and ending with a homeomorphism onto $N(X)$.

Thus we have $X \simeq N(X) = D^3 \simeq \mathit{point}$, so X is contractible since homotopy equivalence is an equivalence relation. In fact, X deformation retracts to a point. For if f_t is a deformation retraction of the ball $N(X)$ to a point $x_0 \in X$ and if $r : N(X) \to X$ is a retraction, for example the end result of a deformation retraction of $N(X)$ to X, then the restriction of the composition rf_t to X is a deformation retraction of X to x_0. However, it is quite a challenging exercise to see exactly what this deformation retraction looks like.

Cell Complexes

A familiar way of constructing the torus $S^1 \times S^1$ is by identifying opposite sides of a square. More generally, an orientable surface M_g of genus g can be constructed from a polygon with $4g$ sides by identifying pairs of edges, as shown in the figure in the first three cases $g = 1, 2, 3$. The $4g$ edges of the polygon become a union of $2g$ circles in the surface, all intersecting in a single point. The interior of the polygon can be thought of as an open disk, or a **2-cell**, attached to the union of the $2g$ circles. One can also regard the union of the circles as being obtained from their common point of intersection, by attaching $2g$ open arcs, or **1-cells**. Thus the surface can be built up in stages: Start with a point, attach 1-cells to this point, then attach a 2-cell.

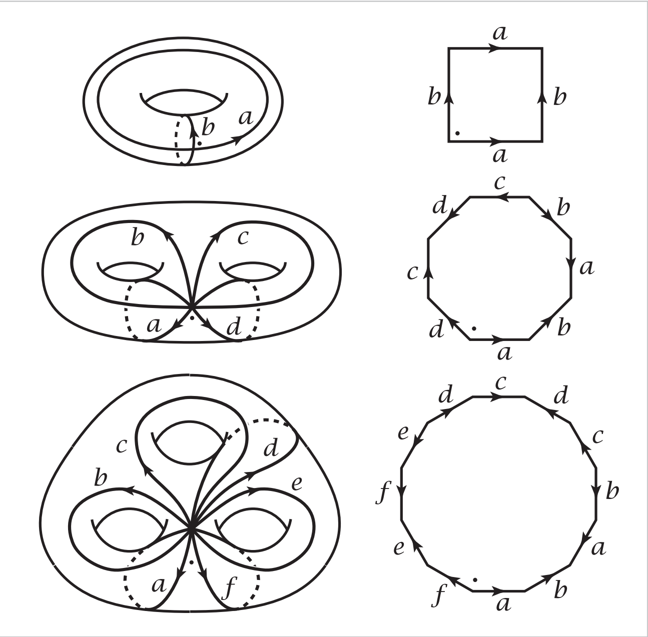

A natural generalization of this is to construct a space by the following procedure:

(1) Start with a discrete set X^0, whose points are regarded as 0-cells.

(2) Inductively, form the ***n*-skeleton** X^n from X^{n-1} by attaching n-cells e^n_α via maps $\varphi_\alpha : S^{n-1} \to X^{n-1}$. This means that X^n is the quotient space of the disjoint union $X^{n-1} \coprod_\alpha D^n_\alpha$ of X^{n-1} with a collection of n-disks D^n_α under the identifications $x \sim \varphi_\alpha(x)$ for $x \in \partial D^n_\alpha$. Thus as a set, $X^n = X^{n-1} \coprod_\alpha e^n_\alpha$ where each e^n_α is an open n-disk.

(3) One can either stop this inductive process at a finite stage, setting $X = X^n$ for some $n < \infty$, or one can continue indefinitely, setting $X = \bigcup_n X^n$. In the latter case X is given the weak topology: A set $A \subset X$ is open (or closed) iff $A \cap X^n$ is open (or closed) in X^n for each n.

A space X constructed in this way is called a **cell complex** or **CW complex**. The explanation of the letters 'CW' is given in the Appendix, where a number of basic topological properties of cell complexes are proved. The reader who wonders about various point-set topological questions lurking in the background of the following discussion should consult the Appendix for details.

If $X = X^n$ for some n, then X is said to be finite-dimensional, and the smallest such n is the **dimension** of X, the maximum dimension of cells of X.

Example 0.1. A 1-dimensional cell complex $X = X^1$ is what is called a **graph** in algebraic topology. It consists of vertices (the 0-cells) to which edges (the 1-cells) are attached. The two ends of an edge can be attached to the same vertex.

