Lecture on CW-complexes

References: Hatcher is our text. Lee1 means Lee's Intro to Topological Manifolds. Lee2 means Lee's Intro to Smooth Manifolds.

1 Discs and spheres

Disks and spheres are the basic building blocks for CW-complexes, so it's worthwhile to lay out a few basic facts about them.

1.1 Discs

By an n-disk I mean what some call an n-ball: $D^n := \{x \in \mathbb{R}^n : |x| \leq 1\}$. Its boundary is the $(n-1)$-sphere S^{n-1}. Of course for many purposes we are willing to replace D^n by any space homeomorphic to it, so we recall:

Proposition 1.1 *Let X be a compact convex subset of $\mathbb{R}^n$ with nonempty interior. Then X is homeomorphic to D^n by a homeomorphism preserving the boundary.*

For a proof see Lee1, Prop. 5.1. Thus the following spaces are all homeomorphic to D^n:

- the n-cube I^n
- more generally any product $D^{m_1} \times ... \times D^{m_k}$ with $\sum m_i = n$
- the n-simplex Δ_n defined as the convex hull of the standard basis vectors in $\mathbb{R}^{n+1}$, i.e. $\{(x_0, ..., x_n) : x_i \geq 0 \sum x_i = 1\}$.

1.2 Spheres

It follows that the boundaries of the above spaces are homeomorphic to S^{n-1}. So we can think of the boundary of an n-simplex as an $(n-1)$-sphere, and so on.

Another thing we want is a more coordinate-free way of looking at spheres. Since S^n is canonically homeomorphic to the one-point compactification of $\mathbb{R}^n$, we can just as well start from an arbitrary n-dimensional vector space V over $\mathbb{R}$ (resp. $\mathbb{C}$); its one-point compactification is denoted S^V and is homeomorphic to S^n (resp. S^{2n}). This gives a nice way of thinking about smash products (Hatcher p.10) of spheres, via the following:

Exercise. Suppose X, Y are locally compact Hausdorff spaces, and let X^+ and so on denote the one-point compactifications. Then there is a natural homeomorphism $(X \times Y)^+ \cong X^+ \wedge Y^+$. (Here we take the points at infinity as our base points.)

In particular, if V, W are finite-dimensional vector spaces then $S^{V \oplus W}$ is naturally homeomorphic to $S^V \wedge S^W$. Even more particularly, $S^m \wedge S^n \cong S^{m+n}$.

On the other hand, we can consider the unit sphere in an n-dimensional real vector space V equipped with an inner product, still deliberately avoiding a choice of basis. This is denoted $S(V)$, and is homeomorphic to S^{n-1}. Then one can ask for a description of $S(V \oplus W)$ in terms of $S(V)$, $S(W)$. Here $V \oplus W$ has the natural inner product coming from the inner products on V, W.

Exercise. $S(V \oplus W)$ is naturally homeomorphic to the join $S(V) \star S(W)$. (For the join see Hatcher p. 9.) In particular $S^{m-1} \star S^{n-1} \cong S^{m+n-1}$.

The one good picture one can draw for the exercise is $S^0 \star S^1 \cong S^2$. Be sure to do this! And by the way, never neglect the 0-sphere. It's true that the poor fellow has only two points ± 1, but it is a perfectly respectable sphere and indeed sometimes plays a crucial role.

Here's one more description of a sphere that we'll use frequently: the quotient space D^n/S^{n-1} is homeomorphic to S^n. A coordinate-free way of stating this is as follows:

Proposition 1.2 *Let V be a finite-dimensional real vector space with inner product, and let $D(V)$ denote the closed unit disc. Then $D(V)/S(V)$ is homeomorphic to S^V.*

Proof: Do it yourself or fill in the details of the following: The argument of Lee1 Example 2.25 shows that the interior $Int\, D(V)$ is homeomorphic to V and hence these two spaces have homeomorphic one-point compactifications. Now recall the useful fact: If X is a compact Hausdorff space and A is a closed subspace, then X/A is homeomorphic to $(X - A)^+$. Applying this to $S(V) \subset D(V)$ completes the proof.

2 CW-complexes

Note: Few examples are given here; for that see Hatcher. The real and complex projective space examples are especially important.

