This is sort of a "uniqueness statement" for a lift. If we know a point in the lift, then we know the whole path. This is since once we've decided our starting point, i.e. which "copy" of X we work in, the rest of $\tilde{f}$ has to follow what f does.

Proof. First we show it is open. Let y be such that $\tilde{f}_1(y) = \tilde{f}_2(y)$. Then there is an evenly covered open neighbourhood $U \subseteq X$ of $f(y)$. Let $\tilde{U}$ be such that $\tilde{f}_1(y) \in \tilde{U}$, $p(\tilde{U}) = U$ and $p|_{\tilde{U}} : \tilde{U} \to U$ is a homeomorphism. Let $V = \tilde{f}_1^{-1}(\tilde{U}) \cap \tilde{f}_2^{-1}(\tilde{U})$. We will show that $\tilde{f}_1 = \tilde{f}_2$ on V.

Indeed, by construction

$$p|_{\tilde{U}} \circ \tilde{f}_1|_V = p|_{\tilde{U}} \circ \tilde{f}_2|_V.$$

Since $p|_{\tilde{U}}$ is a homeomorphism, it follows that

$$\tilde{f}_1|_V = \tilde{f}_2|_V.$$

Now we show S is closed. Suppose not. Then there is some $y \in \bar{S} \setminus S$. So $\tilde{f}_1(y) \neq \tilde{f}_2(y)$. Let U be an evenly covered neighbourhood of $f(y)$. Let $p^{-1}(U) = \coprod U_\alpha$. Let $\tilde{f}_1(y) \in U_\beta$ and $\tilde{f}_2(y) \in U_\gamma$, where $\beta \neq \gamma$. Then $V = \tilde{f}_1^{-1}(U_\beta) \cap \tilde{f}_2^{-1}(U_\gamma)$ is an open neighbourhood of y, and hence intersects S by definition of closure. So there is some $x \in V$ such that $\tilde{f}_1(x) = \tilde{f}_2(x)$. But $\tilde{f}_1(x) \in U_\beta$ and $\tilde{f}_2(x) \in U_\gamma$, and hence U_β and U_γ have a non-trivial intersection. This is a contradiction. So S is closed. $\square$

We just had a uniqueness statement. How about existence? Given a map, is there guarantee that we can lift it to something? Moreover, if I have fixed a "copy" of X I like, can I also lift my map to that copy? We will later come up with a general criterion for when lifts exist. However, it turns out homotopies can always be lifted.

Lemma (Homotopy lifting lemma)**.** Let $p : \tilde{X} \to X$ be a covering space, $H : Y \times I \to X$ be a homotopy from f_0 to f_1. Let $\tilde{f}_0$ be a lift of f_0. Then there exists a *unique* homotopy $\tilde{H} : Y \times I \to \tilde{X}$ such that

(i) $\tilde{H}(\,\cdot\,, 0) = \tilde{f}_0$; and

(ii) $\tilde{H}$ is a lift of H, i.e. $p \circ \tilde{H} = H$.

This lemma might be difficult to comprehend at first. We can look at the special case where $Y = *$. Then a homotopy is just a path. So the lemma specializes to

Lemma (Path lifting lemma)**.** Let $p : \tilde{X} \to X$ be a covering space, $\gamma : I \to X$ a path, and $\tilde{x}_0 \in \tilde{X}$ such that $p(\tilde{x}_0) = x_0 = \gamma(0)$. Then there exists a *unique* path $\tilde{\gamma} : I \to \tilde{X}$ such that

(i) $\tilde{\gamma}(0) = \tilde{x}_0$; and

(ii) $\tilde{\gamma}$ is a lift of γ, i.e. $p \circ \tilde{\gamma} = \gamma$.

This is exactly the picture we were drawing before. We just have to start at a point $\tilde{x}_0$, and then everything is determined because locally, everything upstairs in $\tilde{X}$ is just like X. Note that we have already proved uniqueness. So we just need to prove existence.

Note that if we cover the words "as paths" and just talk about homotopies, then this is just the homotopy lifting lemma. So we can view this as a stronger form of the homotopy lifting lemma.

Proof. The homotopy lifting lemma gives us an $\tilde{H}$, a lift of H with $\tilde{H}(\,\cdot\,, 0) = \tilde{\gamma}$.

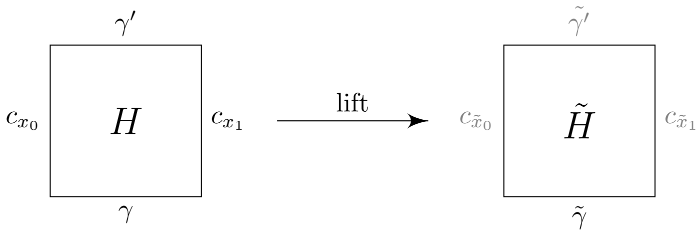

In this diagram, we by assumption know the bottom of the $\tilde{H}$ square is $\tilde{\gamma}$. To show that this is a path homotopy from $\tilde{\gamma}$ to $\tilde{\gamma}'$, we need to show that the other edges are $c_{\tilde{x}_0}$, $c_{\tilde{x}_1}$ and $\tilde{\gamma}'$ respectively.

