I apologize for any errors and typos present in these notes. They serve as an outline to the lecture.

Contents

0.1 Project Ideas

Here are a few links to topics in the lecture notes that could make good projects:

- Max Min results 1.1.2
- Graph Polynomials 1.2.1
- Probabilistic Methods 1.5.1
- Algebraic Graph Theory 1.3.1
- Algorithms in Graph Theory
- Persistent Homology
- Discrete Morse Theory

Chapter 1

Graphs and Surfaces

1.1 January 21, January 23

1.1.1 Basic Definitions

Definition 1.1.1 Graph

A *simple graph* G is a collection of *vertices* V and a set of *edges* $E \subset V \times V$ which are unordered in the sense that if $(v, w) \in E$, then $(w, v) \notin E$.

We visually represent a graph by drawing dots for each of the vertices, and drawing an edge between two vertices if the edge is contained in the set E. For notation, if v and w are the vertices at the ends of an edge e, we may refer to e by vw. The set of vertices that share an edge with v are called the *neighbors* of v.

Here are a few examples of graphs.

- The graph with four vertices
$$V = \{1, 2, 3, 4\}$$
and edges given by
$$E = \{12, 23, 13, 14\}$$
Here the graph G may be represented with the following diagram:
$$\begin{tikzcd} 4 \arrow[r, dash] & 1 \arrow[r, dash] \arrow[d, dash] & 2 \arrow[dl, dash] \\ & 3 & \end{tikzcd}$$

- Let $n \in \mathbb{N}$ be a natural number. Then the *complete graph on n vertices* (written K_n) is the graph with n vertices, and an edge between every edge. Here is K_5:

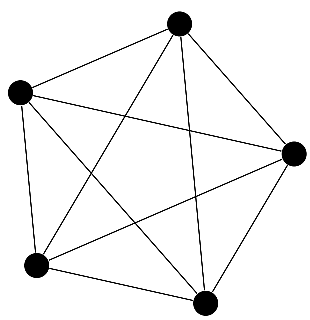

A complete graph has $\frac{1}{2}n(n-1)$ edges in it.

- Every polyhedra in $\mathbb{R}^3$ gives a graph with its edges and vertices. Here is the net of an octahedron:

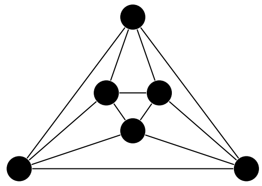

The first example is that of a graph which can be drawn in the plane with no edges crossing. The second is an example of one that cannot be drawn in the plane without edges crossing.

Definition 1.1.2 degree

Let $v \in V$ be a vertex of G. The number of edges connecting to v is called the *degree* of v and is denoted $\deg(v)$.

For example, every vertex of the octahedron has degree 4 to it. The graph K_n is uniquely identified to be the graph that has n vertices of degree $n-1$.

Claim 1.1.3 Let $d(G)$ be the average degree of the vertices of a graph. Then $2|E| = d(G)|V|$.

Proof. By definition, the average degree of the vertices of a graph is given by

$$d(G) = \frac{1}{|V|} \sum_{v \in V} \deg(v).$$

Each edge contributes $+1$ to the degree of 2 different vertices. Therefore, the sum of all of the degrees is equal to half the number of edges. Therefore $2|E| = d(G)|V|$. ■

In particular, this proves the claim above about complete graphs, which is that the number of edges in a complete graph is $n(n-1)/2$.

Corollary 1.1.4

The number of odd degree vertices in any graph is even.

This simple fact about graphs comes into use again and again for some reason.

Definition 1.1.5 Paths and Cycles

Let G be a graph. A *path* in G is a sequence of distinct vertices $\{v_i\}$ such that $v_i v_{i+1}$ is an edge in G for every i. The length of the path is the number of edges in it. A *cycle* in G is a path P such that the first and last vertex of P share a common edge.

Paths and cycles are the simplest building blocks to a graph.

Claim 1.1.6 Let δ be the minimal degree of the vertices in G. There is a path of length at least δ in G.

Proof. Let P be a path of maximal length. Let v_i be the vertices of the path. Since P is maximal, every neighbor of the v_i are contained in P. Since every vertex has degree at least δ, this means that P must contain δ different vertices. ■

1.1.2 Connectivity

The first topological property that we are going to explore is connectedness.

Definition 1.1.7 Vertex Connectedness

A graph G is called *connected* if for any two vertices v and w, there exists a path from v to w. It is called

k-connected if whenever you remove less than k vertices, the resulting graph remains connected. The greatest k such that G is k-connected is called the *connectivity* of G and is denoted $\kappa(G)$.

Connectedness is a topological property of a graph. Determining the connectedness of a graph is a not a quick thing to do, however, we can get estimates on the connectedness of average components of the graph by knowing the average degree of the vertices of a graph.

Theorem 1.1.8 Mader 1972

Every graph with average vertex degree $4k$ contains a k connected subgraph.

Proof. We instead prove that the graph is k connected if $|V| \geq 2k-1$ and $|E| \geq (2k-3)(|V|-k+1)+1$, which follows from the average vertex requirement. Our argument will be by induction on the number of vertices in the graph.

- (Base Case) In the case when the first inequalities is sharp, then $|V| = 2k-1$ and $|E| = \frac{n(n-1)}{2}$. Therefore, G is a complete graph on $2k-1$ vertices, and we can find a K_k in this graph.
- (Inductive Step, type 1) Suppose that the inequality is not sharp. Then pick a vertex of degree less than 2k-3. Then when we remove this vertex, the induction hypothesis still holds.
- (Inductive Step, type 2) Suppose that the inequality is not sharp, and we cannot find a vertex of degree $2k-3$. Then assume for contradiction that G is not k-connected. Then there are two subgraphs $G_1, G_2 \subset G$ such that $G_1 \cap G_2$ has fewer than k vertices. Every vertex in $G_1 \setminus G_2$ has neighbors only in G_1. Since the minimal degree of each vertex is $2k-2$, we have that G_1 has at least $2k-2$ vertices. Similarly, G_2 has $2k-2$ vertices. However, one of the G_i must satisfy the induction hypothesis! This is because $|E| \leq |E_1| + |E_2|$, and so

  $$|E| \leq (2k-3)(n-k+1)$$

  if neither E_1 or E_2 satisfy the induction hypothesis.

