Appendix

Topology of Cell Complexes

Here we collect a number of basic topological facts about CW complexes for convenient reference. A few related facts about manifolds are also proved.

Let us first recall from Chapter 0 that a CW complex is a space X constructed in the following way:

(1) Start with a discrete set X^0, the 0-cells of X.

(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 $X^{n-1} \coprod_\alpha D^n_\alpha$ under the identifications $x \sim \varphi_\alpha(x)$ for $x \in \partial D^n_\alpha$. The cell e^n_α is the homeomorphic image of $D^n_\alpha - \partial D^n_\alpha$ under the quotient map.

(3) $X = \bigcup_n X^n$ with 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.

Note that condition (3) is superfluous when X is finite-dimensional, so that $X = X^n$ for some n. For if A is open in $X = X^n$, the definition of the quotient topology on X^n implies that $A \cap X^{n-1}$ is open in X^{n-1}, and then by the same reasoning $A \cap X^{n-2}$ is open in X^{n-2}, and similarly for all the skeleta X^{n-i}.

Each cell e^n_α has its **characteristic map** Φ_α, which is by definition the composition $D^n_\alpha \hookrightarrow X^{n-1} \coprod_\alpha D^n_\alpha \to X^n \hookrightarrow X$. This is continuous since it is a composition of continuous maps, the inclusion $X^n \hookrightarrow X$ being continuous by (3). The restriction of Φ_α to the interior of D^n_α is a homeomorphism onto e^n_α.

An alternative way to describe the topology on X is to say that a set $A \subset X$ is open (or closed) iff $\Phi_\alpha^{-1}(A)$ is open (or closed) in D^n_α for each characteristic map Φ_α. In one direction this follows from continuity of the Φ_α's, and in the other direction, suppose $\Phi_\alpha^{-1}(A)$ is open in D^n_α for each Φ_α, and suppose by induction on n that $A \cap X^{n-1}$ is open in X^{n-1}. Then since $\Phi_\alpha^{-1}(A)$ is open in D^n_α for all α, $A \cap X^n$ is open in X^n by the definition of the quotient topology on X^n. Hence by (3), A is open in X.

A consequence of this characterization of the topology on X is that X is a quotient space of $\coprod_{n,\alpha} D^n_\alpha$.

A **subcomplex** of a CW complex X is a subspace $A \subset X$ which is a union of cells of X, such that the closure of each cell in A is contained in A. Thus for each cell in A, the image of its attaching map is contained in A, so A is itself a CW complex. Its CW complex topology is the same as the topology induced from X, as one sees by noting inductively that the two topologies agree on $A^n = A \cap X^n$. It is easy to see by induction over skeleta that a subcomplex is a closed subspace. Conversely, a subcomplex could be defined as a closed subspace which is a union of cells.

A finite CW complex, that is, one with only finitely many cells, is compact since attaching a single cell preserves compactness. A sort of converse to this is:

Proposition A.1. *A compact subspace of a CW complex is contained in a finite subcomplex.*

Proof: First we show that a compact set C in a CW complex X can meet only finitely many cells of X. Suppose on the contrary that there is an infinite sequence of points $x_i \in C$ all lying in distinct cells. Then the set $S = \{x_1, x_2, \cdots\}$ is closed in X. Namely, assuming $S \cap X^{n-1}$ is closed in X^{n-1} by induction on n, then for each cell e^n_α of X, $\varphi_\alpha^{-1}(S)$ is closed in ∂D^n_α, and $\Phi_\alpha^{-1}(S)$ consists of at most one more point in D^n_α, so $\Phi_\alpha^{-1}(S)$ is closed in D^n_α. Therefore $S \cap X^n$ is closed in X^n for each n, hence S is closed in X. The same argument shows that any subset of S is closed, so S has the discrete topology. But it is compact, being a closed subset of the compact set C. Therefore S must be finite, a contradiction.

Since C is contained in a finite union of cells, it suffices to show that a finite union of cells is contained in a finite subcomplex of X. A finite union of finite subcomplexes is again a finite subcomplex, so this reduces to showing that a single cell e^n_α is contained in a finite subcomplex. The image of the attaching map φ_α for e^n_α is compact, hence by induction on dimension this image is contained in a finite subcomplex $A \subset X^{n-1}$. So e^n_α is contained in the finite subcomplex $A \cup e^n_\alpha$. □

Now we can explain the mysterious letters 'CW,' which refer to the following two properties satisfied by CW complexes:

(1) Closure-finiteness: The closure of each cell meets only finitely many other cells. This follows from the preceding proposition since the closure of a cell is compact, being the image of a characteristic map.