Now $\tilde{H}(\,\cdot\,, 1)$ is a lift of $H(\,\cdot\,, 1) = \gamma'$, starting at $\tilde{x}_0$. Since lifts are unique, we must have $\tilde{H}(\,\cdot\,, 1) = \tilde{\gamma}'$. So this is indeed a homotopy between $\tilde{\gamma}$ and $\tilde{\gamma}'$. Now we need to check that this is a homotopy of paths.

We know that $\tilde{H}(0, \cdot\,)$ is a lift of $H(0, \cdot\,) = c_{x_0}$. We are aware of one lift of c_{x_0}, namely $c_{\tilde{x}_0}$. By uniqueness of lifts, we must have $\tilde{H}(0, \cdot\,) = c_{\tilde{x}_0}$. Similarly, $\tilde{H}(1, \cdot\,) = c_{\tilde{x}_1}$. So this is a homotopy of paths. □

So far, our picture of covering spaces is like this:

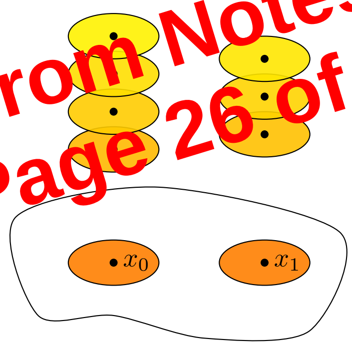

Except... is it? Is it possible that we have four copies of x_0 but just three copies of x_1? This is obviously possible if X is not path connected — the component containing x_0 and the one containing x_1 are completely unrelated. But what if X is path connected?

Corollary. If X is a path connected space, $x_0, x_1 \in X$, then there is a bijection $p^{-1}(x_0) \to p^{-1}(x_1)$.

Proof. Let $\gamma : x_0 \rightsquigarrow x_1$ be a path. We want to use this to construct a bijection between each preimage of x_0 and each preimage of x_1. The obvious thing to do is to use lifts of the path γ.

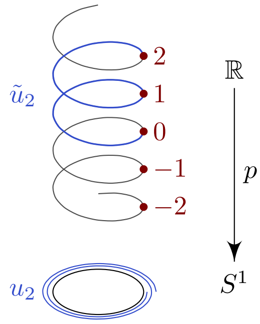

Let $\tilde{u}_n : I \to \mathbb{R}$ be defined by $t \mapsto nt$, and let $u_n = p \circ \tilde{u}_n$. Since $\mathbb{R}$ is simply connected, there is a unique homotopy class between any two points. So for any $[\gamma] \in \pi_1(S^1, 1)$, if $\tilde{\gamma}$ is the lift to $\mathbb{R}$ at 0 and $\tilde{\gamma}(1) = n$, then $\tilde{\gamma} \simeq \tilde{u_n}$ as paths. So $[\gamma] = [u_n]$.

To show that this has the right group operation, we can easily see that $\widetilde{u_m \cdot u_n} = \tilde{u}_{m+n}$, since we are just moving by $n + m$ in both cases. Therefore

$$\ell([u_m][u_n]) = \ell([u_m \cdot u_m]) = m + n = \ell([u_{m+n}]).$$

So ℓ is a group isomorphism. □

What have we done? In general, we might be given a horrible, crazy loop in S^1. It would be rather difficult to work with it directly in S^1. So we pull it up to the universal covering $\mathbb{R}$. Since $\mathbb{R}$ is nice and simply connected, we can easily produce a homotopy that "straightens out" the path. We then project this homotopy down to S^1, to get a homotopy from γ to u_n.

It is indeed possible to produce a homotopy directly inside S^1 from each loop to some u_n, but that would be tedious work that involves messing with a lot of algebra and weird convoluted formulas.

With the fundamental group of the circle, we do many things. An immediate application is that we can properly define the "winding number" of a closed curve. Since $\mathbb{C} \setminus \{0\}$ is homotopy equivalent to S^1, its fundamental group is $\mathbb{Z}$ as well. Any closed curve $S^1 \to \mathbb{C} \setminus \{0\}$ thus induces a group homomorphism $\mathbb{Z} \to \mathbb{Z}$. Any such group homomorphism must be of the form $t \mapsto nt$, and the winding number is given by n. If we stare at it long enough, it is clear that this is exactly the number of times the curve winds around the origin.

Also, we have the following classic application:

Theorem (Brouwer's fixed point theorem)**.** Let $D^2 = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$ be the unit disk. If $f : D^2 \to D^2$ is continuous, then there is some $x \in D^2$ such that $f(x) = x$.