CW-complexes were invented by J.H.C. Whitehead (not to be confused with George Whitehead). The idea is to provide a class of topological spaces which is broad enough to include most examples of interest and yet well-behaved enough to have manageable homotopy and homology. The intuition is simple: When you order a CW-complex from Spaces-R-Us, you get a package of 0-discs (points/vertices), 1-discs (unit intervals), 2-discs and so on, together with assembly instructions telling you how to build your complex inductively by attaching each 1-disc to either one or two vertices along its boundary, thereby obtaining the "1-skeleton" of your complex, then attaching each 2-disc to the 1-skeleton via a continuous map defined on its boundary S^1, thereby obtaining the 2-skeleton, and so on. There might be infinitely many discs, possibly of arbitrarily high dimension, but other than increasing the shipping costs this doesn't cause any problems.

The rigorous definition given in Hatcher is based directly on this intuitive idea. But it's a bit awkward to have the definition itself based on an inductive construction. It's better, in my opinion, to give the definition in one step. Then one can go back and show the "one step" and inductive definitions agree. So here's our official definition:

Definition. A CW-complex is a Hausdorff space X partitioned into subspaces $e_{n,\alpha}$ called n-cells, where $n \geq 0$ and α ranges over some index set, together with *characteristic maps* $\phi_{n,\alpha} : D^n \longrightarrow X$ subject to the following conditions.

A1. The restriction of $\phi_{n,\alpha}$ to $Int\, D^n$ is a homeomorphism onto $e_{n,\alpha}$. The *dimension* of the cell is n.

A2. The boundary of each n-cell (i.e. its closure minus the cell itself) is contained in a finite union of cells of dimension $< n$.

A3. A subset of X is closed if and only if its intersection with the closure of each cell is closed.

This definition is equivalent to the one given in Hatcher, Appendix A. The only tricky part of the proof is showing that Hatcher's definition implies the Hausdorff condition in ours; see Hatcher Proposition A3. Some basic terminology for CW-complexes:

- A *subcomplex* of a CW-complex X is a closed subspace A that is a union of cells of X.
- A *finite CW-complex* is a CW-complex with finitely many cells.
- The *dimension* of a CW-complex X is the maximal n (if it exists) such that X has an n-cell. If no such n exists then X is *infinite-dimensional*.
- The n-skeleton of X is the subcomplex X^n consisting of all cells of dimension $\leq n$.

Next, some comments on conditions A1-3 are in order.

A1. Since X is Hausdorff and D^n is compact, it is only necessary to show that $\phi_{n,\alpha}$ restricted to the interior is a bijection onto $e_{n,\alpha}$; then this restriction is automatically a homeomorphism. Notice however that it is not enough to show that each $e_{n,\alpha}$ is homeomorphic to an open n-disc. It is essential that the homeomorphism in question comes from a characteristic map of the closed disc as above. The restriction of $\phi_{n,\alpha}$ to S^{n-1} (where $n > 0$) is the *attaching map* for the cell. Note that the attaching maps need not be injective, and can even be constant maps.

A2. The boundary of a cell is equal to the image of its attaching map. So this condition says that for $n > 0$, $\phi_{n,\alpha}(S^{n-1})$ is contained in a finite union of cells of dimension $< n$. The requirement that it be contained in a *finite* union of cells is sometimes called "closure-finite"; this where the "C" in "CW" come from. This condition is essential if we want the 0-skeleton of our complex to be discrete. And we certainly do, as otherwise we would lose the intuitive picture of what a complex should be. For example (cf. Hatcher p. 521) if we omit this condition we could take D^2 and make it a pseudo-CW-complex with every (!) point of the boundary a 0-cell, the usual 2-cell, and no other cells.

Exercise: Show from the above definition that the 0-skeleton of a CW-complex has the discrete topology.

A3. This condition is sometimes referred to as the "weak topology"; this is where the "W" in "CW" comes from. But since the "weak" topology corresponds to what functional analysts call the "strong" topology, it seems best to avoid such terminology altogether. Also, a possibly clearer equivalent way of stating A3 is that a subset of X is closed if and only if its intersection with every finite subcomplex is closed. (Check this!) Then we call this the "direct limit topology" defined by the finite subcomplexes. The way to think of it is that a map out of X is continuous if and only if its restriction to every finite subcomplex is continuous. (In other words, noting that the finite subcomplexes of X form a directed set under inclusion, X is the direct limit of its finite subcomplexes in the category of topological spaces.) In fact if X_α is any collection of subcomplexes forming a directed set under inclusion and whose union is X, then X is the direct limit of the X_α's.