(2) Weak topology: A set is closed iff it meets the closure of each cell in a closed set. For if a set meets the closure of each cell in a closed set, it pulls back to a closed set under each characteristic map, hence is closed by an earlier remark.

In J. H. C. Whitehead's original definition of CW complexes these two properties played a more central role. The following proposition contains essentially this definition.

Proposition A.2. *Given a Hausdorff space X and a family of maps $\Phi_\alpha : D^n_\alpha \to X$, then these maps are the characteristic maps of a CW complex structure on X iff:*

(i) *Each Φ_α is injective on $\operatorname{int} D^n_\alpha$, hence Φ_α restricts to a homeomorphism from $\operatorname{int} D^n_\alpha$ onto a cell $e^n_\alpha \subset X$. All these cells are disjoint, and their union is X.*

(ii) *For each cell e^n_α, $\Phi_\alpha(\partial D^n_\alpha)$ is contained in the union of a finite number of cells of dimension less than n.*

(iii) *A subset of X is closed iff it meets the closure of each cell of X in a closed set.*

The 'hence' in (i) follows from the fact that Φ_α maps the compact set D^n_α to a Hausdorff space, so since Φ_α takes compact sets to compact sets, it takes closed sets to closed sets, which means that $\Phi^{-1}_\alpha : e^n_\alpha \to \operatorname{int} D^n_\alpha$ is continuous. By the same compactness argument, condition (iii) can be restated as saying that a set $C \subset X$ is closed iff $\Phi^{-1}_\alpha(C)$ is closed in D^n_α for all α. In particular, (iii) is automatic if there are only finitely many cells since the projection $\coprod_\alpha D^n_\alpha \to X$ is a map from a compact space onto a Hausdorff space, hence is a quotient map.

For an example where all the conditions except the finiteness hypothesis in (ii) are satisfied, take X to be D^2 with its interior as a 2-cell and each point of ∂D^2 as a 0-cell. The identity map of D^2 serves as the Φ_α for the 2-cell. Condition (iii) is satisfied since it is a nontrivial condition only for the 2-cell.

Proof: We have already taken care of the 'only if' implication. For the converse, suppose inductively that X^{n-1}, the union of all cells of dimension less than n, is a CW complex with the appropriate Φ_α's as characteristic maps. The induction can start with $X^{-1} = \varnothing$. Let $f : X^{n-1} \coprod_\alpha D^n_\alpha \to X^n$ be given by the inclusion on X^{n-1} and the maps Φ_α for all the n-cells of X. This is a continuous surjection, and if we can show it is a quotient map, then X^n will be obtained from X^{n-1} by attaching the n-cells e^n_α. Thus if $C \subset X^n$ is such that $f^{-1}(C)$ is closed, we need to show that $C \cap \overline{e}^m_\beta$ is closed for all cells e^m_β of X, the bar denoting closure.

There are three cases. If $m < n$ then $f^{-1}(C)$ closed implies $C \cap X^{n-1}$ closed, hence $C \cap \overline{e}^m_\beta$ is closed since $\overline{e}^m_\beta \subset X^{n-1}$. If $m = n$ then e^m_β is one of the cells e^n_α, so $f^{-1}(C)$ closed implies $f^{-1}(C) \cap D^n_\alpha$ is closed, hence compact, hence its image $C \cap \overline{e}^n_\alpha$ under f is compact and therefore closed. Finally there is the case $m > n$. Then $C \subset X^n$ implies $C \cap \overline{e}^m_\beta \subset \Phi_\beta(\partial D^m_\beta)$. The latter space is contained in a finite union of $\overline{e}^\ell_\gamma$'s with $\ell < m$. By induction on m, each $C \cap \overline{e}^\ell_\gamma$ is closed. Hence the intersection of C with the union of the finite collection of $\overline{e}^\ell_\gamma$'s is closed. Intersecting this closed set with $\overline{e}^m_\beta$, we conclude that $C \cap \overline{e}^m_\beta$ is closed.