Example 0.2. The house with two rooms, pictured earlier, has a visually obvious 2-dimensional cell complex structure. The 0-cells are the vertices where three or more of the depicted edges meet, and the 1-cells are the interiors of the edges connecting these vertices. This gives the 1-skeleton X^1, and the 2-cells are the components of the remainder of the space, $X - X^1$. If one counts up, one finds there are 29 0-cells, 51 1-cells, and 23 2-cells, with the alternating sum $29 - 51 + 23$ equal to 1. This is the **Euler characteristic**, which for a cell complex with finitely many cells is defined to be the number of even-dimensional cells minus the number of odd-dimensional cells. As we shall show in Theorem 2.44, the Euler characteristic of a cell complex depends only on its homotopy type, so the fact that the house with two rooms has the homotopy type of a point implies that its Euler characteristic must be 1, no matter how it is represented as a cell complex.

Example 0.3. The sphere S^n has the structure of a cell complex with just two cells, e^0 and e^n, the n-cell being attached by the constant map $S^{n-1} \to e^0$. This is equivalent to regarding S^n as the quotient space $D^n/\partial D^n$.

Example 0.4. **Real projective n-space** $\mathbb{R}P^n$ is defined to be the space of all lines through the origin in $\mathbb{R}^{n+1}$. Each such line is determined by a nonzero vector in $\mathbb{R}^{n+1}$, unique up to scalar multiplication, and $\mathbb{R}P^n$ is topologized as the quotient space of $\mathbb{R}^{n+1} - \{0\}$ under the equivalence relation $v \sim \lambda v$ for scalars $\lambda \neq 0$. We can restrict to vectors of length 1, so $\mathbb{R}P^n$ is also the quotient space $S^n/(v \sim -v)$, the sphere with antipodal points identified. This is equivalent to saying that $\mathbb{R}P^n$ is the quotient space of a hemisphere D^n with antipodal points of ∂D^n identified. Since ∂D^n with antipodal points identified is just $\mathbb{R}P^{n-1}$, we see that $\mathbb{R}P^n$ is obtained from $\mathbb{R}P^{n-1}$ by attaching an n-cell, with the quotient projection $S^{n-1} \to \mathbb{R}P^{n-1}$ as the attaching map. It follows by induction on n that $\mathbb{R}P^n$ has a cell complex structure $e^0 \cup e^1 \cup \cdots \cup e^n$ with one cell e^i in each dimension $i \leq n$.

Example 0.5. Since $\mathbb{R}P^n$ is obtained from $\mathbb{R}P^{n-1}$ by attaching an n-cell, the infinite union $\mathbb{R}P^\infty = \bigcup_n \mathbb{R}P^n$ becomes a cell complex with one cell in each dimension. We can view $\mathbb{R}P^\infty$ as the space of lines through the origin in $\mathbb{R}^\infty = \bigcup_n \mathbb{R}^n$.

Example 0.6. **Complex projective n-space** $\mathbb{C}P^n$ is the space of complex lines through the origin in $\mathbb{C}^{n+1}$, that is, 1-dimensional vector subspaces of $\mathbb{C}^{n+1}$. As in the case of $\mathbb{R}P^n$, each line is determined by a nonzero vector in $\mathbb{C}^{n+1}$, unique up to scalar multiplication, and $\mathbb{C}P^n$ is topologized as the quotient space of $\mathbb{C}^{n+1} - \{0\}$ under the equivalence relation $v \sim \lambda v$ for $\lambda \neq 0$. Equivalently, this is the quotient of the unit sphere $S^{2n+1} \subset \mathbb{C}^{n+1}$ with $v \sim \lambda v$ for $|\lambda| = 1$. It is also possible to obtain $\mathbb{C}\mathrm{P}^n$ as a quotient space of the disk D^{2n} under the identifications $v \sim \lambda v$ for $v \in \partial D^{2n}$, in the following way. The vectors in $S^{2n+1} \subset \mathbb{C}^{n+1}$ with last coordinate real and nonnegative are precisely the vectors of the form $(w, \sqrt{1-|w|^2}\,) \in \mathbb{C}^n \times \mathbb{C}$ with $|w| \leq 1$. Such vectors form the graph of the function $w \mapsto \sqrt{1-|w|^2}$. This is a disk D^{2n}_+ bounded by the sphere $S^{2n-1} \subset S^{2n+1}$ consisting of vectors $(w, 0) \in \mathbb{C}^n \times \mathbb{C}$ with $|w| = 1$. Each vector in S^{2n+1} is equivalent under the identifications $v \sim \lambda v$ to a vector in D^{2n}_+, and the latter vector is unique if its last coordinate is nonzero. If the last coordinate is zero, we have just the identifications $v \sim \lambda v$ for $v \in S^{2n-1}$.