Yet another useful equivalent to A3: The map $\coprod_{n,\alpha} D^n \longrightarrow X$ given by the disjoint union of all the characteristic maps is a quotient map. (Check this!) Your assembly instructions from Spaces-R-Us probably include an "exploded" diagram showing all the pieces, nuts and bolts at once with arrows indicating the attachments; this neatly reflects the quotient map point of view.

Further discussion of A3 can be found in the next section.

3 Compactly-generated spaces and CW-complexes

Let X be a Hausdorff space. We say that X *compactly generated* if a subset of X is closed if and only if its intersection with every compact subset is closed. (In other words, X is the direct limit of its compact subspaces.) Note that in order to check this condition, it is enough to check it for a cofinal family of compact subsets (i.e. a family of compact subsets K_α such that every compact subset is contained in some K_α). Since finite subcomplexes of a CW-complex are cofinal in this sense (check this!), we see that every CW-complex is compactly-generated.

On the other hand, one can't replace A3 by the assumption that X is compactly-generated. For example, take X to be the subspace of $\mathbb{R}$ consisting of 0 and the points $1/n$, $n \geq 1$. Then X is compact and so is certainly compactly-generated, but it doesn't admit any CW-structure at all. For clearly the only conceivable choice of cells is to make every point a 0-cell, but then $X - \{0\}$ shows that A3 fails. A more interesting example is given by the "Hawaiian earring", namely the subspace of $\mathbb{R}^2$ consisting of the union of the circles C_n having center $(1/n, 0)$ and radius $1/n$. Again this space X is compactly-generated, indeed compact. Moreover we can partition it into one 0-cell and a countable set of 1-cells in the obvious way, so that A1 and A2 are satisfied. But this is not a CW-decomposition because in the CW-topology any subset consisting of one point from each 1-cell is closed, whereas in the given topology it has the origin as a limit point. In fact from our point of view the Hawaiian earring and its higher-dimensional analogs are extremely "pathological", a value judgement to be sure but the surprising fact is that our methods will have almost nothing to say about such spaces.

Given any Hausdorff X we can always retopologize it to get a new space κX: the underlying set is the same, but in the new topology a set is closed if and only if its intersection with with every compact subset is closed. Thus the identity map $\eta : \kappa X \longrightarrow X$ is a continuous bijection but not necessarily a homeomorphism, i.e. the compactly-generated topology may be strictly finer than the given topology (an example will be given shortly).

An interesting example of this phenomenon arises in connection with products. Suppose X and Y are CW-complexes, and consider $X \times Y$ with the product topology. It has obvious cells and characteristic maps, using the fact that a product of discs is homeomorphic to a disc, that clearly satisfy the conditions for a CW-complex—with the possible exception of

*A*3. Now

$$\eta : \kappa(X \times Y) \longrightarrow X \times Y$$

is a continuous bijection. It is a homeomorphism provided that either (a) X and Y both have a countable number of cells, or (b) at least one of X, Y is locally compact (see Hatcher Appendix A). In particular, it is convenient that for the cylinder $X \times I$ used to define homotopies, the two topologies coincide. Note also that the CW-topology on $X \times Y$ is the same as $\kappa(X \times Y)$; this just amounts to the fact that as K, L range over finite subcomplexes of X, Y, the products $K \times L$ are cofinal among all compact subspaces of $X \times Y$.

A counterexample. Suppose X (resp. Y) is a wedge of countably many (resp. cardinality of the continuum) copies of the unit interval I. Then $\eta : \kappa(X \times Y) \longrightarrow X \times Y$ is not a homeomorphism; in other words, the CW-topology on $X \times Y$ is strictly finer than the product topology obtained from the original CW-topologies of the factors. For the short but non-trivial proof, see the original 1952 paper "The topology of metric complexes" by Dowker, or try it yourself.

On the other hand, for the purposes of algebraic topology it is rarely worth fussing over this phenomenon, for the following reason: The standard algebraic invariants such as homotopy groups and homology groups (we haven't defined these yet, but think of the fundamental group) all involve mapping compact spaces such as discs, spheres and simplices into a given space X. If K is compact, then a map $K \longrightarrow X$ is continuous if and only if it is continuous as a map into κX. It follows that $\eta : \kappa X \longrightarrow X$ induces an isomorphism on homotopy groups and homology groups, so we are generally willing to simply replace X by κX whenever convenient.