It remains only to check that X has the weak topology with respect to the X^n's, that is, a set in X is closed iff it intersects each X^n in a closed set. The preceding argument with $C = X^n$ shows that X^n is closed, so a closed set intersects each X^n in a closed set. Conversely, if a set C intersects X^n in a closed set, then C intersects each $\overline{e}^n_\alpha$ in a closed set, so C is closed in X by (iii). □

Next we describe a convenient way of constructing open neighborhoods $N_\varepsilon(A)$ of subsets A of a CW complex X, where ε is a function assigning a number $\varepsilon_\alpha > 0$ to each cell e^n_α of X. The construction is inductive over the skeleta X^n, so suppose we have already constructed $N^n_\varepsilon(A)$, a neighborhood of $A \cap X^n$ in X^n, starting the process with $N^0_\varepsilon(A) = A \cap X^0$. Then we define $N^{n+1}_\varepsilon(A)$ by specifying its preimage under the characteristic map $\Phi_\alpha : D^{n+1} \to X$ of each cell e^{n+1}_α, namely, $\Phi^{-1}_\alpha(N^{n+1}_\varepsilon(A))$ is the union of two parts: an open ε_α-neighborhood of $\Phi^{-1}_\alpha(A) - \partial D^{n+1}$ in $D^{n+1} - \partial D^{n+1}$, and a product $(1 - \varepsilon_\alpha, 1] \times \Phi^{-1}_\alpha(N^n_\varepsilon(A))$ with respect to 'spherical' coordinates (r, θ) in D^{n+1}, where $r \in [0,1]$ is the radial coordinate and θ lies in $\partial D^{n+1} = S^n$. Then we define $N_\varepsilon(A) = \bigcup_n N^n_\varepsilon(A)$. This is an open set in X since it pulls back to an open set under each characteristic map.

Proposition A.3. *CW complexes are normal, and in particular, Hausdorff.*

Proof: Points are closed in a CW complex X since they pull back to closed sets under all characteristic maps Φ_α. For disjoint closed sets A and B in X, we show that $N_\varepsilon(A)$ and $N_\varepsilon(B)$ are disjoint for small enough ε_α's. In the inductive process for building these open sets, assume $N^n_\varepsilon(A)$ and $N^n_\varepsilon(B)$ have been chosen to be disjoint. For a characteristic map $\Phi_\alpha : D^{n+1} \to X$, observe that $\Phi^{-1}_\alpha(N^n_\varepsilon(A))$ and $\Phi^{-1}_\alpha(B)$ are a positive distance apart, since otherwise by compactness we would have a sequence in $\Phi^{-1}_\alpha(B)$ converging to a point of $\Phi^{-1}_\alpha(B)$ in ∂D^{n+1} of distance zero from $\Phi^{-1}_\alpha(N^n_\varepsilon(A))$, but this is impossible since $\Phi^{-1}_\alpha(N^n_\varepsilon(B))$ is a neighborhood of $\Phi^{-1}_\alpha(B) \cap \partial D^{n+1}$ in ∂D^{n+1} disjoint from $\Phi^{-1}_\alpha(N^n_\varepsilon(A))$. Similarly, $\Phi^{-1}_\alpha(N^n_\varepsilon(B))$ and $\Phi^{-1}_\alpha(A)$ are a positive distance apart. Also, $\Phi^{-1}_\alpha(A)$ and $\Phi^{-1}_\alpha(B)$ are a positive distance apart. So a small enough ε_α will make $\Phi^{-1}_\alpha(N^{n+1}_\varepsilon(A))$ disjoint from $\Phi^{-1}_\alpha(N^{n+1}_\varepsilon(B))$ in D^{n+1}. □

Proposition A.4. *Each point in a CW complex has arbitrarily small contractible open neighborhoods, so CW complexes are locally contractible.*

Proof: Given a point x in a CW complex X and a neighborhood U of x in X, we can choose the ε_α's small enough so that $N_\varepsilon(x) \subset U$ by requiring that the closure of $N^n_\varepsilon(x)$ be contained in U for each n. It remains to see that $N_\varepsilon(x)$ is contractible. If $x \in X^m - X^{m-1}$ and $n > m$ we can construct a deformation retraction of $N^n_\varepsilon(x)$ onto $N^{n-1}_\varepsilon(x)$ by sliding outward along radial segments in cells e^n_β, the images under the characteristic maps Φ_β of radial segments in D^n. A deformation retraction of $N_\varepsilon(x)$ onto $N^m_\varepsilon(x)$ is then obtained by performing the deformation retraction of $N^n_\varepsilon(x)$ onto $N^{n-1}_\varepsilon(x)$ during the t-interval $[1/2^n, 1/2^{n-1}]$, points of $N^n_\varepsilon(x) - N^{n-1}_\varepsilon(x)$ being stationary outside this t-interval. Finally, $N^m_\varepsilon(x)$ is an open ball about x, and so deformation retracts onto x. □

In particular, CW complexes are locally path-connected. So a CW complex is path-connected iff it is connected.