Proof. Suppose not. So $x \neq f(x)$ for all $x \in D^2$.

Then what does it *feel* like to live in the torus? If you live in a torus and look around, you don't see a boundary. The space just extends indefinitely for ever, somewhat like $\mathbb{R}^2$. The difference is that in the torus, you aren't actually seeing free space out there, but just seeing copies of the same space over and over again. If you live inside the square, the universe actually looks like this:

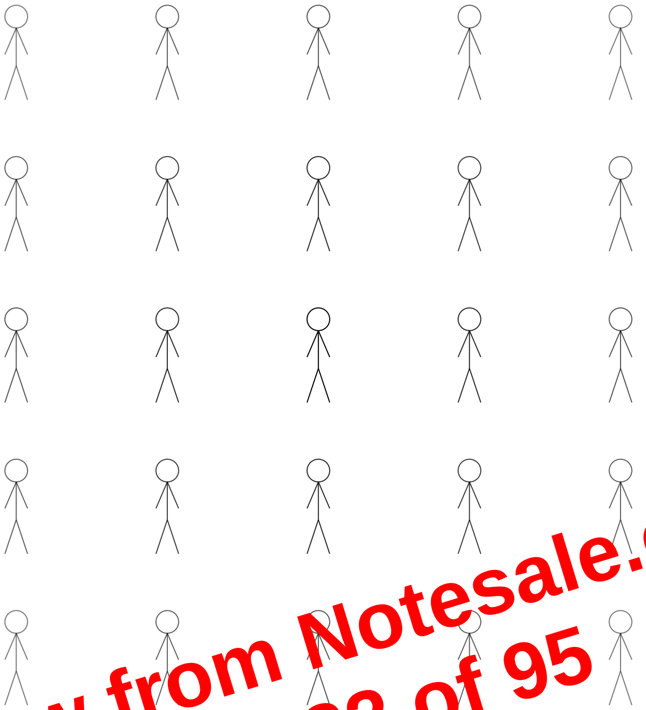

As we said, this looks somewhat like $\mathbb{R}^2$, but we *know* that this is not $\mathbb{R}^2$, since we can see some symmetry in this space. Whenever we move one unit horizontally or vertically, we get back to "the same place". In fact, we can move horizontally by n units and vertically by m units, for any $n, m \in \mathbb{Z}$, and still get back to the same place. This space has a huge translation symmetry. What is this symmetry? It is exactly $\mathbb{Z} \times \mathbb{Z}$.

We see that if we live inside the torus $S^1 \times S^1$, it feels like we are actually living in the universal covering space $\mathbb{R} \times \mathbb{R}$, except that we have an additional symmetry given by the fundamental group $\mathbb{Z} \times \mathbb{Z}$.

Hopefully, you are convinced that universal covers are nice. We would like to say that universal covers always exist. However, this is not always true.

Firstly, we should think — what would having a universal cover imply? Suppose X has a universal cover $\tilde{X}$. Pick any point $x_0 \in X$, and pick an evenly covered neighbourhood U in X. This lifts to some $\tilde{U} \subseteq \tilde{X}$. If we draw a teeny-tiny loop γ around x_0 inside U, we can lift this γ to $\tilde{\gamma}$ in $\tilde{U}$. But we know that $\tilde{X}$ is simply connected. So $\tilde{\gamma}$ is homotopic to the constant path. Hence γ is also homotopic to the constant path. So all loops (contained in U) at x_0 are homotopic to the constant path.

It seems like for every $x_0 \in X$, there is some neighbourhood of x_0 that is simply connected. Except that's not what we just showed above. The homotopy from $\tilde{\gamma}$ to the constant path is a homotopy in $\tilde{X}$, and can pass through anything

Proof. Since X is a path connected, locally path connected and semi-locally simply connected space, let $\bar{X}$ be a universal covering. We have an intermediate group H such that $\pi_1(\tilde{X}, \tilde{x}_0) = 1 \leq H \leq \pi_1(X, x_0)$. How can we obtain a corresponding covering space?

Note that if we have $\bar{X}$ and we want to recover X, we can quotient $\bar{X}$ by the action of $\pi_1(X, x_0)$. Since $\pi_1(X, x_0)$ acts on $\bar{X}$, so does $H \leq \pi_1(X, x_0)$. Now we can define our covering space by taking quotients. We define $\sim_H$ on $\bar{X}$ to be the orbit relation for the action of H, i.e. $\tilde{x} \sim_H \tilde{y}$ if there is some $h \in H$ such that $\tilde{y} = h\tilde{x}$. We then let $\tilde{X}$ be the quotient space $\bar{X}/{\sim_H}$.

We can now do the messy algebra to show that this is the covering space we want. $\square$

We have just showed that every subgroup comes from some covering space, i.e. the map from the set of covering spaces to the subgroups of π_1 is surjective. Now we want to prove injectivity. To do so, we need a generalization of the homotopy lifting lemma.