■

Yet another notion of connectivity is edge connectivity. This is the minimal number of edges that you must remove to disconnect the graph. However, the vertex connectivity and the edge connectivity are usually not well related to each other. Take for instance this graph which is only 1-vertex connected, but is highly edge connected:

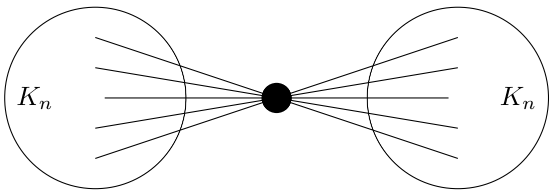

While there is not an easy relation between vertex connectivity and edge connectivity, there is a relation between path connectivity and vertex connectivity. A graph is *k-path connected* if any two vertices v and w can be joined by at least k disjoint path.

Theorem 1.1.9 Menger

The path connectedness of a graph is equal to the connectedness of the graph.

Notice the contrast between the maximality of the path connectedness, and the minimality of the vertex set disconnecting the sets A and B. Exploring Max-Min results would be an interesting project idea.

Proof. We prove a stronger statement: Let G be k-connected, and let A and B be 2 subgraphs. Suppose that B has at least k vertices. Let $b_1, \ldots, b_n \in B$, with $n < k$. Let $\mathcal{P}$ be a collection of n disjoint paths in $G \setminus (B \setminus \{b_1, \ldots b_n\})$ with endpoints $b_1, \ldots b_n$ in B, where $n < k$. Then there exists a collection $\mathcal{P}'$ of $n+1$ disjoint paths with endpoints $b_1, \ldots b_n, y \in B$.

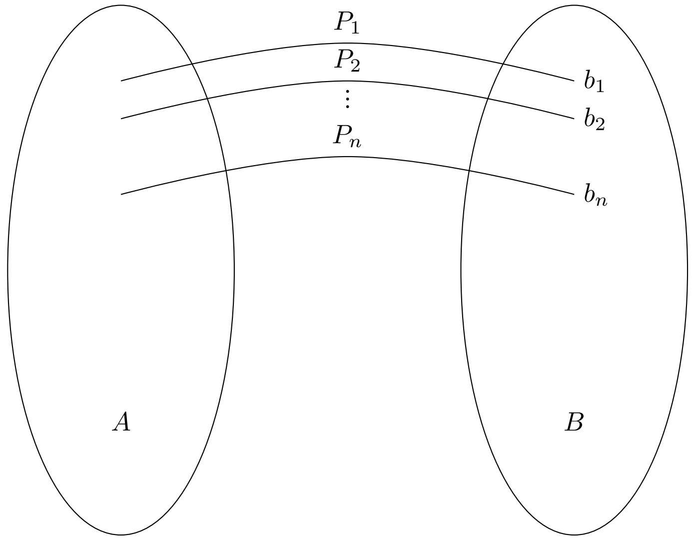

We induct on the size of $G \setminus B$. Clearly, when B is all of G, this statement is easy to prove. Now, let's assume that the statement is true for all B' containing B. Look at our selection $\mathcal{P}$, and take any new path connecting A to B.

- It may be the case that this path is disjoint from the paths of $\mathcal{P}$. In this case, we are done.

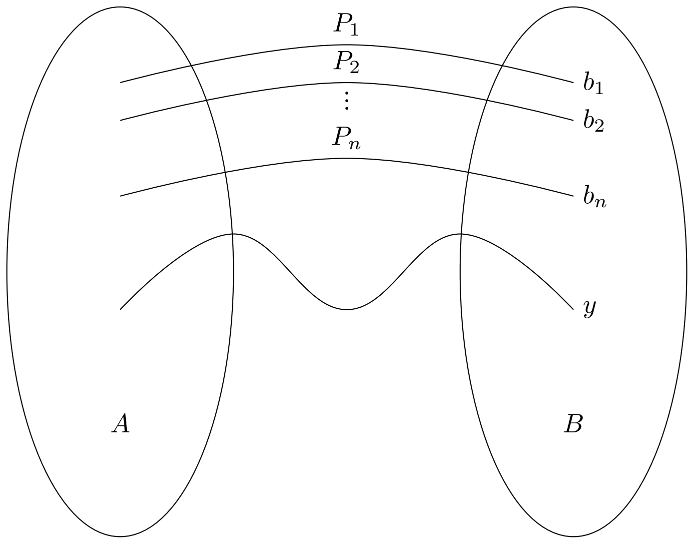

- It may be the case that our path intersects a bunch of the other paths. Let x be the last point at which the new path intersects the old paths. Without loss of generality, we may assume that the intersecting path is P_n.

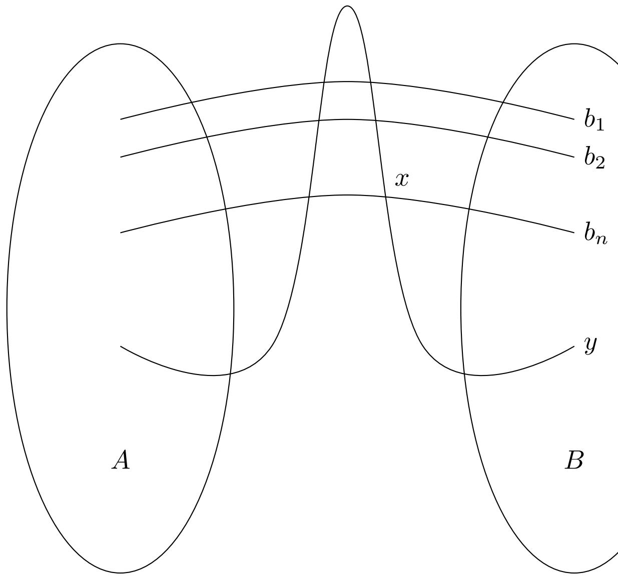

Then we will use the inductive hypothesis to form a set of new paths. Let B' be the subgraph with the paths x to b_n and y thrown in. By the inductive hypothesis, there exists independent paths from A to $b_1, \ldots b, x, z$ for some new point z. There is an additional path from x to z and x to b_n.