Proposition A.5. *For a subcomplex A of a CW complex X, the open neighborhood $N_\varepsilon(A)$ deformation retracts onto A if $\varepsilon_\alpha < 1$ for all α.*

Proof: In each cell of $X - A$, $N_\varepsilon(A)$ is a product neighborhood of the boundary of this cell, so a deformation retraction of $N_\varepsilon(A)$ onto A can be constructed just as in the previous proof. □

Note that for subcomplexes A and B of X, we have $N_\varepsilon(A) \cap N_\varepsilon(B) = N_\varepsilon(A \cap B)$. This implies for example that the van Kampen theorem and Mayer-Vietoris sequences hold for decompositions $X = A \cup B$ into subcomplexes A and B as well as into open sets A and B.

A map $f : X \to Y$ with domain a CW complex is continuous iff its restrictions to the closures $\overline{e}^n_\alpha$ of all cells e^n_α are continuous, and it is useful to know that the same is true for homotopies $f_t : X \to Y$. With this objective in mind, let us introduce a little terminology. A topological space X is said to be **generated** by a collection of subspaces X_α if $X = \bigcup_\alpha X_\alpha$ and a set $A \subset X$ is closed iff $A \cap X_\alpha$ is closed in X_α for each α. Equivalently, we could say 'open' instead of 'closed' here, but 'closed' is more convenient for our present purposes. As noted earlier, though not in these words, a CW complex X is generated by the closures $\overline{e}^n_\alpha$ of its cells e^n_α. Since every finite subcomplex of X is a finite union of closures $\overline{e}^n_\alpha$, X is also generated by its finite subcomplexes. It follows that X is also generated by its compact subspaces, or more briefly, X is **compactly generated.**

Proposition A.15 later in the Appendix asserts that if X is a compactly generated Hausdorff space and Z is locally compact, then $X \times Z$, with the product topology, is compactly generated. In particular, $X \times I$ is compactly generated if X is a CW complex. Since every compact set in $X \times I$ is contained in the product of a compact subspace of X with I, hence in the product of a finite subcomplex of X with I, such product subspaces also generate $X \times I$. Since such a product subspace is a finite union of products $\overline{e}^n_\alpha \times I$, it is also true that $X \times I$ is generated by its subspaces $\overline{e}^n_\alpha \times I$. This implies that a homotopy $F : X \times I \to Y$ is continuous iff its restrictions to the subspaces $\overline{e}^n_\alpha \times I$ are continuous, which is the statement we were seeking.

Products of CW Complexes

There are some unexpected point-set-topological subtleties that arise with products of CW complexes. As we shall show, the product of two CW complexes does have a natural CW structure, but its topology is in general finer, with more open sets, than the product topology. However, the distinctions between the two topologies are rather small, and indeed nonexistent in most cases of interest, so there is no real problem for algebraic topology.

Given a space X and a collection of subspaces X_α whose union is X, these subspaces generate a possibly finer topology on X by defining a set $A \subset X$ to be open iff $A \cap X_\alpha$ is open in X_α for all α. The axioms for a topology are easily verified for this definition. In case $\{X_\alpha\}$ is the collection of compact subsets of X, we write X_c for this new compactly generated topology. It is easy to see that X and X_c have the same compact subsets, and the two induced topologies on these compact subsets coincide. If X is compact, or even locally compact, then $X = X_c$, that is, X is compactly generated.

Theorem A.6. *For CW complexes X and Y with characteristic maps Φ_α and Ψ_β, the product maps $\Phi_\alpha \times \Psi_\beta$ are the characteristic maps for a CW complex structure on $(X \times Y)_c$. If either X or Y is compact or more generally locally compact, then $(X \times Y)_c = X \times Y$. Also, $(X \times Y)_c = X \times Y$ if both X and Y have countably many cells.*

Proof: For the first statement it suffices to check that the three conditions in Proposition A.2 are satisfied when we take the space 'X' there to be $(X \times Y)_c$. The first two conditions are obvious. For the third, which says that $(X \times Y)_c$ is generated by the products $\overline{e}^m_\alpha \times \overline{e}^n_\beta$, observe that every compact set in $X \times Y$ is contained in the product of its projections onto X and Y, and these projections are compact and hence contained in finite subcomplexes of X and Y, so the original compact set is contained in a finite union of products $\overline{e}^m_\alpha \times \overline{e}^n_\beta$. Hence the products $\overline{e}^m_\alpha \times \overline{e}^n_\beta$ generate $(X \times Y)_c$.

