Contents

1. Wed 1/22/14

1.1. **Quotient Spaces and Attaching.** Let X be a topological space, let Z be any set, and let $f : X \to Z$ be an onto function. Then we can define a topology on Z by declaring $U \subset Z$ to be open iff $f^{-1}(U)$ is open in X. (The proof is pretty much immediate from the definition of a topology.)

> **Definition:**
> X = topological space
> $\sim$ = equivalence relation on points of X
> $X/\sim$ = set of equivalence classes $f : X \to X/\sim$ = map taking each point to its equivalence class
>
> The quotient topology is the finest[1] topology on $X/\sim$ that makes f continuous. I.e., U is open iff $f^{-1}(U)$ is open in X.

Variant: **Gluing.** Suppose we have two spaces X, Y and we want to glue them together by identifying some points of X with points of Y. Given a map $f : A \to Y$, where $A \subset X$, we can **glue X to Y along A via f** to obtain the space

$$(X \dot\cup Y)/(a \sim f(a)\ \forall a \in A)$$

or $(X \dot\cup Y)/f$ for short.

Example: Let X and Y be the line segments $\overline{ab}$ and $\overline{cd}$. Let $A = \{a, b\} \subseteq X$ and define f by $f(a) = c$, $f(b) = d$. The space $X \dot\cup Y/f$ is a circle. There's nothing to prevent us from mapping a and b to other points on Y, in which case we would get different-looking spaces.

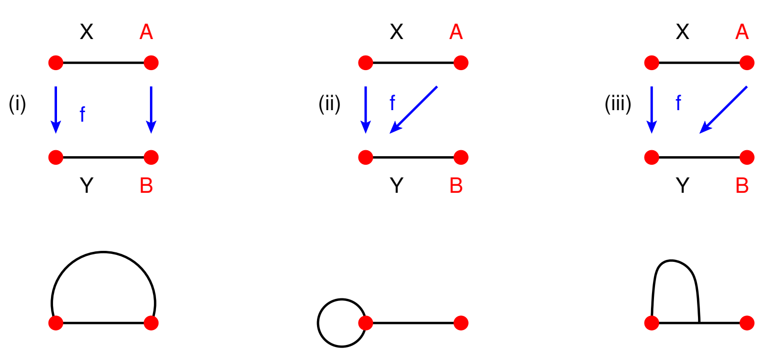

Example: Graphs. A graph consists of a bunch of vertices connected by edges; each edge is attached to two vertices called its endpoints (which need not be distinct). Such a space can be constructed as a union of line segments, one for each edge, quotiented out by attaching maps which glue vertices together. The only restriction on the attaching maps is that endpoints must be mapped to endpoints (so in the last figure, the first two are graphs but the third isn't).

Example: Let $D^2 = \{\mathbf{v} \in \mathbb{R}^2 \mid \|\mathbf{v}\| \leq 1\}$, the closed unit disk, and let $S^1 = \{\mathbf{v} \in \mathbb{R}^2 \mid \|\mathbf{v}\| = 1\}$, its boundary circle. Take two copies of D^2 and glue them together along S^1. (To be precise, the gluing map is a homeomorphism between the two copies of S^1 — we may as well consider the gluing map to be the identity.) The result is a 2-sphere. More generally, gluing together two copies of the closed unit ball $D^n \subseteq \mathbb{R}^n$ along their boundary $(n-1)$-spheres gives a n-sphere.

Example: Here is a weirder space. Take a copy of D^2 and a copy of S^1, and map the boundary of the disc to the S^1 by the attaching map

$$f(\cos\theta, \sin\theta) = (\cos 2\theta, \sin 2\theta).$$

In other words, wrap the disc twice around the circle. This is a perfectly well-defined (and extremely important) topological space, although it's not so easy to draw!

Example: A cylinder can be obtained by starting with a closed square and identifying its left and right sides L, R. We have to be careful about how we do this, however: if we first give the square a twist, we will get a Möbius strip instead of a cylinder. To put it another way, the kind of space we get depends on the choice of homeomorphism $L \to R$.

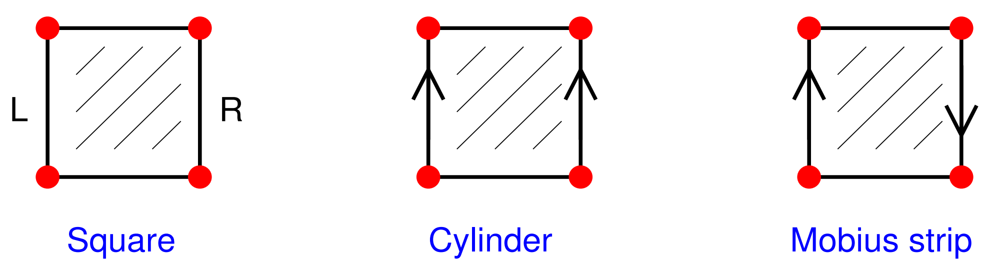

Rather than trying to write down these kinds of quotient spaces in some algebraic way, it's clearer and more convenient to draw them pictorially. We distinguish the two different possibilities by the directions of the arrows in L and R — the gluing map should make the arrows line up.

Question for future thought: How can we intrinsically distinguish the cylinder from the Möbius strip?

Example: We could identify the top and bottom edges of the square as well. Again, the gluing direction matters. There are three possibilities, depending on whether 2, 1 or 0 pairs of opposite edges "line up correctly".