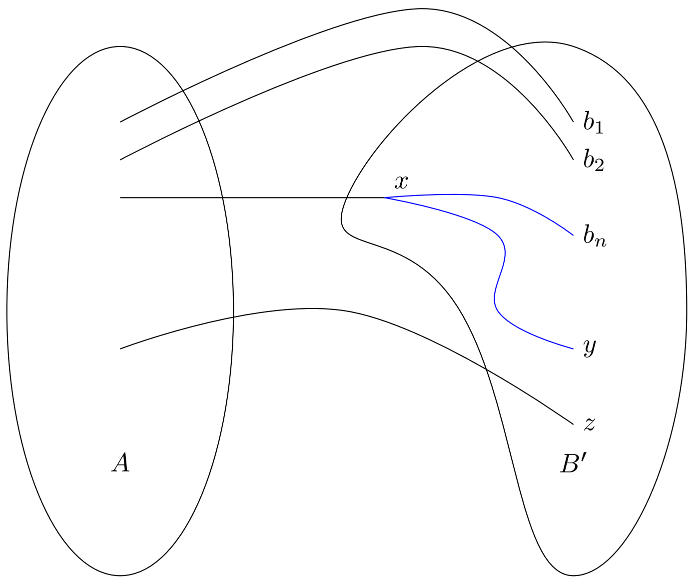

Now, we are interested where z is in the set B'. There are three cases.

1. $z \in B$. In this case, there is no interaction between z and the paths. Then we are done. (replace y with z)
2. The point z lies on the path x to b_n. Then extend on to b_n, and route the path from A to x into a path from A to y.
3. The point z lies on the path x to y. Then extend the path A to z to y.

■

1.2 January 30

Lemma 1.2.1

Suppose that G is 2 connected. Let v w be two vertices in G. Then if $H = G \cup P$, where P is a path from v to w, then H is still 2-connected.

Proof. Suppose not. Then G is 2-connected, $G \cup P$ is not 2 connected. Let v be a *cut vertex* of $G \cup P$. Clearly, $v \in P$, as if $v \in G$, then v would be a cut vertex of $G \cup P$. If $v \in P$, it separates P into 2 components. However, each of those components is connected to G by their ends, and therefore $G \cup P \setminus \{v\}$ is still connected. ∎

Theorem 1.2.2

Let G be a 2-connected graph. Then either G contains a 2-connected subgraph H so that $G = H \cup P$, where P goes between two vertices of H, or G is a cycle.

Proof. G contains a cycle, so it contains a maximal 2-connected proper subgraph H. Look at a vertex $v \in G \setminus H$, and a vertex $w \in H$. As G is 2-connected, by Menger's theorem there exist disjoint paths P_1, P_2 from v to w. This gives two paths from w to that are contained in H. This is equivalent from a path from G to itself that goes through H. This contradicts the maximality of H. ∎

Corollary 1.2.3

Every 2-connected graph is generated from a cycle with the subsequent addition of paths.

1.2.1 Minors and Topological Minors

The first question we should answer are which graphs we should consider topologically equivalent. Topologists are interested in classifying objects up to continuous deformations. While we will not define what the word "continuously deform" means, it has the intuition of bending objects like they were made of play-doh. For us, graphs already have a bit of this property. When we draw the visual representation of a graph, we do not care how we draw the edges of the graph. For example, here are a couple depictions of the same graph:

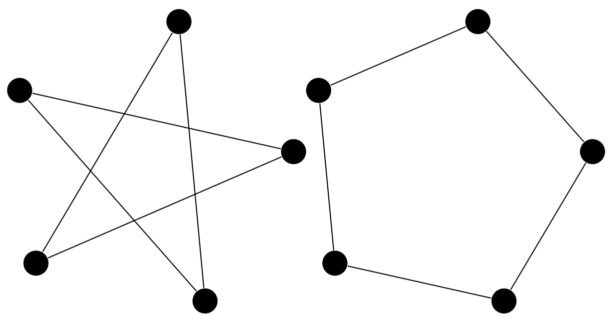

Both of these describe a circle topologically. Therefore, two graphs being isomorphic mean that the realizations describe the same object.

Definition 1.2.4
Graph Isomorphic

We say that two graphs G and H are *graph isomorphic* if there exists a bijection $\phi : V_G \to V_H$ which induces a bijection on the edges of the graphs.

Subdivision and Topological Minors

For topological comparisons, graph isomorphism cannot describe the whole entire story. In particular, one should believe that all cycles topologically describe the same shape, regardless of the number of vertices in the shape.

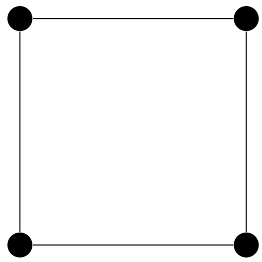

One method for discretizing the notion of topological continuity is that of subdivision.

Definition 1.2.5 **Subdivision**

Let G be a graph, and let e be an edge in G. Then the *subdivision* of G at e is the graph $G \div e$, where the edge e has been replaced by a new vertex v_e and two edges $e_\pm$. If H is obtained from G by a series of subdivisions, we say that H is a subdivision of G. If a subdivision of G is contained in H, we say that G is a *topological minor* of H. The set of all graphs that contain G as a topological minor is denoted TG.

For example, the cycle on 5 vertices is a subdivision of the cycle on 4 vertices. Here is a graph that contains every graph on 5 or fewer vertices as a topological minor:

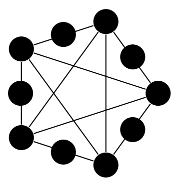

For those of you who have seen some point set topology, if H is a subdivision of G, then H and G are homeomorphic. If G is a topological minor of H, then G is continuously embeddable into H. Notice that this is not the same thing as homotopy type, as all graphs are homotopic to a bouquet on k loops, where k is governed by the edge connectivity of the graph G. Later we will see a deformation that does this kind of thing. For those of you who want a catchphrase: "subdivision preserves the topological properties of a graph completely, but doesn't generally play well with topological properties." For instance, the connectivity of a graph is completely independent of its subdivision type. This is really sad!

Exercise 1.2.6 Rigorously formulate the statement "subdivision preserves the topological properties of graphs completely" and prove it.

Contraction and Minors

A slightly stronger type of graph deformation is contraction along and edge, which is a bit like the inverse operation of subdivision.

Definition 1.2.7

Let G be a graph, and xy an edge in G. Then define the *contraction* G/xy to be the graph with vertices of G less x and y and a new vertex v_{xy}, whose neighbors are those of x and y.

Claim 1.2.8 H is a subdivision of G, then G is a contraction of H.

Notice that the converse is not true– every graph can be contracted to a point.
Similarly, we have the notion of a *minor* as oppose to a topological minor. A graph G is a topological minor of H is G is a subgraph of a contraction of H. While it is clear that every topological minor is a minor, it is not the case that ever minor is a topological minor. An example of such a graph is the *Peterson graph*, which

contains K_5 as a minor, but not a topological minor.