The second assertion of the theorem is a special case of Proposition A.15, having nothing to do with CW complexes, which says that a product $X \times Y$ is compactly generated if X is compactly generated Hausdorff and Y is locally compact.

For the last statement of the theorem, suppose X and Y each have at most countably many cells. For an open set $W \subset (X \times Y)_c$ and a point $(a,b) \in W$ we need to find a product $U \times V \subset W$ with U an open neighborhood of a in X and V an open neighborhood of b in Y. Choose finite subcomplexes $X_1 \subset X_2 \subset \cdots$ of X with $X = \bigcup_i X_i$, and similarly for Y. We may assume $a \in X_1$ and $b \in Y_1$. Since the two topologies agree on $X_1 \times Y_1$, there is a compact product neighborhood $K_1 \times L_1 \subset W$ of (a,b) in $X_1 \times Y_1$. Assuming inductively that $K_i \times L_i \subset W$ has been constructed in $X_i \times Y_i$, we would like to construct $K_{i+1} \times L_{i+1} \subset W$ as a compact neighborhood of $K_i \times L_i$ in $X_{i+1} \times Y_{i+1}$. To do this, we first choose for each $x \in K_i$ compact neighborhoods K_x of x in X_{i+1} and L_x of L_i in Y_{i+1} such that $K_x \times L_x \subset W$, using the compactness of L_i. By compactness of K_i, a finite number of the K_x's cover K_i. Let K_{i+1} be the union of these K_x's and let L_{i+1} be the intersection of the corresponding L_x's. This defines the desired $K_{i+1} \times L_{i+1}$. Let U_i be the interior of K_i in X_i, so $U_i \subset U_{i+1}$ for each i. The union $U = \bigcup_i U_i$ is then open in X since it intersects each X_i in a union of open sets and the X_i's generate X. In the same way the L_i's yield an open set V in Y. Thus we have a product of open sets $U \times V \subset W$ containing (a,b). ◻

We will describe now an example from [Dowker 1952] where the product topology on $X \times Y$ differs from the CW topology. Both X and Y will be graphs consisting of infinitely many edges emanating from a single vertex, with uncountably many edges for X and countably many for Y.

Let $X = \bigvee_s I_s$ where I_s is a copy of the interval $[0,1]$ and the index s ranges over all infinite sequences $s = (s_1, s_2, \cdots)$ of positive integers. The wedge sum is formed at the 0 endpoint of I_s. Similarly we let $Y = \bigvee_j I_j$ but with j varying just over positive integers. Let p_{sj} be the point $(1/s_j, 1/s_j) \in I_s \times I_j \subset X \times Y$ and let P be the union of all these points p_{sj}. Thus P consists of a single point in each 2-cell of $X \times Y$, so P is closed in the CW topology on $X \times Y$. We will show it is not closed in the product topology by showing that (x_0, y_0) lies in its closure, where x_0 is the common endpoint of the intervals I_s and y_0 is the common endpoint of the intervals I_j.

A basic open set containing (x_0, y_0) in the product topology has the form $U \times V$ where $U = \bigvee_s [0, a_s)$ and $V = \bigvee_j [0, b_j)$. It suffices to show that P has nonempty intersection with $U \times V$. Choose a sequence $t = (t_1, t_2, \cdots)$ with $t_j > j$ and $t_j > 1/b_j$ for all j, and choose an integer $k > 1/a_t$. Then $t_k > k > 1/a_t$ hence $1/t_k < a_t$. We also have $1/t_k < b_k$. So $(1/t_k, 1/t_k)$ is a point of P that lies in $[0, a_t) \times [0, b_k)$ and hence in $U \times V$.

Euclidean Neighborhood Retracts

At certain places in this book it is desirable to know that a given compact space is a retract of a finite simplicial complex, or equivalently (as we shall see) a retract of a neighborhood in some Euclidean space. For example, this condition occurs in the Lefschetz fixed point theorem, and it was used in the proof of Alexander duality. So let us study this situation in more detail.

Theorem A.7. *A compact subspace K of $\mathbb{R}^n$ is a retract of some neighborhood iff K is locally contractible in the weak sense that for each $x \in K$ and each neighborhood U of x in K there exists a neighborhood $V \subset U$ of x such that the inclusion $V \hookrightarrow U$ is nullhomotopic.*

Note that if K is a retract of some neighborhood, then it is a retract of every smaller neighborhood, just by restriction of the retraction. So it does not matter if we require the neighborhoods to be open. Similarly it does not matter if the neighborhoods U and V in the statement of the theorem are required to be open.