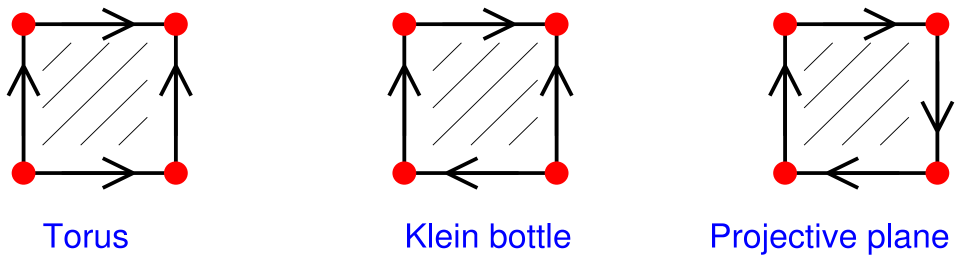

The torus is the easy one — it looks like a donut. You can make it from a cylinder by gluing the two boundary circles together. The Klein bottle is weirder. You could make it by starting with a cylinder and gluing the two boundary circles together — but somehow reversing the orientation of one of them, which you can't do in $\mathbb{R}^3$.

The projective plane is possibly the weirdest of all. You have started with a disk and identified every pair of opposite points on the boundary circle. (We've already seen this construction!) It turns out that this is the same thing as starting with a 2-sphere, and identifying every pair of antipodally opposite points on it.

All of these spaces are *2-manifolds*. A (topological) n-manifold is a space in which every point has a neighborhood homeomorphic to $\mathbb{R}^n$. (This basic definition can be souped up to define differentiable, smooth, holomorphic, and algebraic manifolds, depending on your category of choice.)

Again, food for thought: How do you distinguish between manifolds of a given dimension? Is there some set of numerical or algebraic invariants that classify manifolds? If so, how do you calculate those invariants for a given manifold?

For example, a 2-sphere and a torus are both clearly 2-manifolds. However, it would be surprising if they were homeomorphic. How do we measure the "holeyness" that a torus has but a sphere lacks?

2. Fri 1/24/14

2.1. Cell Complexes. A more careful definition of a cell complex.

A cell complex X is a topological space of the form $\bigcup_{n=0}^{\infty} X^n$, where the X^n are spaces called the *skeletons* of X and are defined as follows:

- X^0 has the discrete topology.
- $X^n = X^{n-1} \cupdot \bigcup_\alpha D^n_\alpha/\varphi$, where:
  - Each D^n_α is (a homeomorphic copy of) the closed n-ball B^n, so that it has a subspace ∂D^n_α corresponding to the $(n-1)$-sphere S^{n-1}. It comes with a continuous *attaching map* $\varphi_\alpha : \partial D^n_\alpha \to X^{d-1}$.
  - $\varphi = \bigcup_\alpha \varphi_\alpha$.

If $X = X^n$ for some n then X is *finite-dimensional* and its topology is determined by attaching. If not, then we have to do some additional work to say what the topology is (see the appendix in Hatcher).

2.2. Homotopy equivalence vs. homeomorphism. The "obvious" question of topology is this: when are two topological spaces the same? I.e., when are they homeomorphic? This turns out to often be the wrong question to ask. For example, no two of the letters H, I, T, X are homeomorphic, but they are all topologically trivial in a sense that we'll make clear soon. A more useful notion of topological sameness is *homotopy equivalence.*

Intuitively, we want to describe how a function $X \to Y$ might continuously evolve over time from a starting function F_0 to a final function F_1. We can describe this as a family of maps $\{F_t\}$ for all $t \in I = [0, 1]$. To say that the evolution is continuous is just to say that the map $F : X \times I \to Y$ is continuous, where $F(x, t) = F_t(x)$. To make this precise:

Definition: Two maps $f : X \to Y$ and $g : X \to Y$ are **homotopic** (written $f \simeq g$) if there is a map $F : X \times I \to Y$ such that $F(x, 0) = f(x)$ and $F(x, 1) = g(x)$ for all x. (I.e., $f = F_0$ and $g = F_1$).

Fact: $\simeq$ is an equivalence relation (proof left to the reader; not hard).

Definition: Two spaces X, Y are **homotopy-equivalent** if there are maps $f : X \to Y$ and $g : Y \to X$ such that $f \circ g \simeq \mathbb{1}_Y$ and $g \circ f \simeq \mathbb{1}_X$.

Example: Let Y be the unit circle and let $X = Y \times I$ be the annulus. Define f, g by $f(x) = (x, 0)$, $g(x, s) = x$. Then $g \circ f = \mathbb{1}_Y$ (no problem there) and $f \circ g$ is the map sending $(x, s) \mapsto (x, 0)$ for every x, s. Consider the homotopy $F : X \times I \to X$ defined by

$$F((x, s), t) = (x, st).$$

Then F is continuous, with

$$F_0(x, s) = F((x, s), 0) = (x, 0), \qquad F_1(x, s) = F((x, s), 1) = (x, s)$$

i.e.

$$F_0 = f \simeq F_1 = \mathbb{1}_X.$$

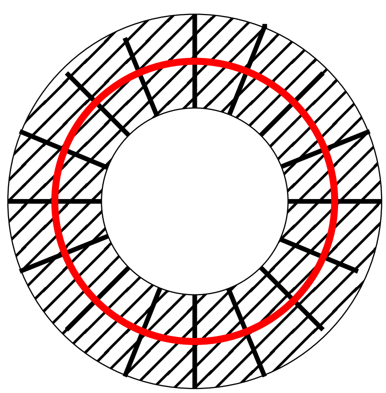

Note that $F_t(x, 0) = (x, 0)$ for every t. That is, the subspace Y is remaining fixed as X gradually contracts onto Y. This is an example of a very common kind of homotopy equivalence, a *deformation retract* (or *deformation retraction*).