Suppose we have path-connected spaces (Y, y_0), (X, x_0) and $(\tilde{X}, \tilde{x}_0)$, with $f : (Y, y_0) \to (X, x_0)$ a continuous map, $p : (\tilde{X}, \tilde{x}_0) \to (X, x_0)$ a covering map. When does a lift of f to $\tilde{f} : (Y, y_0) \to (\tilde{X}, \tilde{x}_0)$ exist? The answer is given by the lifting criterion.

Lemma (Lifting criterion)**.** Let $p : (\tilde{X}, \tilde{x}_0) \to (X, x_0)$ be a covering map of path-connected based spaces, and (Y, y_0) a path-connected, locally path connected based space. If $f : (Y, y_0) \to (X, x_0)$ is a continuous map, then there is a (unique) lift $\tilde{f} : (Y, y_0) \to (\tilde{X}, \tilde{x}_0)$ such that the diagram below commutes (i.e. $p \circ \tilde{f} = f$):

$$\begin{tikzcd}
& (\tilde{X}, \tilde{x}_0) \arrow[d, "p"] \\
(Y, y_0) \arrow[r, "f"'] \arrow[ur, dashed, "\tilde{f}"] & (X, x_0)
\end{tikzcd}$$

if and only if the following condition holds:

$$f_*\pi_1(Y, y_0) \leq p_*\pi_1(\tilde{X}, \tilde{x}_0).$$

Note that uniqueness comes from the uniqueness of lifts. So this lemma is really about existence.

Also, note that the condition holds trivially when Y is simply connected, e.g. when it is an interval (path lifting) or a square (homotopy lifting). So paths and homotopies can always be lifted.

Proof. One direction is easy: if $\tilde{f}$ exists, then $f = p \circ \tilde{f}$. So $f_* = p_* \circ \tilde{f}_*$. So we know that $\operatorname{im} f_* \subseteq \operatorname{im} p_*$. So done.

In the other direction, uniqueness follows from the uniqueness of lifts. So we only need to prove existence. We define $\tilde{f}$ as follows:

Given a $y \in Y$, there is some path $\alpha_y : y_0 \rightsquigarrow y$. Then f maps this to $\beta_y : x_0 \rightsquigarrow f(y)$ in X. By path lifting, this path lifts uniquely to $\tilde{\beta}_y$ in $\tilde{X}$. Then we set $\tilde{f}(y) = \tilde{\beta}_y(1)$. Note that if $\tilde{f}$ exists, then this *must* be what $\tilde{f}$ sends y to. What we need to show is that this is well-defined.

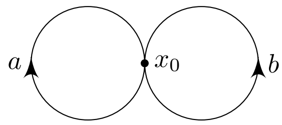

We call this a "rose with 2 petals". This is a cell complex, with one 0-cell and $|S|$ 1-cells. For each $s \in S$, we have one 1-cell, e_s, and we fix a path $\gamma_s : [0, 1] \to e_s$ that goes around the 1-cell once. We will call the 0-cells and 1-cells *vertices* and *edges*, and call the whole thing a *graph*.

What's the universal cover of X? Since we are just lifting a 1-complex, the result should be a 1-complex, i.e. a graph. Moreover, this graph is connected and simply connected, i.e. it's a tree. We also know that every vertex in the universal cover is a copy of the vertex in our original graph. So it must have 4 edges attached to it. So it has to look like something this:

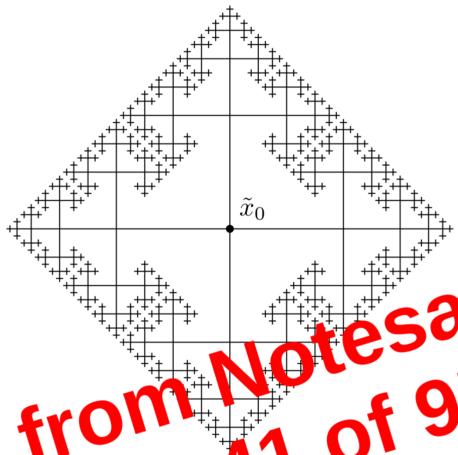

In X, we know that at each vertex, there should be an edge labeled a going in; an edge labeled a going out; an edge labeled b going in; an edge labeled b going out. This should be the case in $\tilde{X}$ as well. So $\tilde{X}$ looks like this:

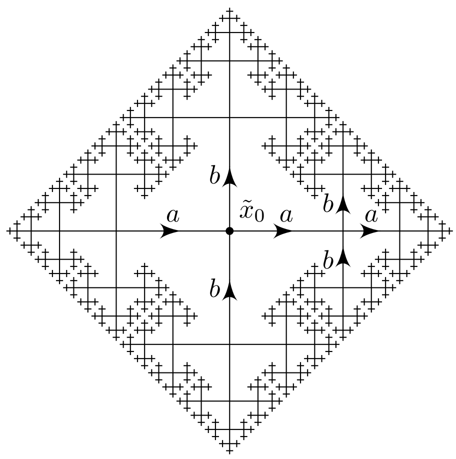

The projection map is then obvious — we send all the vertices in $\tilde{X}$ to $x_0 \in X$, and then the edges according to the labels they have, in a way that respects the direction of the arrow. It is easy to show this is really a covering map.