From this description of $\mathbb{C}\mathrm{P}^n$ as the quotient of D^{2n}_+ under the identifications $v \sim \lambda v$ for $v \in S^{2n-1}$ it follows that $\mathbb{C}\mathrm{P}^n$ is obtained from $\mathbb{C}\mathrm{P}^{n-1}$ by attaching a cell e^{2n} via the quotient map $S^{2n-1} \to \mathbb{C}\mathrm{P}^{n-1}$. So by induction on n we obtain a cell structure $\mathbb{C}\mathrm{P}^n = e^0 \cup e^2 \cup \cdots \cup e^{2n}$ with cells only in even dimensions. Similarly, $\mathbb{C}\mathrm{P}^\infty$ has a cell structure with one cell in each even dimension.

After these examples we return now to general theory. Each cell e^n_α in a cell complex X has a **characteristic map** $\Phi_\alpha : D^n_\alpha \to X$ which extends the attaching map φ_α and is a homeomorphism from the interior of D^n_α onto e^n_α. Namely, we can take Φ_α to be the composition $D^n_\alpha \hookrightarrow X^{n-1} \coprod_\alpha D^n_\alpha \to X^n \hookrightarrow X$ where the middle map is the quotient map defining X^n. For example, in the canonical cell structure on S^n described in Example 0.3, a characteristic map for the n-cell is the quotient map $D^n \to S^n$ collapsing ∂D^n to a point. For $\mathbb{R}\mathrm{P}^n$ a characteristic map for the cell e^i is the quotient map $D^i \to \mathbb{R}\mathrm{P}^i \subset \mathbb{R}\mathrm{P}^n$ identifying antipodal points of ∂D^i, and similarly for $\mathbb{C}\mathrm{P}^n$.

A **subcomplex** of a cell complex X is a closed subspace $A \subset X$ that is a union of cells of X. Since A is closed, the characteristic map of each cell in A has image contained in A, and in particular the image of the attaching map of each cell in A is contained in A, so A is a cell complex in its own right. A pair (X, A) consisting of a cell complex X and a subcomplex A will be called a **CW pair**.

For example, each skeleton X^n of a cell complex X is a subcomplex. Particular cases of this are the subcomplexes $\mathbb{R}\mathrm{P}^k \subset \mathbb{R}\mathrm{P}^n$ and $\mathbb{C}\mathrm{P}^k \subset \mathbb{C}\mathrm{P}^n$ for $k \leq n$. These are in fact the only subcomplexes of $\mathbb{R}\mathrm{P}^n$ and $\mathbb{C}\mathrm{P}^n$.

There are natural inclusions $S^0 \subset S^1 \subset \cdots \subset S^n$, but these subspheres are not subcomplexes of S^n in its usual cell structure with just two cells. However, we can give S^n a different cell structure in which each of the subspheres S^k is a subcomplex, by regarding each S^k as being obtained inductively from the equatorial S^{k-1} by attaching two k-cells, the components of $S^k - S^{k-1}$. The infinite-dimensional sphere $S^\infty = \bigcup_n S^n$ then becomes a cell complex as well. Note that the two-to-one quotient map $S^\infty \to \mathbb{R}\mathrm{P}^\infty$ that identifies antipodal points of S^∞ identifies the two n-cells of S^∞ to the single n-cell of $\mathbb{R}\mathrm{P}^\infty$.

In the examples of cell complexes given so far, the closure of each cell is a subcomplex, and more generally the closure of any collection of cells is a subcomplex. Most naturally arising cell structures have this property, but it need not hold in general. For example, if we start with S^1 with its minimal cell structure and attach to this a 2-cell by a map $S^1 \to S^1$ whose image is a nontrivial subarc of S^1, then the closure of the 2-cell is not a subcomplex since it contains only a part of the 1-cell.

Operations on Spaces

Cell complexes have a very nice mixture of rigidity and flexibility, with enough rigidity to allow many arguments to proceed in a combinatorial cell-by-cell fashion and enough flexibility to allow many natural constructions to be performed on them. Here are some of those constructions.