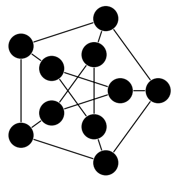

Clearly K_5 is a minor of this, however, this graph doe not contain a subdivision which is K_5 as every vertex has degree 3. Asking for topological minor is a stronger property than just a minor.
Contraction doesn't preserve many of the topological properties of graphs– in fact, there are no topological invariants based on the contraction type of a graph besides the number of connected components.

Exercise 1.2.9 Formulate and prove the statements "there are no topological invariants based on contraction type."

However, contraction usually plays well with combinatorial properties of graphs. It is frequently used in induction arguments of graphs.

Proposition 1.2.10 Let G be a graph. The reliability polynomial $R_G(p)$ returns for every value $p \in [0,1]$, the probability that G is connected if every edge of G is disconnected with a probability of $1-p$. It is a polynomial in p of degree at most $|E|$.

Proof. We prove by induction on the number of edges. Take any edge e in G.

- With probability p, the edge e is not deleted. Then the probability that G is connected after removing other edges is $R_{G/e}(p)$.
- With probability $(1-p)$, the edge e is deleted. Then the probability that G is connected after removing other edges is $R_{G\setminus e}(p)$.

These two above cases are disjoint. Therefore the probability that G is connected after removing each edge with probability $(1-p)$ is $pR_{G/e}(p) + (1-p)R_{G\setminus}$.
Since G/e and $G \setminus e$ each have one fewer edge, their reliability polynomials are both polynomials of degree at most $|E_{G/e}| = |E_{G\setminus e}| = |E_G| - 1$. Therefore, $R_G(p)$ is a polynomial of degree at most p. ∎

3-connectedness

Now, let's do a more sophisticated induction argument to prove a classification result for 3-connected graphs.

Lemma 1.2.11 Let G be a graph which is 3-connected with more than 4 vertices. Then there exists an edge $e \in G$ such that G/e is still 3 connected.

Proof. Suppose for contradiction there is no such edge. Define the function $f : E \times V \to V$ by

$$f(xy, z) = \{\text{Size of the smallest component of } G \setminus \{x, y, z\}\}\}$$

Given our hypotheses, we will prove that f is nonconstant, and has no minimal value, which is a contradiction. First to show that f is nonconstant.

- Pick any edge xy. Since contraction preserves connectivity G/xy is connected. There always exists a vertex whose removal keeps the graph connected. Therefore one possible value of f is $|V| - 2$.
- On the flip side, given any edge xy, we know that G/xy is 2 connected. Since G/xy is 2 connected but G is not, we know that v_{xy} is part of a separating set of G/xy. Let z be any vertex such that $(G/xy) \setminus \{v_{xy}, z\}$ is disconnected. Then x, y, z separate G, which means that $f(xy, z) \le |V| - 3$.

Now to show that f has no minimal value. Suppose xy, z is a minimizing set. Then let C be the smallest component of $G \setminus \{x, y, z\}$. We know that z has a neighbor in the component C. Call this vertex w. Notice that the neighbors of w are entirely contained in C.

We also know that G/zw is 2-connected. By similar argument, this means that the removal of w, z and some additional vertex u will separate G. Let D be a component of $G \setminus \{w, z, u\}$ which does not contain xy. We have that w must have a neighbor in D (otherwise the removal of w would not be necessary). Therefore $D \cap C$ is nonempty. Furthermore, every vertex of D is contained in C, because it is disjoint from all the components of $G \setminus \{x, y, z\}$.

It does not contain w, so D is a strict subset of C. Therefore D is smaller size than C and is disconnected by vertices w, z, u. So $f(wz, u) < f(xy, z)$. This is a contradiction! ■

We can use this lemma to prove a result on the structure of 3-connected graphs:

Theorem 1.2.12 Tutte

A graph G is three connected if and only if there is a series of contractions of G to K_4, where each contraction is on edges with vertices of degree greater than 3.

Proof. The forward direction is easy. The other direction is tricky.
Now, suppose that G/xy is 3 connected, and x and y both have degree greater than or equal to 3. Suppose that G is not 3-connected. Let u, v separate G into two sets. Then it must be the case that xy are in the same connected component C of $G \setminus \{u, v\}$.

- Suppose C contains both x and y. Then G/xy would still be separated by u and v. Contradiction.
- Suppose $u, v \in \{x, y\}$. Then G/xy is disconnected by v_{xy}, a contradiction.
- Suppose $u \in \{x, y\}$, and suppose C contains a vertex $w \neq x, y$. Then G/xy is still separated by u and v. Contradiction.
- Suppose $u \in \{x, y\}$, and $C \subset \{x, y\}$. Then either x or y has degree less than 3. Contradiction.

■

1.3 February 2

1.3.1 Algebraic Graph Theory

In this section we introduce a common technique which we will see later employed as a topological method. Let $\mathbb{F}_2$ be the field of two elements. Given a graph G, define the *edge space* $\mathcal{E}$ to be the $\mathbb{F}_2$ vector space generated on a basis E. In other words, elements of $\mathcal{E}$ correspond to subsets of G, and the vector addition corresponds to symmetric difference. Define $\mathcal{C}$, the *cycle space*, to be the subspace spanned by all the edge sets giving cycles in $\mathcal{E}$. We will be interested in understanding the algebraic structure of $\mathcal{C}$. In particular, we are interested in generating sets of $\mathcal{C}$. While this seems like a cumbersome object, it is a convenient tool for storing a lot of data about G.
For example, an *induced cycle* of G is a cycle which cannot be made into 2 smaller cycles with the addition of an edge in G.

Claim 1.3.1 The induced cycles of G generate the space $\mathcal{C}$.

In the case of 3-connected graphs, we can do better.

Theorem 1.3.2 Tutte

The cycle space of 3-connected graphs is generated by non-separating induced cycles.

First, let's look at a 2-connected graph where this is not true.

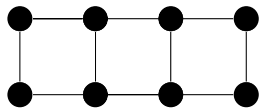

However, this theorem is not an if and only if. For instance, the requirement of 3 connectedness is not necessary. For instance, any tree will trivially satisfy the requirement that its cycle space is generated by non-separating induced cycles.

Proof. We use 1.2.1 to prove this by induction on the number of edges.
In order to prove this, we will need to add some terminology. Let e be an edge of G.

- For notation, we will say that a non-separating induced cycles are "basic", and all cycles generated by them are "good".
- Notice that $\mathcal{E}_{G/e}$ can be identified with a subspace of $\mathcal{E}_G$- however, this identification is not canonical!
- Given a cycle C in G, C/e is identified with a cycle in G/e unless $C = C_3$. If this is the case, we call C a *fundamental triangle*. Note that C/e does not necessarily identify to the edge set as C in $\mathcal{E}_G$, even if C does not contain the contracted edge e.