5 Seifert-van Kampen theorem

5.1 Seifert-van Kampen theorem

The Seifert-van Kampen theorem is the theorem that tells us what happens when we glue spaces together.

Here we let $X = A \cup B$, where $A, B, A \cap B$ are path-connected.

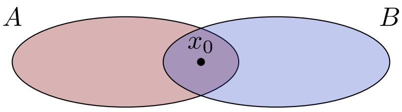

We pick a basepoint $x_0 \in A \cap B$ for convenience. Since we like diagrams, we can write this as a commutative diagram:

$$\begin{tikzcd} A \cap B \arrow[r, hook] \arrow[d, hook] & B \arrow[d, hook] \\ A \arrow[r, hook] & X \end{tikzcd}$$

where all arrows are inclusion (i.e. injective) maps. We can consider what happens when we take the fundamental groups of each space. Then we have the induced homomorphisms

$$\begin{tikzcd} \pi_1(A \cap B, x_0) \arrow[r] \arrow[d] & \pi_1(B, x_0) \arrow[d] \\ \pi_1(A, x_0) \arrow[r] & \pi_1(X, x_0) \end{tikzcd}$$

We might guess that $\pi_1(X, x_0)$ is just the free product with amalgamation

$$\pi_1(X, x_0) = \pi_1(A, x_0) \underset{\pi_1(A \cap B, x_0)}{*} \pi_1(B, x_0).$$

The Seifert-van Kampen theorem says that, under mild hypotheses, this guess is correct.

Theorem (Seifert-van Kampen theorem)**.** Let A, B be open subspaces of X such that $X = A \cup B$, and $A, B, A \cap B$ are path-connected. Then for any $x_0 \in A \cap B$, we have

$$\pi_1(X, x_0) = \pi_1(A, x_0) \underset{\pi_1(A \cap B, x_0)}{*} \pi_1(B, x_0).$$

Note that by the universal property of the free product with amalgamation, we by definition know that there is a unique map $\pi_1(A, x_0) \underset{\pi_1(A \cap B, x_0)}{*} \pi_1(B, x_0) \to \pi_1(X, x_0)$. The theorem asserts that *this* map is an isomorphism.

Proof is omitted because time is short.

Example. Consider a higher-dimensional sphere $S^n = \{\mathbf{v} \in \mathbb{R}^{n+1} : |\mathbf{v}| = 1\}$ for $n \geq 2$. We want to find $\pi_1(S^n)$.

The idea is to write S^n as a union of two open sets. We let $n = \mathbf{e}_1 \in S^n \subseteq \mathbb{R}^{n+1}$ be the North pole, and $s = -\mathbf{e}_1$ be the South pole. We let $A = S^n \setminus \{n\}$, If we just wanted a torus, we are done (after closing the loop), but now we want a surface with genus 2, so we add another torus:

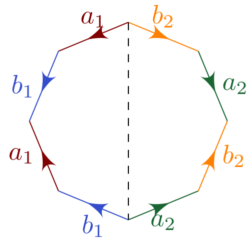

To visualize how this works, imagine cutting this apart along the dashed line. This would give two tori with a hole, where the boundary of the holes are just the dashed line. Then gluing back the dashed lines would give back our orientable surface with genus 2.

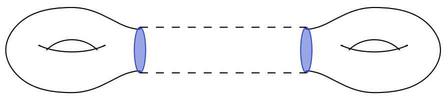

In general, to produce Σ_g, we produce a polygon with $4g$ sides. Then we get

$$\pi_1\Sigma_g = \langle a_1, b_1, \cdots, a_g, b_g \mid a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} \rangle.$$

We do we care? The classification theorem tells us that each surface is homeomorphic to *some* of these orientable and non-orientable surfaces, but it doesn't tell us there is no overlap. It might be that $\Sigma_6 \cong \Sigma_{241}$, via some weird homeomorphism that destroys some holes.

However, this result lets us know that all these orientable surfaces are genuinely different. While it is difficult to stare at this fundamental group and say that $\pi_1\Sigma_g \not\cong \pi_1\Sigma_{g'}$ for $g \neq g'$, we can perform a little trick. We can take the abelianization of the group $\pi_1\Sigma_g$, where we further quotient by all commutators. Then the abelianized fundamental group of Σ_g will simply be $\mathbb{Z}^{2g}$. These are clearly distinct for different values of g. So all these surfaces are distinct. Moreover, they are not even homotopy equivalent.

The fundamental groups of the non-orientable surfaces is left as an exercise for the reader.

The relevance is that these can be used to define simplices (which are simple, as opposed to complexes).