One final note: Re-topologizing with the compactly-generated topology is of no use at all in examples such as the Hawaiian earring, because the latter space is compact Hausdorff and so already is compactly-generated. You could re-topologize with a CW-topology making it a countable wedge of circles, but this topology is wildly different from the original one. In fact the fundamental group of a CW-complex with countably many cells is countable (exercise), but it can be shown that π_1 of the Hawaiian earring is uncountable. The moral, once again, is that we'll be better off ignoring such spaces.

4 Exercises

Note for exercise 1. Let $\pi(X)$ denote the set of path-components of a space X. This is the same thing as the set of homotopy classes of maps from a point to X. If we fix a basepoint $x_0 \in X$, this in turn is identified in the evident way with the set of pointed (or "based") homotopy classes of maps $S^0 \longrightarrow X$ ($1 \in S^0$ is the basepoint), denoted $\pi_0(X, x_0)$. This might seem a pointless (ha ha) complication of the simpler $\pi(X)$, but it fits with $\pi_1(X, x_0)$ and the higher homotopy groups $\pi_n(X, x_0)$ to be considered later.

1. Let X be a CW-complex, and chose a basepoint $x_0 \in X^0$. Prove the following theorem (the basepoint x_0 is implicit):

a) $\pi_0 X^k \longrightarrow \pi_0 X$ is surjective for $k \geq 0$ and bijective for $k \geq 1$.
b) $\pi_1 X^k \longrightarrow \pi_1 X$ is surjective for $k \geq 1$ and an isomorphism for $k \geq 2$.

Suggestions: Reduce to the case of a finite complex and then use induction on the number of cells. The Lebesgue lemma and the Seifert-Kampen theorem also come to mind. Pay attention to basepoints.

Notes: 1. The above theorem foreshadows the much more general Cellular Approximation Theorem, to be proved later.

2. It follows for example that $\pi_1 \mathbb{R}P^n \cong \mathbb{Z}/2$ for $n \geq 2$ (generated by the inclusion $S^1 = \mathbb{R}P^1 \subset \mathbb{R}P^n$), and that $\mathbb{C}P^n$ is simply-connected for all n. Why?

2. Let X be a CW-complex with a countable number of cells. Show that for any choice of basepoint x_0, $\pi_1(X, x_0)$ is countable.

3. We will see later that for all $n \geq 0$, S^n is not contractible (obvious for $n = 0$; true for $n = 1$ by the fundamental group). Use the following approach, which is a sort of topological manifestation of "Cantor's/Hilbert's Hotel" to show that S^∞ *is* contractible. Let $S : \mathbb{R}^\infty \longrightarrow \mathbb{R}^\infty$ denote the shift map sending the i-th standard basis vector e_i to e_{i+1}. Show that S is nullhomotopic and also homotopic to the identity map, and deduce the result.

Recall that S^∞ is topologized as the direct limit of the S^n's, and be sure to justify continuity of all maps used.

5 Solutions to some exercises

First of all let me emphasize again a key point: A subspace K of a CW-complex X is compact if and only if it is contained in a finite subcomplex (the "if" is easy; for the "only if" see Hatcher A1). Be sure you understand the proof and how the direct limit topology is used.

1a. Surjectivity: It suffices to prove the case $k = 0$. In other words, we must show that every point $x \in X$ can be connected by a path to a point $y \in X^0$. By definition x lies in a unique cell $e_{n,\alpha}$. If $n = 0$, i.e. $x \in X^0$, we are done. So suppose $n > 0$ and $x = \phi_{n,\alpha}(a)$, where $a \in Int\, D^n$. Choose any path μ from a to a point of S^{n-1}. Then the path $\phi_{n,\alpha} \circ \mu$ joins x to a point of the $(n-1)$-skeleton, and we are done by induction on n.

Injectivity: In view of part (a), it suffices to show that if two points $x_0, x_1 \in X^1$ are joined by a path $\lambda : I \longrightarrow X$, then they are joined by a path $\mu : I \longrightarrow X^1$. For use in part (b), however, we prove something stronger.

Lemma 5.1 *With the above notation, λ is path-homotopic to a path $\mu : I \longrightarrow X^1$.*

Proof: Since I is compact, its image under λ lies in a finite subcomplex; hence we reduce at once to the case that X is a finite complex. Fix a cell $e := e_{n,\alpha}$ of maximal dimension, where we may assume $n \geq 2$. Then e is an open subset of X (why?). We will find a new path σ path-homotopic to λ such that $\sigma(I) \cap e = \emptyset$; then we are done by induction on the number of cells of X.