Diestel's exposition proves this in 3 steps:

1. Every fundamental triangle is basic.
   Assume for contradiction that this is not true. Then the removal of the edge from the contraction of the fundamental triangle would separate G/e, which is a contradiction!
2. If C is an induced cycle in G, but not fundamental, then $C + C/e + D = 0$ or $C + C/e + D = e$ for some good D.
   This says that the cycle space of G differs from the cycle space of G/e by at most an edge and a good cycle. Where the edge comes from is obvious. Suppose upon contracting our cycle and lifting it back up, that we have something which is not a cycle. A quick case check shows that $C/e + C$ is a combination of fundamental triangles and edges.
3. For every basic cycle $C' \subset G/e$, there exists a basic cycle $C \subset G$ such that $C/e = C$.
   The proof of this is a little trickier. Assume that we are in the case where the graph G looks something like the picture below, and the edge being contracted is xy.

   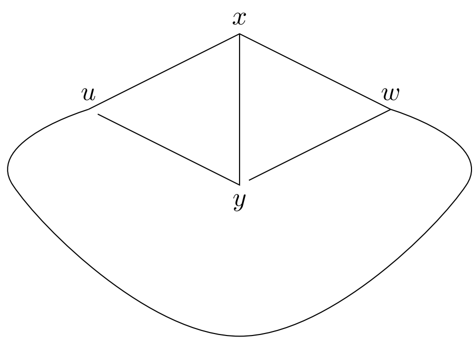

   Suppose that cycle C' is the one that "goes around" the graph and then through the contracted vertex v_{xy}. Then we have two candidate cycles– the one that goes through the x vertex, and the one that goes through the y vertex.

   - If C_x separates G, then y is in its own component.
   - If C_y separates G, then x is in its own component.
   - If both x and y are in their own components, then the only neighbors of x and y are u and w. However, u and w then separate G,contradicting the connectivity of G.

Similar arguments can be made for the cases where G does not have the edges ux or uy.

At this point we are able to finish the proof. Let C be any cycle. By induction hypotheses, $C/e = C'_1 + \ldots + C'_k$, where each of the C'_i are basic in G/e. For each i, we have a basic cycle C_i so that $C_i/e = C'_i$. This cycle differs from C_i by possibly a good cycle and an edge by the second claim. So we have that

$$\sum C_i = \sum C'_i + \text{ Good cycles and possibly } e = C/e + \text{ Good cycles and possibly } e$$

Likewise we have that

$$C = C/e + \text{ Good cycles and possibly } e$$

This gives us that

$$C = \sum C_i + \text{ Good cycles and possibly } e$$

However, each of the C_i is basic and a cycle in G. So the sum must not include the set e. The sum of basic cycles and good cycles is good, therefore C is good. ■

1.4 February 4

We will use some algebraic criterion to classify other sets types of graphs.

1.4.1 Planar Graphs

We are now going to explore embedding of the graph.

Definition 1.4.1 A *planar representation* of a graph G is a set of vertices $V \subset \mathbb{R}^2$ and a set of continuous arcs $\{f_e : I \to \mathbb{R}^2\}_{e \in E}$ indexed by E such that

- If $f_e(t) = f_{t'}(x')$ and $t, t' \neq 0, 1$, then $e' = e$ and $x = x'$.
- If xy is an edge, then $f_{xy}(0) = x$ and $f_{xy}(1) = y$.

Note that a graph G has many different planar representations, and it can be difficult to tell what it means for two planar representations to be the same. We will leave this problem until later- first, we will explore some properties of planar representations.
Perhaps most importantly, a planar representation of a graph gives us a new object to keep track of: the *faces* of a graph. Each connected component of $\mathbb{R}^2 \setminus G$ gives us a face whose boundary components are cycles in G. Notice that this definition gives us a large "outer face" to the graph. The information contained by the faces is much more than just their bounding cycles: take for instance these two graphs, which have 1 face each, are not the same planar graph:

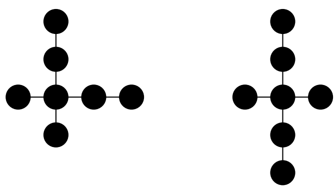

Even without having a rigorous idea of what isomorphic planar graphs should be, we can explore some of the properties of planar graphs.

1.5 February 5

1.5.1 Graph Colorings I

We're going to take a slight detour to talk about graph colorings, which are a major tool in graph theory.

Definition 1.5.1 Colorings

Let G be a graph. A *k-graph coloring* of G is an assignment $f : V \to \{1, 2, \ldots, k\}$ such that if $xy \in E$, $f(x) \neq f(y)$. The minimal k such that a k coloring of G exists is called the *chromatic number* of G and is denoted $\gamma(G)$.

Colorings, like connectivity, are both influenced by local properties of the graph and by global properties of the graph. For instance, a local result is:

Claim 1.5.2 Let $\Delta(G)$ be the maximal degree of vertices in G. Then G admits a $\Delta + 1$ coloring.

A global result on coloring is:

Claim 1.5.3 Let G be a graph. Then

$$\gamma(G) \leq \frac{1}{2} + \sqrt{2|E| + \frac{1}{4}}$$

Proof. Create a on $\gamma(G)$ vertices, where each vertex represents one color class from G, and you draw an edge between two vertices if the color classes have an edge between them. If this is an efficient coloring, then this multigraph must have edges between every two color classes; otherwise those classes could be labeled the same way. Since this is a complete graph on $\gamma(k)$ vertices, G must have at least $\frac{1}{2}k(k-1)$ edges. ∎

In fact, we can improve our local result on coloring by 1.

Theorem 1.5.4 Brooks

Let G be a graph which is not an odd cycle or a complete graph. Then $\gamma(G) \leq \Delta(G)$.

Proof. We induct on the number of vertices. Suppose G does not satisfy the theorem. Then $G \setminus v$ does, and therefore v must have degree $\Delta(G)$. Given a coloring of $G \setminus v$, we get a Δ-coloring of G where the neighbors of v_i must use each color exactly once– let's index these neighbors by their color label. Let $H_{i,j}$ be the subgraph of $G \setminus v$ which uses the i and j color.