Definition: Let X be a space and $A \subseteq X$ a subspace. A **deformation retract** from X to A is a homotopy $F : X \times I \to X$ such that $F_0 = \mathbb{1}_X$; $F_1(X) = A$; and $F_t|_A = \mathbb{1}_A$ for all times t.

Definition: A space X is *contractible* if it is homotopy-equivalent to a point.

Example: A subspace $X \subset \mathbb{R}^n$ is called *star-shaped* if there exists some $x_0 \in X$ such that for every $y \in X$, the line segment from x_0 to y is contained in X. It's not hard to show that every star-shaped set (such as a line segment and the letters T and X) are contractible. In particular, all convex sets are star-shaped, hence contractible.

Contractibility turns out to be an excellent notion of topological triviality. Note that all contractible spaces are homotopy-equivalent (because homotopy-equivalence is an equivalence relation).

So, how do you prove that a space is *not* contractible (or more generally, determine when two spaces are not homotopy-equivalent)? Hold that thought.

3. Mon 1/27/14

Definition: Two maps $f : X \to Y$ and $g : X \to Y$ are **homotopic** (written $f \simeq g$) if there is a map $F : X \times I \to Y$ such that $F(x, 0) = f(x)$ and $F(x, 1) = g(x)$ for all x. (I.e., $f = F_0$ and $g = F_1$).

3.1. **More on homotopy equivalence.** One of the most important cases is when $X = S^1$. A map $S^1 \to Y$ is then a **closed path** in Y — think of it as a rubber band in Y. A homotopy is something that continuously deforms the rubber band through a family of closed curves.

Definition: Two spaces X, Y are **homotopy-equivalent** if there are maps $f : X \to Y$ and $g : Y \to X$ such that $f \circ g \simeq \mathbb{1}_Y$ and $g \circ f \simeq \mathbb{1}_X$.

Definition: Let X be a space and $A \subseteq X$ a subspace. A **deformation retract** from X to A is a homotopy $F : X \times I \to X$ such that $F_0 = \mathbb{1}_X$; $F_1(X) = A$; and $F_t|_A = \mathbb{1}_A$ for all times t.

Example: For every X, the **cylinder** over X is $X \times I$. We can identify X homeomorphically with any slice $X \times \{t\}$ of the cylinder. Consider the map $F : (X \times I) \times I \to X \times I$ given by $F((x, s), t) = (x, s(1-t))$. Then $F_0 = \mathbb{1}_{X \times I}$ and F_1 is projection $X \times I \to X$. Moreover, the slice $X \times \{0\}$ is fixed by every F_t¿ Therefore F is a deformation retraction of $X \times I$ onto X.

Proposition: If $F : X \times I \to X$ is a deformation retraction of X onto A, then $X \simeq A$.

Proof: Define $f : X \to A$ and $g : A \to X$ by $f = F_1$ and $g = \mathbb{1}_A$ (regarded as the inclusion $A \hookrightarrow X$). Then $f \circ g = \mathbb{1}_A$ and $g \circ f = f = F_1 \simeq F_0 = \mathbb{1}_X$.

Definition: A space X is **contractible** if it is homotopy-equivalent to a point.

Example: $X = I \times I$. Pick any point $p \in X$. For $x \in X$ and $t \in I$, define $F(x, t)$ to be the point you get by constructing the line segment $\overline{xp}$

and walking t of the way along it from x. In linear coordinates, $F(x,t) = (1-t)x + tp$. This is a deformation retraction from X onto $\{p\}$, so X is contractible.

Example: A subspace $X \subset \mathbb{R}^n$ is called *starlike* if there exists some $x_0 \in X$ such that for every $y \in X$, the line segment from x_0 to y is contained in X. It's not hard to show that every starlike set (such as a line segment and the letters T and X) is contractible. In particular, all convex sets are starlike, hence contractible.

Example: Contractible implies path-connected (and therefore connected). Suppose that $F : X \times I \to X$ is a homotopy with $F(x,0) = x$ and $F(x,1) = x_0$. For every point $z \in X$, the function $\gamma(t) = F(z,t)$ is a path from z to x_0.

Example: The **cone** over X is the space $CX = X \times I/(x,0) \sim (y,0)\ \forall x, y$. In other words, take a cylinder over X and squash all points at time 0 to a point, called the **cone point**. Every cone is contractible, because the map

$$F : CX \times I \to CX, \quad ((x,s),t) \mapsto (x, s(1-t))$$

is a deformation retraction of CX onto the cone point.

Example: The circle S^1 is not contractible, although this will take some proving. Therefore, neither is the cylinder or the Möbius strip. Neither is S^2, although this will take even more proving.

Proposition: Two spaces X, Y are homotopy-equivalent if and only if there exists a third space that deformation-retracts onto both X and Y.

The $\Longleftarrow$ direction is obvious, since deformation retraction induces a homotopy equivalence. For the forward direction, we need a new construction.

Definition: The **mapping cylinder** of the map $f : X \to Y$ is

$$M_f = (X \times I) \cup Y \ \ / \ \ (x,1) \sim f(x).$$

That is, we attach the cylinder $X \times I$ to Y by gluing along the time-1 slice, using f as the gluing map. The mapping cylinder deformation-retracts to Y

by the homotopy $H : M_f \times I \to Y$ defined by

$$H(p,t) = \begin{cases} (x, st) & \text{if } p = (x,s) \in X \times I, \\ p & \text{if } p \in Y. \end{cases}$$

Notice that the two cases agree when $s = 0$, so H is well-defined. Also, note that $M_f \simeq Y$ for any continuous map f.

We want to show that if f is a homotopy equivalence, then also $M_f \simeq X$.; that is, there is a function $g : Y \to X$ such that $f \circ g \simeq \mathbb{1}_Y$, say via a homotopy G.