Definition (n-simplex)**.** An *n-simplex* is the convex hull of $(n+1)$ affinely independent points $a_0, \cdots, a_n \in \mathbb{R}^m$, i.e. the set

$$\sigma = \langle a_0, \cdots, a_n \rangle = \left\{ \sum_{i=0}^{n} t_i a_i : \sum_{i=0}^{n} t_i = 1, t_i \geq 0 \right\}.$$

The points $a_0, \cdots, a_n$ are the *vertices*, and are said to *span* σ. The $(n+1)$-tuple $(t_0, \cdots, t_n)$ is called the *barycentric coordinates* for the point $\sum t_i a_i$.

Example. When $n = 0$, then our 0-simplex is just a point:

When $n = 1$, then we get a line:

When $n = 2$, we get a triangle:

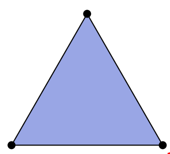

When $n = 3$, we get a tetrahedron:

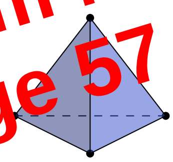

The key motivation of this is that simplices are determined by their vertices. Unlike arbitrary subspaces of $\mathbb{R}^n$, they can be specified by a finite amount of data. We can also easily extract the faces of the simplices.

Definition (Face, boundary and interior)**.** A *face* of a simplex is a subset (or subsimplex) spanned by a subset of the vertices. The *boundary* is the union of the proper faces, and the *interior* is the complement of the boundary.

The boundary of σ is usually denoted by $\partial\sigma$, while the interior is denoted by $\mathring{\sigma}$, and we write $\tau \leq \sigma$ when τ is a face of σ.

In particular, the interior of a vertex is the vertex itself. Note that this notions of interior and boundary are distinct from the topological notions of interior and boundary.

Example. The *standard n-simplex* is spanned by the basis vectors $\{\mathbf{e}_0, \cdots, \mathbf{e}_n\}$ in $\mathbb{R}^{n+1}$. For example, when $n = 2$, we get the following:

So g is a simplicial approximation of f.

The last part follows from the observation that if f is a simplicial map, then it maps vertices to vertices. So we can pick $g(v) = f(v)$. $\square$

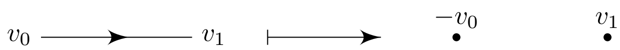

If we take a triangle, we get

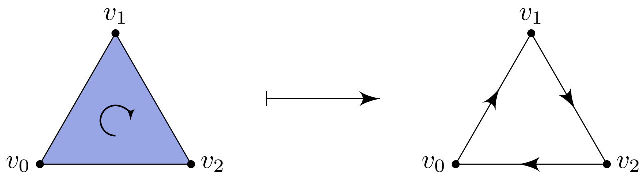

An important property of the boundary map is that the boundary of a boundary is empty:

Lemma. $d_{n-1} \circ d_n = 0$.

In other words, $\operatorname{im} d_{n+1} \subseteq \ker d_n$.

Proof. This just involves expanding the definition and working through the mess. □

With this in mind, we will define the homology groups as follows:

Definition (Simplicial homology group $H_n(K)$)**.** The *nth simplicial homology group* $H_n(K)$ is defined as

$$H_n(K) = \frac{\ker d_n}{\operatorname{im} d_{n+1}}.$$

This is a nice, clean definition, but what does this mean geometrically?

Somehow, $H_k(K)$ describes all the "k-dimensional holes" in $|K|$. Since we are going to draw pictures, we are going to start with the easy case of $k = 1$. Our $H_1(K)$ is made from the kernel of d_1 and the image of d_2. First, we give these things names.

Definition (Chains, cycles and boundaries)**.** The elements of $C_k(K)$ are called *k-chains* of K, those of $\ker d_k$ are called *k-cycles* of K, and those of $\operatorname{im} d_{k+1}$ are called *k-boundaries* of K.

Suppose we have some $c \in \ker d_k$. In other words, $dc = 0$. If we interpret c as a "path", if it has no boundary, then it represents some sort of loop, i.e. a cycle. For example, if we have the following cycle:

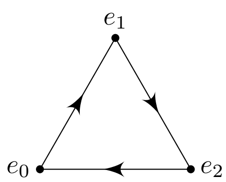

We have

$$c = (e_0, e_1) + (e_1, e_2) + (e_2, e_0).$$

We can then compute the boundary as

$$dc = (e_1 - e_0) + (e_2 - e_1) + (e_0 - e_2) = 0.$$

We have now everything we need to know about the homology groups, and we just need to do some linear algebra to figure out the image and kernel, and thus the homology groups. We have

$$H_0(K) = \frac{\ker(d_0 : C_0(K) \to C_{-1}(K))}{\mathrm{im}(d_1 : C_1(K) \to C_0(K))} \cong \frac{C_0(K)}{\mathrm{im}\, d_1} \cong \frac{\mathbb{Z}^3}{\mathrm{im}\, d_1}.$$

After doing some row operations with our matrix, we see that the image of d_1 is a two-dimensional subspace generated by the image of two of the edges. Hence we have

$$H_0(K) = \mathbb{Z}.$$

What does this $H_0(K)$ represent? We initially said that $H_k(K)$ should represent the k-dimensional holes, but when $k = 0$, this is simpler. As for π_0, H_0 just represents the path components of K. We interpret this to mean K has one path component. In general, if K has r path components, then we expect $H_0(K)$ to be $\mathbb{Z}^r$.