- For all i and j, v_i and v_j lie in the same component of $H_{i,j}$.
- For all i and j, the component of $H_{i,j}$ that contains i and j is a path from i to j.
- For all i, j, k the paths from v_i to v_j and v_i to v_k of $H_{i,j}$ and $H_{i,k}$ only intersect in the vertex v_i. (As any other common vertex 2 neighbors colored i and 2 neighbors colored j)

Without loss of generality assume that v_1v_2 is not an edge of G. Let u be on the path from v_1 to v_2. It must be colored 2. Pick a new coloring for G by switching colors 1 and 3 in the path from v_1 to v_3. Then u is on the path from $v_1 = v_{(3)}$ to v_2, and is still colored 2. Since the path from v_1 to v_2 is still the same color except at the end, we have that u is in the path from $v_{(1)}$ to v_2 and from v_2 to $v_{(3)}$. This contradicts the last little statement we made. ∎

It turns out that neither global or local results allow us to pin down the colorability of a graph. Neither will information on the connectivity of a graph. For example:

- Average degree does not give us a good pin on chromaticity. Take, for instance, a graph that has a K_n subgraph but then a huge number of free-floating vertices. This is n colorable, although the average degree is very small.
- Connectivity does not give us a good pin on chromaticity. Take, for instance the complete bipartite graph on n and n vertices: this is a graph with vertices $v_1, \ldots v_n, w_1, \ldots w_n$ with edges w_iv_j for every i, j. This graph is n connected, but is 2-colorable. On the other hand, a tree is a 2 connected graph which is also 2 colorable.
- Local structure does not give us a good pin on chromaticity. Take for instance this theorem of Erdös:

Theorem 1.5.5 Erdös

There exists a graph G which contains no cycles of size smaller than n, but have chromatic number at least k, for every n and k.

In order to prove this, Erdös used a technique called the probabilistic method. For all of these difficulties, we expect that chromatic number of graph is a rather mysterious value to get a good pin on. However, colorings of graphs have some unexpected properties.
2-colorability is such a nice property that it gets its own special name : bipartite.

Claim 1.5.6 A graph G is bipartite if and only if all of its cycles are even length.

Proof. The right direction is easy. Take any cycle, and try to color it with two colors– this is only possible if the cycle has an even number of vertices.
For the left direction, take any vertex $v \in G$. Without loss of generality, let G be connected. For any vertex w, take a path P_w from v to w. If P_w has even length, color w white, otherwise color w black. As all cycles are of even length, the parity of the length of P_w is independent of path chosen. ∎

Claim 1.5.7 The maximal number of edges in a bipartite graph with at most n vertices if $\lfloor n/2 \rfloor \lceil n/2 \rceil$.

In bipartite graphs, we have another min-max result similar to the min-max result that we saw for connectivity.

Theorem 1.5.8 König's Theorem

Let G be a bipartite graph. Let $U \subset V$ be a subset minimal with respect to the property that every vertex of v has a neighbor in U. Let $F \subset E$ be a subset of edges maximal with respect to the property that all edges of f are disjoint. Then $|U| = |F|$.

Usually, this is stated as "the maximal matching is the minimal vertex cover."

Proof. We will construct a special graph that reduces this to the connectivity statement that we had before. So we are looking for a graph where:

- The path connectedness gives disjoint edges.
- Vertex connectivity corresponds to vertex covers.

Let $V_G = V_1 \sqcup V_2$ be a coloring of the vertices of G. Then add additional vertices w_1, w_2 so that the neighborhood of w_1 is all of V_1 and the neighborhood of w_2 is all of V_2. Then it is the case that a separating set of vertices for w_1 and w_2 exactly corresponds to a vertex cover, the maximum number of disjoint paths from w_1 to w_2 is exactly the larges set of disjoint edges in G. ∎

In fact, König's theorem is equivalent to Menger's theorem.
Two colorability is a rather easy property to look for. Given a graph G with n vertices, determining if G is 2-colorable is an algorithm that runs in n steps. However, the algorithm to check if a graph is even 3-colorable runs in exponential time with respect to the number of vertices, and is one of the hallmark "NP-complete" problems.
Colorings of graphs follow some nice algebraic properties:

Definition 1.5.9 Chromatic Polynomial

Let G be a graph. For $k \geq 0$, define $P_G(k)$ to be the number of k colorings of G.

As the name suggest, the chromatic polynomial is a polynomial.

Theorem 1.5.10

$P_G(k)$ is a polynomial with integer coefficients of degree $|V_G|$.

Proof. We prove a stronger statement, that the chromatic polynomial satisfies the recursive relation:

$$P_G(k) = P_G G \setminus xy(K) - P_{G/xy}(K)$$

for any edge xy. This along, with the knowledge that if H_n is the graph with n vertices and no edges, then $P_{H_n}(k) = k^n$, will prove that $P_G(k)$ is a polynomial with integer coefficients.
The proof of the statement is now quite easy– note that colorings of G are in bijection with colorings of $G \setminus xy$ where x and y do not have the same color. Further, colorings of $G \setminus xy$ where x and y do have the same color are in bijection with G/xy. This finishes the proof. ∎

The chromatic polynomial satisfies some very unusual properties:

Theorem 1.5.11 Jackson, Thomassen

The zeros of the chromatic polynomials cannot occur in the interval $[-\infty, 32/27] \setminus \{0, 1\}$, and are dense everywhere else.

Interestingly, the roots of chromatic polynomials are not even the algebraic numbers. For instance, $\phi + 1$ may never be the root of a chromatic polynomial. However, the set of all roots of chromatic polynomials is dense in C. The properties of chromatic polynomials is still not well understood.

1.6 February 6

Claim 1.6.1 Let G be a connected planar graph, and let V, E, F be the set of vertices, edges and faces.

- $2|E| \geq |V|$
- $2|E| \geq 3|F|$
- $|V| - |E| + |F| = 2$ (Euler's Formula)
- If $|V| \geq 3$, then $|E| \leq 3|V| - 6$.

Proof. The first claim follows from our argument about average degree and edges. Since G is connected, the average degree is at least 1.
The second claim is similar: every edge belongs to at least 1 face, and every face has at least 3 edges.
The third claim follows from induction on the number of edges. The base case is when G is a tree- then G has 1 face, n vertices and $n - 1$ edges. By induction, the addition of every edge creates a new face and a few edge, so the quantity $|V| - |E| + |F|$ remains constant. The third claim follows from triangulation. Fix a graph G. You can always add an edge to G and keep it planar whenever G has a face which is not a triangle. When all of the faces of G are triangular, the graph is *maximally planar*. This means that $2|E| = 3|F|$. From Euler's formula, we have

$$|V| - |E| + 2/3|E| = 2$$

from which the claim follows. ∎

Corollary 1.6.2

The only Platonic solids are the tetrahedron, cube, octahedron, dodecahedron or icosahedron.