To this end, let p denote the center point of e (i.e. $\phi_{n,\alpha}(0)$) and consider the open cover of $X = U \cup V$ where $U = X - p$ and $V = e$. We will use such open covers frequently; note that the subcomplex $Y = X - e$ is a deformation retract of U (sometimes only the retraction is needed, not the deformation) via the obvious radial deformation, V is contractible, and $U \cap V$ is homeorphic to $Int\, D^n - 0$ and homotopy-equivalent to S^{n-1}.

By the Lebesgue covering lemma we may subdivide I so that on each interval $I_k = [t_k, t_{k+1}]$ of the subdivision, $\lambda(I_k)$ is contained in U or in V. Furthermore we may assume that adjacent intervals I_k, I_{k+1} do not both map into U and do not both map into V (if they did, just concatenate them to an interval $[t_k, t_{k+2}]$, and repeat as necessary). Now suppose $\lambda(I_k) \subset V$. Since the adjacent intervals (there are either one or two of them) map into U, the endpoints t_k, t_{k+1} are in U, i.e. $\neq p$. Clearly there exists a path (e.g. a "piecewise-linear" one) $\mu : I_k \longrightarrow V$ with $p \notin \mu(I_k)$. Since V is homeomorphic to a disc, any two paths with the same endpoints are path-homotopic. Moreover any path-homotopy from $\lambda|_{I_k}$ to μ extends to a path-homotopy of λ, by taking the stationary homotopy on the complement of I_k. Repeating this construction for each subinterval that maps into V, we obtain a path τ that is path-homotopic to λ and has $\tau(I) \subset U$. Applying the retraction of U onto Y yields the desired σ, completing the proof of the lemma and hence of the injectivity.

b) By the lemma, any loop $\alpha : I \longrightarrow X$ is path-homotopic to a loop lying in X^1. This proves the surjectivity.

For the injectivity, first observe that we may assume X is path-connected (since π_1 depends only on the path-component of the basepoint). Next, we reduce to the case that X is a finite complex. To see this, consider a general X and suppose $\alpha : I \longrightarrow X^2$ is a loop such that $i \circ \alpha$ is nullhomotopic, where $i : X^2 \longrightarrow X$ is the inclusion. A nullhomotopy is a map from the compact space $I \times I$ and therefore has image contained in some finite subcomplex K. So if we know the result for K, then α is nullhomotopic.

From now on we assume X is a finite path-connected CW-complex, and proceed by induction on the number of cells. At the inductive step we use the same notation as in the injectivity for 1a, except that now the cell e has dimension $n \geq 3$. We will apply the Seifert-van Kampen theorem to the open cover U, V. Let's check the hypotheses of S-vK, though, namely that $U, V, U \cap V$ are path-connected. This is obvious for V, and holds for $U \cap V$ (a disc with a point removed) since $n \geq 2$. Since U deformation retracts onto Y and Y contains the 1-skeleton, U is path-connected by 1a.

There is one other condition required for S-vK, namely that the basepoint used must lie in $U \cap V$, whereas the basepoint we started with does not. Using the Change of Basepoint Theorem (7.13 in Lee1), however, we see that it now suffices to choose a basepoint $q \in U \cap V$ and show that $\pi_1(U, q) \longrightarrow \pi_1(X, q)$ is an isomorphism. Since $n \geq 3$, $U \cap V = V - p$ is simply-connected, so S-vK tells us that $\pi_1 X$ is the free product of $\pi_1 U$ and $\pi_1 V$. But V is also simply-connected, indeed contractible, so $\pi_1 U \longrightarrow \pi_1 X$ is an isomorphism as claimed. QE freakin' D!

2. First of all, we may assume X is path-connected, with basepoint x_0. Next we reduce to the case X is a finite complex. Consider the evident map of sets

$$\coprod_{\alpha} \pi_1(X_\alpha, x_0) \longrightarrow \pi_1(X, x_0),$$

where X_α ranges over all path-connected finite subcomplexes containing x_0. This map is clearly surjective, by the same argument used in surjectivity for 1b. (In fact it induces an isomorphism from the direct limit of the $\pi_1 X_\alpha$'s; we'll consider direct limits systematically later). Since there are countably many X_α's, this reduces to the case X a finite complex. Moreover, by 1b surjectivity we reduce further to the case X 1-dimensional, i.e. a "graph". In this case $\pi_1 X$ is a finitely-generated free group (Lee1 Thm. 10.12; we'll prove this a different way soon, using the homotopy extension property). Any finitely-generated group, indeed any countably generated group, is countable; the proof is left to the reader. Done!