The map $f \circ g$ can be viewed as mapping Y to $X \times \{0\}$ (identifying $f(g(y))$ with $(f(g(y)), 0) \sim g(y)$). We know we can homotope $f \circ g$ to $\mathbb{1}_Y$. How can we extend this homotopy to the whole mapping cylinder so as to fix X? (One problem is that we don't know that this homotopy fixes $X \times \{0\}$ pointwise).

Some of you may have heard of a thing called a **mapping cone**. This is defined as

$$C_f = CX \cup Y/(x,1) \sim f(x) = M_f/(x,0) \sim (y,0).$$

In other words, take the mapping cylinder and squash the time-0 copy of X to a point. It is not in general homotopy-equivalent to either X or Y; it corresponds to contracting the subspace $f(X)$ of Y.

4. Wed 1/29/14

4.1. **Criteria for Homotopy Equivalence.** The definition of homotopy and homotopy equivalence can be a hassle to work with directly. We often just care that two spaces are homotopy-equivalent and may not want to construct a homotopy equivalence explicitly.

Homotopy Equivalence Criterion 1: Let $A \subset X$, so that we have a quotient map $q : X \to X/A$. If A is contractible, then q is a homotopy equivalence.

Homotopy Equivalence Criterion 2: Let $A \subset X$ and $f, g : A \to Y$. If $f \simeq g$, then $X \cup Y/f \simeq X \cup Y/g$.

Intuition: (1) If A was topologically trivial to begin with, then squashing it down to a point shouldn't change homotopy type.

(2) If f can be morphed into g, then using them as attaching maps should produce spaces that can be morphed into each other.

Cell complexes have these properties (as we will prove soon); this is one good reason to work with them. But first some examples.

Example: Graphs. Let G be a finite path-connected graph with v vertices and e edges.

Let x, y be distinct vertices that share an edge a. The operation of **contraction** shrinks the edge gradually until the two endpoints are identified. Topologically, we are passing from G to the quotient space G/a. This is a homotopy equivalence, although it's not clear what the homotopy inverse should look like — G/a is not in general a subgraph of G.

This is the tough part of proving HEC2 directly — how do you lift X/A back to X? But if we assume HEC2, can we classify connected graphs by homotopy type?

Definition: Let X, Y be two path-connected spaces that are disjoint. The **wedge (sum)** $X \vee Y$ is the space obtained by identifying a point of X with a point of Y. (Which point? Up to homotopy, it doesn't matter, by HEC2 — here you need path-connectedness. Check this!)

In particular, if we have a space that is a wedge of spheres, it can be specified up to homotopy by counting the spheres of each dimension in the wedge.

Theorem: $G \simeq \bigvee^{e-v+1} S^1$ (i.e., a wedge, or "bouquet", of $e - v + 1$ circles).

Example: The torus with meridional discs. Figure from Hatcher, p.12.

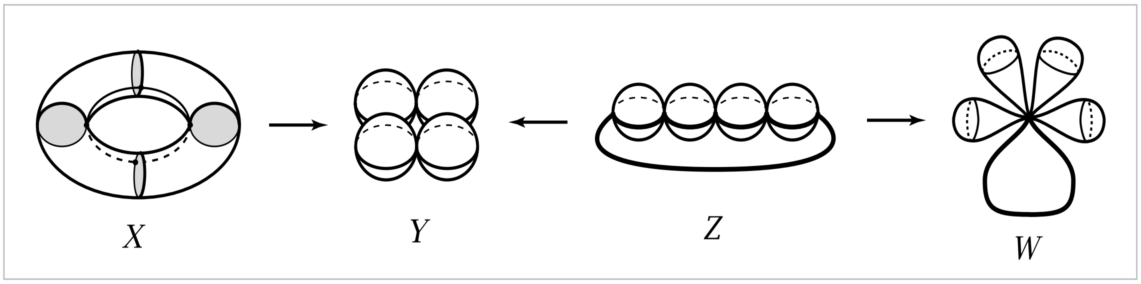

X is the torus with n meridional discs attached. (Here $n = 4$.)

Each of those discs is contractible. Contracting them all to points gives a quotient space Y that looks like a necklace of spheres. By HEC1, the quotient map $X \to Y$ is a homotopy equivalence.

Draw a circle around all the spheres. "Uncontract" part of the circle to get Z. In other words, $Y \to Z$ is a quotient map obtained by shrinking the outside curve to a point.

Now slide the spheres around so that they are all attached at the same point. (This is really an application of HEC1. If X is one of the spheres, A is a point in X, and Y is another sphere, then we are using the fact that all maps $A \to Y$ are homotopic, precisely because Y is path-connected.) Alternately, you could contract the back part of the circle to a point, which would not mess up any of the spheres by HEC1.

We have now shown that $X \simeq W = S^1 \vee \bigvee^4 S^2$, although we haven't constructed an explicit homotopy equivalence.

What happens if you contract a non-contractible subspace?[2] It depends. For example if $A = X = S^1$ then of course X/A is a point. (here contracting a circle makes things simpler.) But if $X = S^2$ and A is the equator, then $X/A = S^2 \vee S^2$ — here the contraction has made things more complicated.

In general, $A \subset X$ and $A \cong S^1$, then X/A is homotopy-equivalent to the space Y obtained by attaching a 2-cell to X along A. (To see this, let D be the closure of that disc in Y; then D is contractible and $Y/D = X/A$.) So the answer depends strongly on how A sits inside D — specifically, is the

[2] Thanks to Leonard Huang for posing this question.

homeomorphism $S^1 \to A$ homotopic to a constant map? If so, contracting A has the effect of attaching a 2-sphere; if not, we are probably removing a 1-sphere. This last sentence is not a formal statement, but we will actually be able to show, eventually, that X/A is **never** homotopy-equivalent to X.