Proof: Let us do the harder half first, constructing a retraction of a neighborhood of K onto K under the local contractibility assumption. The first step is to put a CW structure on the open set $X = \mathbb{R}^n - K$, with the size of the cells approaching zero near K. Consider the subdivision of $\mathbb{R}^n$ into unit cubes of dimension n with vertices at the points with integer coordinates. Call this collection of cubes C_0. For an integer $k > 0$, we can subdivide the cubes of C_0 by taking n-dimensional cubes of edgelength $1/2^k$ with vertices having coordinates of the form $i/2^k$ for $i \in \mathbb{Z}$. Denote this collection of cubes by C_k. Let $A_0 \subset C_0$ be the set of cubes disjoint from K, and inductively, let $A_k \subset C_k$ be the set of cubes disjoint from K and not contained in cubes of A_j for $j < k$. The open set X is then the union of all the cubes in the combined collection $A = \bigcup_k A_k$. Note that the collection A is locally finite: Each point of X has a neighborhood meeting only finitely many cubes in A, since the point has a positive distance from the closed set K.

If two cubes of A intersect, their intersection is an i-dimensional face of one of them for some $i < n$. Likewise, when two faces of cubes of A intersect, their intersection is a face of one of them. This implies that the open faces of cubes of A that are minimal with respect to inclusion among such faces form the cells of a CW structure on X, since the boundary of such a face is a union of such faces. The vertices of this CW structure are thus the vertices of all the cubes of A, and the n-cells are the interiors of the cubes of A.

Next we define inductively a subcomplex Z of this CW structure on X and a map $r: Z \to K$. The 0-cells of Z are exactly the 0-cells of X, and we let r send each 0-cell to the closest point of K, or if this is not unique, any one of the closest points of K. Assume inductively that Z^k and $r: Z^k \to K$ have been defined. For a cell e^{k+1} of X with boundary in Z^k, if the restriction of r to this boundary extends over e^{k+1} then we include e^{k+1} in Z^{k+1} and we let r on e^{k+1} be such an extension that is not too large, say an extension for which the diameter of its image $r(e^{k+1})$ is less than twice the infimum of the diameters for all possible extensions. This defines Z^{k+1} and $r: Z^{k+1} \to K$. At the end of the induction we set $Z = Z^n$.

It remains to verify that by letting r equal the identity on K we obtain a continuous retraction $Z \cup K \to K$, and that $Z \cup K$ contains a neighborhood of K. Given a point $x \in K$, let U be a ball in the metric space K centered at x. Since K is locally contractible, we can choose a finite sequence of balls in K centered at x, of the form $U = U_n \supset V_n \supset U_{n-1} \supset V_{n-1} \supset \cdots \supset U_0 \supset V_0$, each ball having radius equal to some small fraction of the radius of the preceding one, and with V_i contractible in U_i. Let $B \subset \mathbb{R}^n$ be a ball centered at x with radius less than half the radius of V_0, and let Y be the subcomplex of X formed by the cells whose closures are contained in B. Thus $Y \cup K$ contains a neighborhood of x in $\mathbb{R}^n$. By the choice of B and the definition of r on 0-cells we have $r(Y^0) \subset V_0$. Since V_0 is contractible in U_0, r is defined on the 1-cells of Y. Also, $r(Y^1) \subset V_1$ by the definition of r on 1-cells and the fact that U_0 is much smaller than V_1. Similarly, by induction we have r defined on Y^i with $r(Y^i) \subset V_i$ for all i. In particular, r maps Y to U. Since U could be arbitrarily small, this shows that extending r by the identity map on K gives a continuous map $r: Z \cup K \to K$. And since $Y \subset Z$, we see that $Z \cup K$ contains a neighborhood of K by the earlier observation that $Y \cup K$ contains a neighborhood of x. Thus $r: Z \cup K \to K$ retracts a neighborhood of K onto K.

Now for the converse. Since open sets in $\mathbb{R}^n$ are locally contractible, it suffices to show that a retract of a locally contractible space is locally contractible. Let $r: X \to A$

be a retraction and let $U \subset A$ be a neighborhood of a given point $x \in A$. If X is locally contractible, then inside the open set $r^{-1}(U)$ there is a neighborhood V of x that is contractible in $r^{-1}(U)$, say by a homotopy $f_t : V \to r^{-1}(U)$. Then $V \cap A$ is contractible in U via the restriction of the composition rf_t. □

A space X is called a **Euclidean neighborhood retract** or **ENR** if for some n there exists an embedding $i : X \hookrightarrow \mathbb{R}^n$ such that $i(X)$ is a retract of some neighborhood in $\mathbb{R}^n$. The preceding theorem implies that the existence of the retraction is independent of the choice of embedding, at least when X is compact.