Products. If X and Y are cell complexes, then $X \times Y$ has the structure of a cell complex with cells the products $e^m_\alpha \times e^n_\beta$ where e^m_α ranges over the cells of X and e^n_β ranges over the cells of Y. For example, the cell structure on the torus $S^1 \times S^1$ described at the beginning of this section is obtained in this way from the standard cell structure on S^1. For completely general CW complexes X and Y there is one small complication: The topology on $X \times Y$ as a cell complex is sometimes finer than the product topology, with more open sets than the product topology has, though the two topologies coincide if either X or Y has only finitely many cells, or if both X and Y have countably many cells. This is explained in the Appendix. In practice this subtle issue of point-set topology rarely causes problems, however.

Quotients. If (X, A) is a CW pair consisting of a cell complex X and a subcomplex A, then the quotient space X/A inherits a natural cell complex structure from X. The cells of X/A are the cells of $X - A$ plus one new 0-cell, the image of A in X/A. For a cell e^n_α of $X - A$ attached by $\varphi_\alpha : S^{n-1} \to X^{n-1}$, the attaching map for the corresponding cell in X/A is the composition $S^{n-1} \to X^{n-1} \to X^{n-1}/A^{n-1}$.

For example, if we give S^{n-1} any cell structure and build D^n from S^{n-1} by attaching an n-cell, then the quotient D^n/S^{n-1} is S^n with its usual cell structure. As another example, take X to be a closed orientable surface with the cell structure described at the beginning of this section, with a single 2-cell, and let A be the complement of this 2-cell, the 1-skeleton of X. Then X/A has a cell structure consisting of a 0-cell with a 2-cell attached, and there is only one way to attach a cell to a 0-cell, by the constant map, so X/A is S^2.

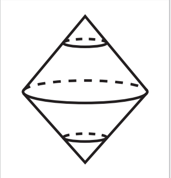

Suspension. For a space X, the **suspension** SX is the quotient of $X \times I$ obtained by collapsing $X \times \{0\}$ to one point and $X \times \{1\}$ to another point. The motivating example is $X = S^n$, when $SX = S^{n+1}$ with the two 'suspension points' at the north and south poles of S^{n+1}, the points $(0, \cdots, 0, \pm 1)$. One can regard SX as a double cone on X, the union of two copies of the **cone** $CX = (X\times I)/(X\times\{0\})$. If X is a CW complex, so are SX and CX as quotients of $X\times I$ with its product cell structure, I being given the standard cell structure of two 0-cells joined by a 1-cell.

Suspension becomes increasingly important the farther one goes into algebraic topology, though why this should be so is certainly not evident in advance. One especially useful property of suspension is that not only spaces but also maps can be suspended. Namely, a map $f:X\to Y$ suspends to $Sf:SX\to SY$, the quotient map of $f\times 1\!\!1:X\times I\to Y\times I$.

Join. The cone CX is the union of all line segments joining points of X to an external vertex, and similarly the suspension SX is the union of all line segments joining points of X to two external vertices. More generally, given X and a second space Y, one can define the space of all line segments joining points in X to points in Y. This is the **join** $X*Y$, the quotient space of $X\times Y\times I$ under the identifications $(x,y_1,0)\sim(x,y_2,0)$ and $(x_1,y,1)\sim(x_2,y,1)$. Thus we are collapsing the subspace $X\times Y\times\{0\}$ to X and $X\times Y\times\{1\}$ to Y. For example, if X and Y are both closed intervals, then we are collapsing two opposite faces of a cube onto line segments so that the cube becomes a tetrahedron. In the general case, $X*Y$ contains copies of X and Y at its two ends, and every other point (x,y,t) in $X*Y$ is on a unique line segment joining the point $x\in X\subset X*Y$ to the point $y\in Y\subset X*Y$, the segment obtained by fixing x and y and letting the coordinate t in (x,y,t) vary.