5. Friday 2/1/14

A loose end from last time: I asserted that contracting a closed loop in $\mathbb{R}^3$ might depend on whether the loop was knotted. I now think this is false. Since $\mathbb{R}^3$ is contractible, the answer to that question is always "yes", so we get $S^2 \vee \mathbb{R}^3$ (which is homotopy-equivalent to just S^2).

5.1. **Back to cell complexes.** Notation: e^n_α is going to denote an n-cell — a homeomorphic copy of $\mathbb{R}^n$. Think of it as the interior of a closed n-disk $D^n_\alpha = \overline{e^n_\alpha}$, so that the boundary ∂D^n_α is an $(n-1)$-sphere.

Definition: A **cell complex** or **CW complex** is a space X constructed as follows.

Let X^0 be a collection of points e^0_α ("vertices"), considered as a space endowed with the discrete topology. This is the **0-skeleton.**

Take a collection of closed line segments $\{D^1_\alpha\}$ and identify each endpoint of each segment with a vertex. In other words, we have *attaching maps* $\varphi_\alpha : \partial D^1_\alpha \to X^0$. The **1-skeleton** is then

$$X^1 = X^0 \cup \coprod_\alpha e^1_\alpha \ / \ \cup \varphi_\alpha.$$

Take a collection of closed 2-discs $\{D^2_\alpha\}$ and a collection of attaching maps $\varphi_\alpha : \partial D^2_\alpha \to X^1$. The **2-skeleton** is then

$$X^2 = X^1 \cup \coprod_\alpha e^2_\alpha \ / \ \cup \varphi_\alpha.$$

Repeat as you like. We then set $X = \bigcup_{n=0}^\infty X^n$. X is d-dimensional if the largest cell attached is of dimension d, or infinite-dimensional if there is no largest cell. In the latter case one has to be careful about how to define the topology; see the Appendix. We'll mostly consider finite-dimensional cell complexes.

Definition: A cell complex is **proper** (my terminology) if the closure of each e_α is a union of cells. This condition is not explicit in Hatcher, but generally seems to be implicit. For example, this rules out the complex obtained by identifying the boundary of a 2-cell with a point that is not a vertex. In general, I will assume that all cell complexes are proper.

Definition: A proper cell complex is **regular** if each attaching map is a homeomorphism. The "obvious" cell structure on S^n obtained by squashing ∂B^n to a point (equivalently, passing from $\mathbb{R}^n$ to its 1-point compactification) is *not* regular for $n > 1$. However, we can construct a regular cell structure on S^n inductively: start with a regular cell structure on S^{n-1}, then take two n-cells and attach their boundaries to S^{n-1} by homeomorphisms.

Why is regularity a good condition? It says that we essentially only need to keep track of the relations between cells of adjacent dimensions.

Every closed cell D^n_α comes with a **characteristic map** $\Phi_\alpha : D^n_\alpha \to X$, which extends the attaching map $\varphi_\alpha : \partial D^n_\alpha \to X^{n-1}$. The characteristic map is nothing fancy; we are just writing down a formal way to identify the closure of a cell. The attaching and characteristic maps often get suppressed: e.g., we would call the closure of a cell $\overline{e^n_\alpha}$ rather than $\overline{\Phi(e^n_\alpha)}$.

5.2. **Real and Complex Projective Spaces. Definition: Real projective space $\mathbb{R}\boldsymbol{P^n}$** is the space of lines through the origin in $\mathbb{R}^{n+1}$.

Note that every nonzero vector $\mathbf{v} = (v_0, \dots, v_n) \in \mathbb{R}^{n+1}$ gives rise to a line $\mathbb{R}\mathbf{v} = \{a\mathbf{v} : \; a \in \mathbb{R}\}$, and that $\mathbb{R}\mathbf{v} = \mathbb{R}\mathbf{w}$ if and only if $\mathbf{v} = \lambda\mathbf{w}$ for some $b \in \mathbb{R} \setminus \{0\}$. Therefore, $\mathbb{R}P^n$ is a topological quotient space:

$$\mathbb{R}P^n = (\mathbb{R}^{n+1} \setminus \{0\})/\mathbf{v} \sim \lambda\mathbf{v}.$$

Each line has a canonical direction vector whose last nonzero coefficient is 1. We can decompose $\mathbb{R}P^n$ into pieces $e^0, \dots, e^n$ depending on the position of that coefficient: defined by

$$e^k = \{\mathbb{R}\mathbf{v} \in \mathbb{R}P^n \mid v_i = 1, \; v_{i+1} = \cdots = v_n = 0\}.$$

Note that $e^k \cong \mathbb{R}^k$ — the coordinates $v_0, v_1, \ldots, v_{k-1}$ provide a homeomorphism. Moreover, the closure of e^k is the subspace $e^0 \cup e^1 \cup \cdots \cup e^n$. To see this, consider the first nontrivial example $k = 1, n = 1$.

In $\mathbb{R}P^1$, the cell e^1 is the set of lines spanned by vectors of the form $(x, 1)$ for $x \in \mathbb{R}$. What happens to these lines as v_0 tends to infinity (from either direction)?

$$\lim_{x \to \pm\infty} \mathbb{R}(x, 1) = \lim_{x \to \pm\infty} \mathbb{R}(1, 1/x) = \mathbb{R}(1, 0).$$

This is just saying that as the slope of a line increases towards infinity, it becomes more and more vertical. So e^0 is in the closure of e^1.