Proof. Recall that a Platonic solid is one where all the vertices have the same degree and all of the faces have the same number of edges. We can take any polyhedron and convert it into a planar graph by stereoscopic projection. Let m be the degree of the vertex, and n the size of each face. Then it is the case that:

$$2|E| = n|F|$$

$$2|E| = m|V|$$

Therefore, $|F| = m/n|V|$. The formula for Euler characteristic tells us that:

$$|V| - m/2|V| + m/n|V| = 2$$

Now, we have some bounds on m and n.

- We know that n is at least 3.

- The quantity $m\left(\frac{1}{n} - \frac{1}{2}\right) > -1$. This means that m cannot be greater than 5.
- m must be at least 3. Valid values of m are now 3, 4, 5.

Tabulating our results we have:

m	n	$\lvert V\rvert$	$\lvert E\rvert$	$\lvert F\rvert$	Shape
3	3	4	6	4	Tetrahedron
3	4	8	12	6	Cube
3	5	20	30	12	Dodecahedron
4	3	6	12	8	Octohedron
5	3	12	30	20	Icosohedron

■

There is a kind of duality that you might notice here on first inspection. First of, in the proof, it seems like the values of m and n are exchangeable. This is reflected in the platonic solids that we've found- they come in pairs where the roles of vertices and faces are reversed. This is a general phenomenon among planar graphs:

Definition 1.6.3 Let G be a connected planar graph. The *dual (multi)-graph* to G, is the graph G^* constructed by taking G, placing a vertex in the middle of each face, and connecting two vertices if their corresponding vertices share an edge.

A multi graph is a graph where vertices can have multiple edges going between them. This is necessary to produce the dual of trees, which become bouquets:

Exercise 1.6.4 Prove that $G = (G^*)^*$.

We can also use these inequalities to show that certain graphs are non-planar:

Lemma 1.6.5 For every $n \geq 5$, the complete graph K_n admits no planar embedding. Similarly, the complete bipartite graph $K_{3,3}$ has no planar embedding.

Whether $K_{3,3}$ is embeddable frequently referred to as the "Utilities Problem," which asks you to connect 3 utilities to 3 apartments without crossing the connections.

Proof. It suffices to show this for $n = 5$. For K_5, we have that $|E| = 10$, and $|V| = 5$. This fails to satisfy $|E| \leq 3|V| - 6$. Therefore K_5 is nonplanar.
For $K_{3,3}$, we know have a similar estimate, with $|E| = 9$ and $|V| = 6$. This means that the number of faces must be 5 by Euler's formula. Since $K_{3,3}$ is bipartite, the minimal face size is 4. This means that there must be at least 10 edges. However, there are only 9 edges, a contradiction. ■

Kuratowski's Theorem

In fact, $K_{3,3}$ and K_5 form the essential obstruction to a graph being planar.

Theorem 1.6.6 (Kuratowski) The following are equivalent:

- G is planar
- G contains $K_{3,3}$ or $K_{5,5}$ as a topological minor.
- G contains $K_{3,3}$ or $K_{5,5}$ as a minor.

We will prove a number of proposition about K_5 and $K_{3,3}$ first.

Proposition 1.6.7 If the $\Delta(G)$, the maximal degree of G is at most 3, then the following are equivalent:

- G is a topological minor of H.
- G is a minor of H.

Proposition 1.6.8 The following are equivalent:

- G contains K_5 or $K_{3,3}$ as a topological minor.
- G contains K_5 or $K_{3,3}$ as a minor.

Proof. First, we show that if G contains $K_{3,3}$ as a minor, it contains it as a topological minor. This follows from a slightly stronger claim: If the maximal degree of H is 3 or less, then G contains H as a topological minor if and only if it contains H has a minor. This follows from the fact that if you ban contracting 2 degree 3 vertices together, graph contraction can be reversed by subdividing edges.
So, we only need to show that if G contains K_5 as a minor then it contains K_5 or $K_{3,3}$ as a topological minor. Let H be a maximal subgraph of G which contains K_5 as a minor. Look at each of the blobs corresponding to a vertex of K_5. If one of these is not a topological 4-claw, then it has exactly 2 vertices of degree 3. Then we can find a topological $K_{3,3}$. ∎

We start with a proof of Kuratowski's theorem for 3-connected graphs using the contraction lemma that we have.

Lemma 1.6.9 Every non planar 3-connected graph contains $K_{3,3}$ or K_5 as a minor.

Proof. We induct on the number of edges in the graph. Suppose that G is 3-connected, contains no K_5 or $K_{3,3}$ as a minor. There exists an edge xy such that G/xy is also three connected– furthermore, G/xy does not contain a K_5 or a $K_{3,3}$ as a minor. By inductive hypothesis this means that G/xy is planar. Take the union of all vertices that share a face with v_{xy}. This forms a cycle as $G/xy - v_{xy}$ is 2 connected. The neighbors of x and y are contained within this cycle.
By viewing the vertex v_{xy} as the vertex x, we get a planar embedding of $G \setminus y$. Now we try to add the vertex y in.
We will only run into a problem if the vertices of y belong to two different faces containing x in $G \setminus y$. One can see that the obstructions to such a construction is a topological $K_{3,3}$ or K_5. ∎

Claim 1.6.10 Suppose G is 3-connected and planar. Then there exists a drawing of G which has all convex faces.

Proof. We induct on the number of edges as before. The base case is K_4, and this graph is clearly connected. Pick xy, an edge of G so that G/xy is still 3-connected. By the induction hypothesis, there exists an embedding of G/xy so that every face is convex. We can treat the vertex v_{xy} as the vertex x. The addition of a vertex and edges to the center of any convex face can be done in such a way that every additional edge creates a convex face. ∎

This means that in particular every single planar graph which is 3 connected can be drawn with straight edges. In fact, every planar graph can be drawn with straight edges.
We now have enough machinery to prove Kuratowski's theorem.

Proof of Theorem 1.6. We may assume that G is 2 connected. Let x, y be a cutset of G so that

$$G = G_1 \cup G_2 \qquad G_1 \cap G_2 = \{x, y\}$$

Then if G does not contain a topological $K_{3,3}$ or a K_5, neither does $G_1 \cup xy$ or $G_2 \cup xy$. This is because there is a path from x to y in both G_1 and G_2. Therefore by induction there exits a planar drawing of both

$G_1 + xy$ and $G_2 + xy$. We can assume that these drawings both have xy as an exterior edge– gluing along this edge yields a planar drawing of G. ■

The search for a topological K_5 or $K_{3,3}$ seems like a difficult thing to do in a graph. There is an algebraic criterion for planarity of a graph as well, which is easier to check.