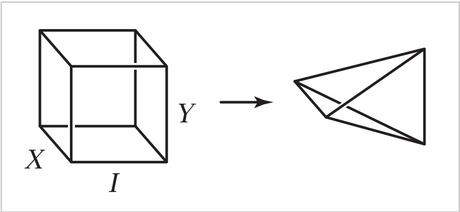

A nice way to write points of $X*Y$ is as formal linear combinations t_1x+t_2y with $0\le t_i\le 1$ and $t_1+t_2=1$, subject to the rules $0x+1y=y$ and $1x+0y=x$ that correspond exactly to the identifications defining $X*Y$. In much the same way, an iterated join $X_1*\cdots*X_n$ can be constructed as the space of formal linear combinations $t_1x_1+\cdots+t_nx_n$ with $0\le t_i\le 1$ and $t_1+\cdots+t_n=1$, with the convention that terms $0x_i$ can be omitted. A very special case that plays a central role in algebraic topology is when each X_i is just a point. For example, the join of two points is a line segment, the join of three points is a triangle, and the join of four points is a tetrahedron. In general, the join of n points is a convex polyhedron of dimension $n-1$ called a **simplex**. Concretely, if the n points are the n standard basis vectors for $\mathbb{R}^n$, then their join is the $(n-1)$-dimensional simplex

$$\Delta^{n-1}=\{\,(t_1,\cdots,t_n)\in\mathbb{R}^n \mid t_1+\cdots+t_n=1 \text{ and } t_i\ge 0\,\}$$

Another interesting example is when each X_i is S^0, two points. If we take the two points of X_i to be the two unit vectors along the i^{th} coordinate axis in $\mathbb{R}^n$, then the join $X_1*\cdots*X_n$ is the union of 2^n copies of the simplex Δ^{n-1}, and radial projection from the origin gives a homeomorphism between $X_1*\cdots*X_n$ and S^{n-1}.

If X and Y are CW complexes, then there is a natural CW structure on $X * Y$ having the subspaces X and Y as subcomplexes, with the remaining cells being the product cells of $X \times Y \times (0,1)$. As usual with products, the CW topology on $X * Y$ may be finer than the quotient of the product topology on $X \times Y \times I$.

Wedge Sum. This is a rather trivial but still quite useful operation. Given spaces X and Y with chosen points $x_0 \in X$ and $y_0 \in Y$, then the **wedge sum** $X \vee Y$ is the quotient of the disjoint union $X \amalg Y$ obtained by identifying x_0 and y_0 to a single point. For example, $S^1 \vee S^1$ is homeomorphic to the figure '8,' two circles touching at a point. More generally one could form the wedge sum $\bigvee_\alpha X_\alpha$ of an arbitrary collection of spaces X_α by starting with the disjoint union $\coprod_\alpha X_\alpha$ and identifying points $x_\alpha \in X_\alpha$ to a single point. In case the spaces X_α are cell complexes and the points x_α are 0-cells, then $\bigvee_\alpha X_\alpha$ is a cell complex since it is obtained from the cell complex $\coprod_\alpha X_\alpha$ by collapsing a subcomplex to a point.

For any cell complex X, the quotient X^n/X^{n-1} is a wedge sum of n-spheres $\bigvee_\alpha S^n_\alpha$, with one sphere for each n-cell of X.

Smash Product. Like suspension, this is another construction whose importance becomes evident only later. Inside a product space $X \times Y$ there are copies of X and Y, namely $X \times \{y_0\}$ and $\{x_0\} \times Y$ for points $x_0 \in X$ and $y_0 \in Y$. These two copies of X and Y in $X \times Y$ intersect only at the point (x_0, y_0), so their union can be identified with the wedge sum $X \vee Y$. The **smash product** $X \wedge Y$ is then defined to be the quotient $X \times Y / X \vee Y$. One can think of $X \wedge Y$ as a reduced version of $X \times Y$ obtained by collapsing away the parts that are not genuinely a product, the separate factors X and Y.

The smash product $X \wedge Y$ is a cell complex if X and Y are cell complexes with x_0 and y_0 0-cells, assuming that we give $X \times Y$ the cell-complex topology rather than the product topology in cases when these two topologies differ. For example, $S^m \wedge S^n$ has a cell structure with just two cells, of dimensions 0 and $m+n$, hence $S^m \wedge S^n = S^{m+n}$. In particular, when $m = n = 1$ we see that collapsing longitude and meridian circles of a torus to a point produces a 2-sphere.

Two Criteria for Homotopy Equivalence

Earlier in this chapter the main tool we used for constructing homotopy equivalences was the fact that a mapping cylinder deformation retracts onto its 'target' end. By repeated application of this fact one can often produce homotopy equivalences between rather different-looking spaces. However, this process can be a bit cumbersome in practice, so it is useful to have other techniques available as well. We will describe two commonly used methods here. The first involves collapsing certain subspaces to points, and the second involves varying the way in which the parts of a space are put together.