Corollary A.8. *A compact space is an ENR iff it can be embedded as a retract of a finite simplicial complex. Hence the homology groups and the fundamental group of a compact ENR are finitely generated.*

Proof: A finite simplicial complex K with n vertices is a subcomplex of a simplex Δ^{n-1}, and hence embeds in $\mathbb{R}^n$. The preceding theorem then implies that K is a retract of some neighborhood in $\mathbb{R}^n$, so any retract of K is also a retract of such a neighborhood, via the composition of the two retractions. Conversely, let K be a compact space that is a retract of some open neighborhood U in $\mathbb{R}^n$. Since K is compact it is bounded, lying in some large simplex $\Delta^n \subset \mathbb{R}^n$. Subdivide Δ^n, say by repeated barycentric subdivision, so that all simplices of the subdivision have diameter less than the distance from K to the complement of U. Then the union of all the simplices in this subdivision that intersect K is a finite simplicial complex that retracts onto K via the restriction of the retraction $U \to K$. □

Corollary A.9. *Every compact manifold, with or without boundary, is an ENR.*

Proof: Manifolds are locally contractible, so it suffices to show that a compact manifold M can be embedded in $\mathbb{R}^k$ for some k. If M is not closed, it embeds in the closed manifold obtained from two copies of M by identifying their boundaries. So it suffices to consider the case that M is closed. By compactness there exist finitely many closed balls $B_i^n \subset M$ whose interiors cover M, where n is the dimension of M. Let $f_i : M \to S^n$ be the quotient map collapsing the complement of the interior of B_i^n to a point. These f_i's are the components of a map $f : M \to (S^n)^m$ which is injective since if x and y are distinct points of M with x in the interior of B_i^n, say, then $f_i(x) \neq f_i(y)$. Composing f with an embedding $(S^n)^m \hookrightarrow \mathbb{R}^k$, for example the product of the standard embeddings $S^n \hookrightarrow \mathbb{R}^{n+1}$, we obtain a continuous injection $M \hookrightarrow \mathbb{R}^k$, and this is a homeomorphism onto its image since M is compact. □

Corollary A.10. *Every finite CW complex is an ENR.*

Proof: Since CW complexes are locally contractible, it suffices to show that a finite CW complex can be embedded in some $\mathbb{R}^n$. This is proved by induction on the number of cells. Suppose the CW complex X is obtained from a subcomplex A by attaching a cell e^k via a map $f:S^{k-1}\to A$, and suppose that we have an embedding $A\hookrightarrow\mathbb{R}^m$. Then we can embed X in $\mathbb{R}^k\times\mathbb{R}^m\times\mathbb{R}$ as the union of $D^k\times\{0\}\times\{0\}$, $\{0\}\times A\times\{1\}$, and all line segments joining points $(x,0,0)$ and $(0,f(x),1)$ for $x\in S^{k-1}$. □

Spaces Dominated by CW Complexes

We have been considering spaces which are retracts of finite simplicial complexes, and now we show that such spaces have the homotopy type of CW complexes. In fact, we can just as easily prove something a little more general than this. A space Y is said to be **dominated** by a space X if there are maps $Y\xrightarrow{i}X\xrightarrow{r}Y$ with $ri\simeq\mathbb{1}$. This makes the notion of a retract into something that depends only on the homotopy types of the spaces involved.

Proposition A.11. *A space dominated by a CW complex is homotopy equivalent to a CW complex.*

Proof: Recall from §3.F that the mapping telescope $T(f_1,f_2,\cdots)$ of a sequence of maps $X_1\xrightarrow{f_1}X_2\xrightarrow{f_2}X_3\longrightarrow\cdots$ is the quotient space of $\coprod_i(X_i\times[i,i+1])$ obtained by identifying $(x,i+1)\in X_i\times[i,i+1]$ with $(f(x),i+1)\in X_{i+1}\times[i+1,i+2]$. We shall need the following elementary facts:

(1) $T(f_1,f_2,\cdots)\simeq T(g_1,g_2,\cdots)$ if $f_i\simeq g_i$ for each i.
(2) $T(f_1,f_2,\cdots)\simeq T(f_2,f_3,\cdots)$.
(3) $T(f_1,f_2,\cdots)\simeq T(f_2f_1,f_4f_3,\cdots)$.