More generally, consider a point $\mathbf{v} = (v_0, \ldots, v_{k-1}, 0, \ldots, 0) \in \mathbb{R}^k \subset \mathbb{R}^n$, which corresponds to the line $\mathbb{R}\mathbf{v} \in e^k$. Let that point waltz off to infinity in some direction $\mathbf{a} \in \mathbb{R}^k$. The limit is in $\overline{e^k}$, and it is

$$\lim_{t \to \pm\infty} \mathbb{R}(v_0 + a_0 t,\ v_1 + a_1 t,\ \ldots,\ v_{k-1} + a_{k-1} t,\ v_k,\ 0,\ \ldots, 0)$$

and the canonical vector of this line will have its final 1 in the j^{th} position, where j is the last nonzero entry of $\mathbf{a}$ — so $0 \leq j < k$. This says that

$$\overline{e^k} = e^k \cup e^{k-1} \cup \cdots \cup e^0$$

and these closures of cells are precisely the CW-subcomplexes of $\mathbb{R}P^n$.

So we have a proper cell complex structure on $\mathbb{R}P^n$. It is not regular, because the attaching maps are all 2-1 (that's the $\pm$). We can now see why this construction of $\mathbb{R}P^2$ agrees with the one we've seen before (take a disk and identify antipodal points on the boundary). Start with a copy of $\mathbb{R}^2$ in which each point $(x, y) \in \mathbb{R}^2$ is identified with the line $\mathbb{R}(x, y, 1)$ (i.e., with $e^2 \subset \mathbb{R}P^2$). Then, for example, walking towards infinity in either direction along the line $y = 2x$ gets us the same line in e^1:

$$\lim_{y=2x \to \infty} \mathbb{R}(x, y, 1) = \lim_{x \to \infty} \mathbb{R}(x, 2x, 1) = \lim_{x \to \infty} \mathbb{R}(1, 2, 1/x) = \mathbb{R}(1, 2, 0).$$

Remark 5.1. More generally, one can put a topological structure on the *Grassmannian* $G(k, \mathbb{R}^n)$, which is the set of k-dimensional vector subspaces of $\mathbb{R}^n$. (Thus $\mathbb{R}P^n = G(1, \mathbb{R}^{n+1})$.) The Grassmannian can be constructed as a quotient space: start with the space of full-rank $n \times k$ matrices over $\mathbb{R}$ (which is an open subset of $\mathbb{R}^{kn}$), then identify matrices with the same column spans. The result is a nice compact space (in fact, a manifold) that has a beautiful combinatorial structure — which is currently outside the scope of these notes, but if you are interested then you should think about how to generalize the cellulation of $\mathbb{R}P^n$.

Definition: Complex projective space $\mathbb{C}P^n$ is constructed in the same way as a quotient space and as a cell complex, replacing $\mathbb{R}$ with $\mathbb{C}$. The difference is that we now obtain cells of even dimensions:

$$\mathbb{C}P^n = e^0 \cup e^2 \cup e^4 \cup \cdots \cup e^{2n}, \qquad \overline{e^{2k}} = \bigcup_{i=0}^{j} e^{2i}.$$

For example, $\mathbb{C}P^2$ is the Riemann sphere.

6. Monday 2/3/14

Definitions: Let X be a cell complex.

Subcomplex of X: closed subspace that is a union of cells.
CW-pair: (X, A) such that X is a cell complex and A is a subcomplex.

The Homotopy Extension Property

Goal: Establish criteria for homotopy equivalence to avoid having to use the definition of homotopy.

1. **Collapsing a contractible subspace.** If (X, A) is a CW-pair and A is contractible, then $X \simeq X/A$.

2. **Homotoping an attaching map.** If (X, A) is a CW-pair and we attach X to Y along A by either of two homotopic maps $f : A \to Y$ and $g : A \to Y$, then

$$Y \sqcup_f X \simeq Y \sqcup_g X.$$

It turns out that both these properties can be proved by studying the more general problem:

When can a homotopy have its domain extended?

I.e., suppose we have spaces $A \subseteq X$, and a homotopy $f_t : A \to Y$. If we can extend f_0 to a function $\overline{f_0} : X \to Y$, can we "stretch f_0 across the homotopy" to get a homotopy $\overline{f_t} : X \to Y$ extending f_t?

Definition: A pair (X, A) has the *homotopy extension property* (HEP) if every homotopy $F : A \to Y$ can be extended to a homotopy $G : X \to Y$.

Equivalently, every continuous map

$$\underbrace{(X \times \{0\}) \cup (A \times I)}_{Z} \xrightarrow{F} Y$$

can be extended to a commutative diagram of continuous maps

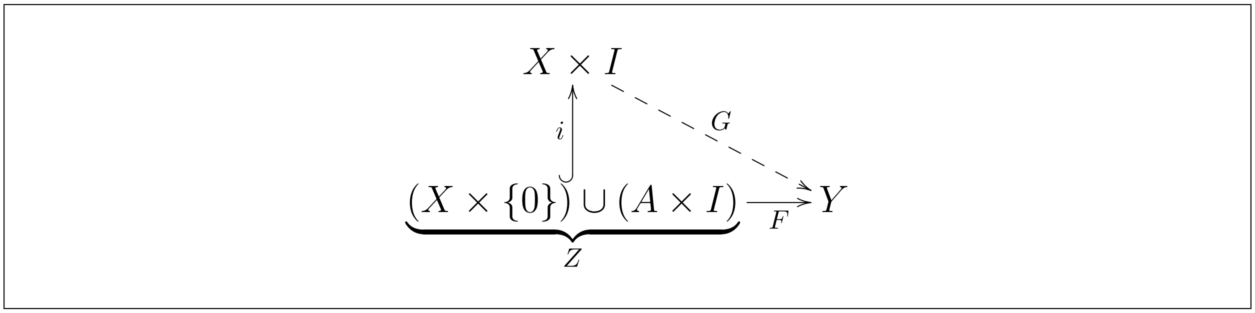

Here i means the inclusion map. To say that the diagram commutes is to say that $G \circ i = F$.