Similarly, we have

$$H_1(K) = \frac{\ker d_1}{\mathrm{im}\, d_2} \cong \ker d_1.$$

It is easy to see that in fact we have

$$\ker d_1 = \langle (e_0, e_1) + (e_1, e_2) + (e_2, e_0) \rangle \cong \mathbb{Z}.$$

So we also have

$$H_1(K) \cong \mathbb{Z}.$$

We see that this $H_1(K)$ is generated by precisely the single loop in the triangle. The fact that $H_1(K)$ is non-trivial means that we do indeed have a hole in the middle of the circle.

Example. Let L be the standard 2-simplex (and all its faces) in $\mathbb{R}^3$.

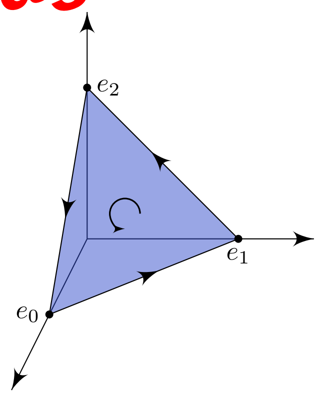

Now our chain groups are

$$\begin{aligned} C_0(L) &= C_0(K) \cong \mathbb{Z}^3 \cong \langle (e_0), (e_1), (e_2) \rangle \\ C_1(L) &= C_1(K) \cong \mathbb{Z}^3 \cong \langle (e_0, e_1), (e_1, e_2), (e_2, e_0) \rangle \\ C_2(L) &\cong \mathbb{Z} = \langle (e_0, e_1, e_2) \rangle. \end{aligned}$$

We will learn how to compute the homology of the union $M \cup N$ in terms of those of M, N and $M \cap N$.

Recall that to state the Seifert-van Kampen theorem, we needed to learn some new group-theoretic notions, such as free products with amalgamation. The situation is somewhat similar here. We will need to learn some algebra in order to state the Mayer-Vietoris theorem. The objects we need are known as *exact sequences*.

Definition (Exact sequence)**.** A pair of homomorphisms of abelian groups

$$A \xrightarrow{f} B \xrightarrow{g} C$$

is *exact (at B)* if

$$\operatorname{im} f = \ker g.$$

A collection of homomorphisms

$$\cdots \xrightarrow{f_{i-1}} A_i \xrightarrow{f_i} A_{i+1} \xrightarrow{f_{i+1}} A_{i+2} \xrightarrow{f_{i+2}} \cdots$$

is *exact at A_i* if

$$\ker f_i = \operatorname{im} f_{i-1}.$$

We say it is *exact* if it is exact at every A_i.

Recall that we have seen something similar before. When we defined the chain complexes, we had $d^2 = 0$, i.e. $\operatorname{im} d \subseteq \ker d$. Here we are requiring exact equivalence, which is something even better.

Algebraically, we can think of an exact sequence as chain complexes with trivial homology groups. Alternatively, we see the homology groups as measuring the failure of a sequence to be exact.

There is a particular type of exact sequences that is important.

Definition (Short exact sequence)**.** A *short exact sequence* is an exact sequence of the form

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

What does this mean?

- The kernel of f is equal to the image of the zero map, i.e. $\{0\}$. So f is injective.
- The image of g is the kernel of the zero map, which is everything. So g is surjective.
- $\operatorname{im} f = \ker g$.

Since we like chain complexes, we can produce short exact sequences of chain complexes.

Definition (Short exact sequence of chain complexes)**.** A *short exact sequence of chain complexes* is a pair of *chain* maps $i_{\bullet}$ and $j_{\bullet}$

$$0 \longrightarrow A_{\bullet} \xrightarrow{i_{\bullet}} B_{\bullet} \xrightarrow{j_{\bullet}} C_{\bullet} \longrightarrow 0$$

By our previous calculation, we know a_* is a map $a_* : \mathbb{Z} \to \mathbb{Z}$. If a is homotopic to the identity, then a_* should be homotopic to the identity map. We will now compute a_* and show that it is multiplication by -1 when n is even.