The second of these is obvious. To prove the other two we will use Proposition 0.18, whose proof applies not just to CW pairs but to any pair (X_1,A) for which there is a deformation retraction of $X_1\times I$ onto $X_1\times\{0\}\cup A\times I$. To prove (1) we regard $T(f_1,f_2,\cdots)$ as being obtained from $\coprod_i(X_i\times\{i\})$ by attaching $\coprod_i(X_i\times[i,i+1])$. Then we can obtain $T(g_1,g_2,\cdots)$ by varying the attaching map by homotopy. To prove (3) we view $T(f_1,f_2,\cdots)$ as obtained from the disjoint union of the mapping cylinders $M(f_{2i})$ by attaching $\coprod_i(X_{2i-1}\times[2i-1,2i])$. By sliding the attachment of $X\times[2i-1,2i]$ to $X_{2i}\subset M(f_{2i})$ down the latter mapping cylinder to X_{2i+1} we convert $M(f_{2i-1})\cup M(f_{2i})$ into $M(f_{2i}f_{2i-1})\cup M(f_{2i})$. This last space deformation retracts onto $M(f_{2i}f_{2i-1})$. Doing this for all i gives the homotopy equivalence in (3).

Now to prove the proposition, suppose that the space Y is dominated by the CW complex X via maps $Y\xrightarrow{i}X\xrightarrow{r}Y$ with $ri\simeq\mathbb{1}$. By (2) and (3) we have $T(ir,ir,\cdots)\simeq T(r,i,r,i,\cdots)\simeq T(i,r,i,r,\cdots)\simeq T(ri,ri,\cdots)$. Since $ri\simeq\mathbb{1}$, $T(ri,ri,\cdots)$ is homotopy equivalent to the telescope of the identity maps $Y\to Y\to Y\to\cdots$, which is $Y\times[0,\infty)\simeq Y$. On the other hand, the map ir is homotopic to a cellular map $f:X\to X$, so $T(ir,ir,\cdots)\simeq T(f,f,\cdots)$, which is a CW complex. □

One might ask whether a space dominated by a finite CW complex is homotopy equivalent to a finite CW complex. In the simply-connected case this follows from Proposition 4C.1 since such a space has finitely generated homology groups. But there are counterexamples in the general case; see [Wall 1965].

In view of Corollary A.10 the preceding proposition implies:

Corollary A.12. *A compact manifold is homotopy equivalent to a CW complex.* □

One could ask more refined questions. For example, do all compact manifolds have CW complex structures, or even simplicial complex structures? Answers here are considerably harder to come by. Restricting attention to closed manifolds for simplicity, the present status of these questions is the following. For manifolds of dimensions less than 4, simplicial complex structures always exist. In dimension 4 there are closed manifolds that do not have simplicial complex structures, while the existence of CW structures is an open question. In dimensions greater than 4, CW structures always exist, but whether simplicial structures always exist is unknown, though it is known that there are n-manifolds not having simplicial structures locally isomorphic to any linear simplicial subdivision of $\mathbb{R}^n$, for all $n \geq 4$. For more on these questions, see [Kirby & Siebenmann 1977] and [Freedman & Quinn 1990].

Exercises

1. Show that a covering space of a CW complex is also a CW complex, with cells projecting homeomorphically onto cells.

2. Let X be a CW complex and x_0 any point of X. Construct a new CW complex structure on X having x_0 as a 0-cell, and having each of the original cells a union of the new cells. The latter condition is expressed by saying the new CW structure is a **subdivision** of the old one.

3. Show that a CW complex is path-connected iff its 1-skeleton is path-connected.

4. Show that a CW complex is locally compact iff each point has a neighborhood that meets only finitely many cells.

5. For a space X, show that the identity map $X_c \to X$ induces an isomorphism on π_1, where X_c denotes X with the compactly generated topology.

The Compact-Open Topology

By definition, the compact-open topology on the space X^Y of maps $f:Y\to X$ has a subbasis consisting of the sets $M(K,U)$ of mappings taking a compact set $K \subset Y$ to an open set $U \subset X$. Thus a basis for X^Y consists of sets of maps taking a finite number of compact sets $K_i \subset Y$ to open sets $U_i \subset X$. If Y is compact, which is the only case we consider in this book, convergence to $f \in X^Y$ means, loosely speaking,