So the HEP means that for every F, there exists G with $G \circ i = F$ in the above diagram.

Note that $Z = X \times \{0\} \cup A \times I \to Y$ is the mapping cylinder M_i of the inclusion $A \hookrightarrow X$.

A good example is the case that X is a disc and A is its boundary circle. The space Z is an empty can with a bottom but no top. The space $X \times I$ is the same can full of soup – a solid cylinder. We want to be able to "fill in the can" by extending any map with domain Z to a map with domain $X \times I$.

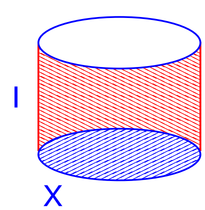

Lemma: (X, A) has the HEP if and only if $Z = X \times \{0\} \cup A \times I \to Y$ is a retract of $X \times I$ — i.e., if there is a continuous map $r : X \times I \to Z$ that fixes Z pointwise.

Proof. ($\Longrightarrow$): Apply the HEP to the identity map $F = \mathbb{1}_Z$. The result is by definition a retraction.

($\Longleftarrow$): Suppose r is a retraction $X \times I \to Z$ and $\varphi : Z \to Y$ is any map. Then $\varphi \circ r$ is a map $X \times I \to Y$, and it extends φ because $(\varphi \circ r)|_Z = \varphi \circ (r|_Z) = \varphi \circ \mathbb{1}_Z = \varphi$. $\square$

Here's how to see this in the case of the can. Fill the can with something mushy like clay. Now pull the sides of the can out a little to get a frustum (in particular, the can should now be the graph of a function $\mathbb{R}^2 \to \mathbb{R}$). Now push the clay straight down. That's certainly a continuous function, and after all the clay has been mashed onto the can we end up with the can itself. This is the desired retraction $X \times I \to Z$. An analogous construction would work for $X = D^n$ and $A = \partial D^n$ (this is the case $n = 2$).

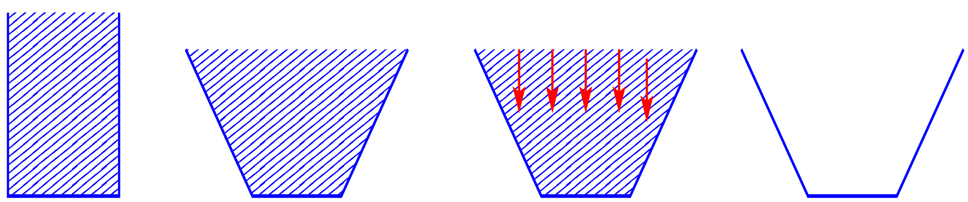

Prop 0.16: If (X, A) is a CW pair with inclusion $i : A \to X$, then the mapping cylinder $Z = M_i$ is a <u>deformation</u> retract of $X \times I$. In particular, the pair (X, A) has the HEP by the Lemma.

Proof. First consider the case that A is the complement in X of a single maximal open n-cell e^n. We already know (by the "frustum argument") that there is a deformation retraction

$$D^n \times I \to D^n \times \{0\} \cup \partial D^n \times I.$$

Identifying $D^n \times I$ with $\overline{e^n} \times I$ by the characteristic map and fixing $A \times I$ pointwise gives a deformation retraction

$$X \times I \to X \times \{0\} \cup A \times I$$

which fixes $A \times I$ pointwise. So (X, A) has the HEP.

Here we are implicitly using the facts that

(1) one can sew continuous maps together: we have two continuous maps on $A \times I$ and $D^n \times I$ that agree on their intersection $\partial D^n \times I$, so they together define a continuous map on the union $X \times I$; and
(2) since e^n is maximal, it is not in the closure of any other cell, so there is no problem extending the deformation retraction.

More generally, we can compose these deformation retractions to shrink one cell of X at a time, from the top dimension down, to deformation-retract $X \times I$ onto the mapping cylinder of any subcomplex. When X is finite (or more generally when $X \setminus A$ has finitely many cells), this is clearly no problem. A small trick (see Hatcher) is needed to verify that this "infinite composition" is possible when X is infinite-dimensional. $\square$

7. Wednesday 2/5/14

From last time:

Definition: A pair (X, A) has the *homotopy extension property* (HEP) if every homotopy $F : A \to Y$ can be extended to a homotopy $G : X \to Y$.

Equivalently, every map $X \times \{0\} \cup A \times I \to Y$ can be extended to a map $X \times I \to Y$.

Every CW-pair (X, A) satisfies the HEP.

Example: Hatcher's example of a pair that does **not** satisfy the HEP: $X = I$, $A = \{0, 1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots\}$. Here $X \times I = I \times I$ is the closed unit square, and the mapping cylinder Z is the "comb space":

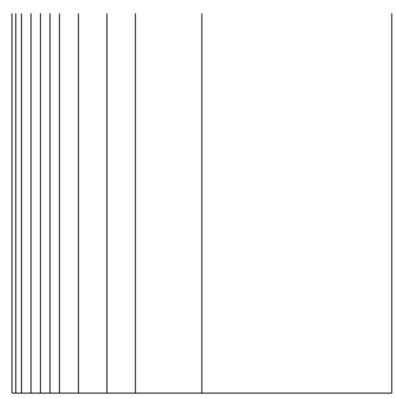

Here the horizontal line segment is $X \times \{0\}$, and the teeth of the comb are $A \times I$. There exists no retraction $X \times I \to Z$ (despite the fact that Z is contractible).