To do this, we want to use a triangulation that is compatible with the antipodal map. The standard triangulation clearly doesn't work. Instead, we use the following triangulation $h : |K| \to S^n$:

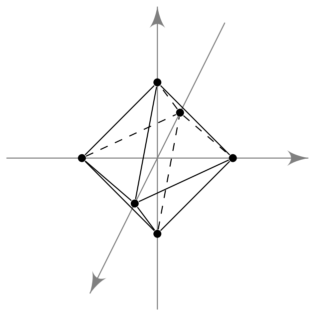

The vertices of K are given by

$$V_K = \{\pm \mathbf{e}_0, \pm \mathbf{e}_1, \cdots, \pm \mathbf{e}_n\}.$$

This triangulation works nicely with the antipodal map, since this maps a vertex to a vertex. To understand the homology group better, we need the following lemma:

Lemma. In the triangulation of S^n given by vertices $V_K = \{\pm \mathbf{e}_0, \pm \mathbf{e}_1, \cdots, \pm \mathbf{e}_n\}$, the element

$$x = \sum_{\varepsilon \in \{\pm 1\}^{n+1}} \varepsilon_0 \cdots \varepsilon_n (\varepsilon_0 \mathbf{e}_0, \cdots, \varepsilon_n \mathbf{e}_n)$$

is a cycle and generates $H_n(S^n)$.

Proof. By direct computation, we see that $dx = 0$. So x is a cycle. To show it generates $H_n(S^n)$, we note that everything in $H_n(S^n) \cong \mathbb{Z}$ is a multiple of the generator, and since x has coefficients ± 1, it cannot be a multiple of anything else (apart from $-x$). So x is indeed a generator. □

Now we can prove our original proposition.

Proposition. If n is even, then the antipodal map $a \not\simeq \mathrm{id}$.

Proof. We can directly compute that $a_* x = (-1)^{n+1} x$. If n is even, then $a_* = -1$, but $\mathrm{id}_* = 1$. So $a \not\simeq \mathrm{id}$. □

7.7 Homology of surfaces

We want to study compact surfaces and their homology groups. To work with the simplicial homology, we need to assume they are triangulable. We will not prove this fact, and just assume it to be true (it really is).

Recall we have classified compact surfaces, and have found the following orientable surfaces Σ_g.

Example. If $\alpha : S^n \to S^n$ is the antipodal map, we saw that $\alpha_* : H_n(S^n) \to H_n(S^n)$ is multiplication by $(-1)^{n+1}$. So

$$L(\alpha) = 1 + (-1)^n(-1)^{n+1} = 1 - 1 = 0.$$

We see that even though the antipodal map has different behaviour for different dimensions, the Lefschetz number ends up being zero all the time. We will soon see why this is the case.

Why do we want to care about the Lefschetz number? The important thing is that we will use this to prove a really powerful generalization of Brouwer's fixed point theorem that allows us to talk about things that are not balls.

Before that, we want to understand the Lefschetz number first. To define the Lefschetz number, we take the trace of f_*, and this is a map of the homology groups. However, we would like to understand the Lefschetz number in terms of the chain groups, since these are easier to comprehend. Recall that the homology groups are defined as quotients of the chain groups, so we would like to know what happens to the trace when we take quotients.

Lemma. Let V be a finite-dimensional vector space and $W \leq V$ a subspace. Let $A : V \to V$ be a linear map such that $A(W) \subseteq W$. Let $B = A|_W : W \to W$ and $C : V/W \to V/W$ the induced map on the quotient. Then

$$\operatorname{tr}(A) = \operatorname{tr}(B) + \operatorname{tr}(C).$$

Proof. In the right basis,

$$A = \begin{pmatrix} B & A' \\ 0 & C \end{pmatrix}. \qquad \square$$

What this allows us to do is to not look at the induced maps on homology, but just the maps on chain complexes. This makes our life much easier when it comes to computation.

Corollary. Let $f_\bullet : C_\bullet(K; \mathbb{Q}) \to C_\bullet(K; \mathbb{Q})$ be a chain map. Then

$$\sum_{i \geq 0} (-1)^i \operatorname{tr}(f_i : C_i(K) \to C_i(K)) = \sum_{i \geq 0} (-1)^i \operatorname{tr}(f_* : H_i(K) \to H_i(K)),$$

with homology groups understood to be over $\mathbb{Q}$.

This is a great corollary. The thing on the right is the conceptually right thing to have — homology groups are nice and are properties of the space itself, not the triangulation. However, to actually do computations, we want to work with the chain groups and actually calculate with chain groups.

Proof. There is an exact sequence

$$0 \longrightarrow B_i(K; Q) \longrightarrow Z_i(K; \mathbb{Q}) \longrightarrow H_i(K; \mathbb{Q}) \longrightarrow 0$$

This is since $H_i(K, \mathbb{Q})$ is defined as the quotient of Z_i over B_i. We also have the exact sequence

$$0 \longrightarrow Z_i(K; Q) \longrightarrow C_i(K; \mathbb{Q}) \xrightarrow{d_i} B_{i-1}(K; \mathbb{Q}) \longrightarrow 0$$