Prop 0.17: If (X, A) satisfies the HEP and A is contractible, then $X \simeq X/A$ via the quotient map $q : X \to X/A$.

Note that it's not clear how to construct a homotopy inverse directly. For example, suppose X is a closed line segment and A is a proper closed subsegment; say $X = [0, 3] \subset \mathbb{R}$ and $A = [1, 2] \subset X$. We know that A is contractible and it is pretty clear that $X \simeq X/A$ (in fact they are homeomorphic). But the quotient map will *not* have a homotopy inverse that fixes all the points in $X \setminus A$, because there would be no way to extend such a map so that is continuous on A.

Proof. Start with a contraction of A, i.e., a homotopy between $\mathbb{1}_A$ and a constant map sending all of A to some fixed point $a \in A$. Since $\mathbb{1}_A$ extends to $\mathbb{1}_X$, we can extend the contraction to a homotopy $f_t : X \to X$, where $f_0 = \mathbb{1}_X$. Note that $f_t(A) \subset A$ for all A. Therefore there is a homotopy $\overline{f_t} : X/A \to X/A$ such that $q \circ f_t = \overline{f_t} \circ q$. (We just define $\overline{f_t}|_{X \setminus A} = f_t|_{X \setminus A}$.) In other words, the following diagram **commutes**.

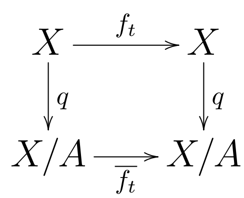

When $t = 1$, we have $f_1(A) = \{a\}$. So there is a well-defined map $g : X/A \to X$ such that $g \circ q = f_1$. Specifically,

$$g(x) = \begin{cases} a & \text{if } x \in A, \\ f_1(x) & \text{if } x \notin A. \end{cases}$$

Stick that in the commutative diagram:

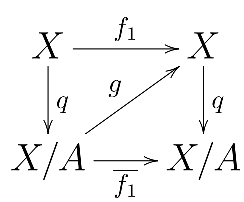

Now $g \circ q = f_1 \simeq f_0 = \mathbb{1}_X$ and $q \circ g = \overline{f_1} \simeq \overline{f_0} = \mathbb{1}_{X/A}$. Therefore g and q are homotopy inverses. $\square$

Prop 0.18: Let (X, A) be a CW pair and let $f, g : A \to Y$ such that $f \simeq g$. Then $X \sqcup_f Y \simeq X \sqcup_g Y$.

Proof. Let $F : A \times I \to Y$ be a homotopy of f with g. By what we have proved about CW pairs, there is a deformation retraction of $X \times I$ onto $M = X \times \{0\} \cup A \times I$.

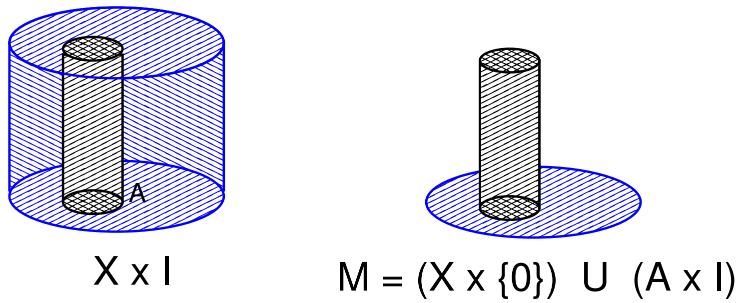

If we attach Y to these spaces via F, we get a deformation retraction of $(X \times I) \sqcup_F Y$ onto $M \sqcup_F Y = (X \times \{0\}) \sqcup_F (A \times I)$, which in turn deformation-retracts onto $X \sqcup_f Y$ by shrinking I down to 0.

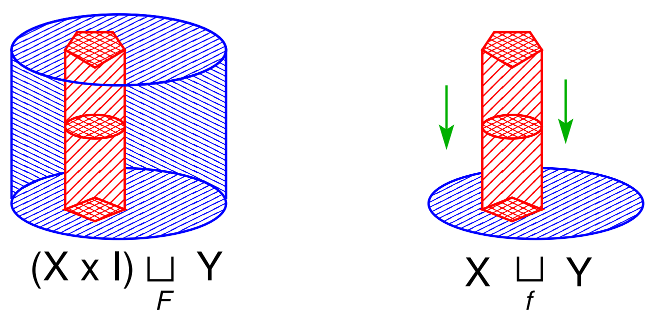

Switching the roles of f and g, we can do the same thing for $X \sqcup_g Y$. ☐

This actually proves something slightly stronger, namely that $X \sqcup_f Y$ and $X \sqcup_g Y$ are homotopic *relative to* Y, i.e., via a homotopy h_t such that $h_t|_Y = \mathbb{1}_Y$ for all t. Notation:

$$X \sqcup_f Y \simeq X \sqcup_g Y \quad \text{rel } Y.$$

Example: : Hatcher p.19 #19. Show that the space obtained by attaching n 2-cells along any collection of circles in S^2 is homotopy-equivalent to the wedge sum of $n+1$ 2-spheres.

Start with one circle.

Lots of disjoint circles.

What if the circles overlap? What if they coincide?

Example: Hatcher p.19 #21. Let X be a connected space that is the union of a finite number of 2-spheres, any two of which intersect in at most one point. Show that X is homotopy equivalent to a wedge sum of S^1's and S^2's.

What if any two spheres intersect in at most a finite number of points?