1.7 February 9

We now look at a criterion for determine planarity algebraically within a graph:

Definition 1.7.1 A subspace $\mathcal{U} \subset \mathcal{E}$ is called *simple* if there exists a basis for $\mathcal{U}$ so that every edge e shows up in each basis element at most twice.

Theorem 1.7.2 A graph is planar if and only if there is a simple basis for the cycle space.

Proof. In the case of planarity, choose a basis given by the faces of some planar embedding of G. This is a basis as every cycle C is given by the sum of cycles on the interior of C with respect to that embedding. As every edge is contained in exactly 2 faces, we have that this satisfies the requirement $(*)$.
For the other direction, we first that the requirement $(*)$ is unchanged by the minor property. This is by the projection of $\mathcal{E}_G$ to $\mathcal{E}_{G/e}$. The only time that the cycle space changes dimension under this projection is when you contract a triangle– but in this case, you remove on cycle, and one basis element which contained the contracted edge. This means that contraction does not change the property of having a simple basis for the cycle space. Similarly, removing edges does not change the property of having a simple basis for the cycle space. By Kuratowski's theorem, it therefore suffices to show that subdivisions of $K_{3,3}$ and K_5 do not satisfy this criteria in order to prove the converse direction.
Let's do K_5 first. We have that $\dim(\mathcal{C}(K_5) = 6$. Therefore any basis has 6 elements in it. There is an additional element C so that the collection of basis elements and C represents every edge exactly two or 0 times. This collection of 7 cycles contains at least 21 edges, however, there are only 10 edges in K_5, a contradiction.
Similarly for $K_{3,3}$, we know that the cycle space has dimension 4. Given some basis of 4 elements, we have an additional element C so that the collection of basis elements and C represents every edge exactly two or 0 times. The collection of 5 cycles contains 20 edges, because each cycle has length at least 4. But there are 9 edges total, so again creates a counting problem. ■

Ok, so we've finished the algebraic criteria for planarity. One unfortunate thing about this proof is that we have to use Kuratowski's theorem to finish it. Is there a way to prove the planarity condition without this awkward middle step? To come in a later lecture– or maybe a project?

1.8 February 11

1.8.1 Abstract Duality

Recall earlier we constructed dual multigraphs for plane graphs. For this section, we are going to need a slightly more rigorous definition of a multigraph.

Definition 1.8.1 (Multigraph) A multigraph on n vertices is a symmetric $n \times n$ matrix with entries in $\mathbb{N}$.

The coefficient in the i, j spot of the matrix denotes the number of edges that lie between the i and j vertex.

Definition 1.8.2

A *cut* of G is a set of edges between two subgraphs H_1, H_2 with the property that $G = H_1 \sqcup H_2$. A cut set is *minimal* if it contains no smaller cutset in it.

The cut space $\mathcal{C}^* \subset E$ is the is the set of all cuts. In many ways it is better behaved than the cycle space, as it is a space of its own right (we don't need to take span). To prove this observe that

Claim 1.8.3 Cuts of the from $E(v, G \setminus v)$ span the cut space. The cut space has a simple basis.

Proof. Given any cut $E(H_1, H_2)$, let v be a vertex of H_1. Observe that $E(H_1 \setminus v, H_2 \cup v) = E(H_1, H_2) + E(v, G \setminus v)$. Also, notice that edges belong to at most 2 sets of the form $E(v, G \setminus v)$. ■

The cut space serves as a natural dual to the cycle space.

Claim 1.8.4 With the standard inner product on $\mathcal{E}$, we have that $(C)^\perp = C^*$.

Proof. To show that $\mathcal{C} \subset (C^*)\perp$, it suffices to show that every cycle shares an even number of edges with each cut. If you think of a cut describing the edges that go between H_1 and H_2, it is clear that a cycle must cross those edges an even number of times. So the inclusion is proved.
Now to show that $\mathcal{C}^* = \mathcal{C}^\perp$, we need to show that $E \notin \mathcal{C}$ implies that $E\dot{F} \neq 0$ for some cut set F. Now, if E is not a sum of cycles, then there is a vertex which shares an odd number of edges with E. This is because every cycle has even degree on every vertex. Now this vertex has a cut set F which has $E\dot{F} = 1$. This shows the other inclusion. ■

Claim 1.8.5 The minimal cuts generate the cut space.

Proof. First, notice that a cut is minimal if and only if the components H_1 and H_2 are connected.
Let a partition H_1, H_2 be given. The cut set $E(H_1, H_2)$ is the disjoint union of the edge sets $E(H_1, H_2^k)$ over all connected subsets H_2^k of H_2. So, this is a sum of minimal cuts.

■

So, when can the cut set of some graph be represented by the cycle space of a different graph? Is this correspondence unique?

Definition 1.8.6

We say that two multigraphs G and G^* are *abstract dual* if there is a bijection of the edges so that $\mathcal{C}_G = \mathcal{C}^*_{G^*}$.

Claim 1.8.7 G and G^* planar dual implies that G and G^* are abstract dual.

Proof. Let $F \subset E_G$ be a minimal edge cut that separates G. This creates 2 components. The intersection of all the faces contained in one region and contained in the other region should have the property that each face only contains 2 edges from the minimal cut set– as only 2 edges are required separate the cycle which is each face. This means that we have a collection of faces which all have degree 2 in the dual. This gives us a collection of cycles. Of course, it can't be a collection of cycles, as this would cause various separate regions by the Jordan curve theorem. ■

The other direction is clear. So this means that G and G^* planar dual imply that G and G^* are abstract dual, and G and G^* are abstract dual. In other words, the abstract dual of an abstract dual is itself.
In fact, abstract duality is equivalent to planarity.

Theorem 1.8.8
Whitney

A graph is planar if and only if it has an abstract dual.

Proof. The forward direction we just proved– the dual multigraph is a candidate dual. For the other direction, notice that the cut space of G^* has a simple basis, therefore the cycle space of G has a simple basis, therefore G is planar. ■

1.9 February 12

1.9.1 Coloring II

Lemma 1.9.1 Every planar graph contains a vertex of degree 5 or fewer.

Proof. We have the following inequality from before:

$$|E| \leq 3|V| - 6$$

Assume that every vertex has degree 6 tells us $|E| \geq 3|V|$. This is a contradiction. ■

Corollary 1.9.2 Every graph is 6 colorable.

Proof. By induction. Remove a vertex of degree 5. The rest of the graph is 6 colorable. Since the vertex has degree 5, there is nothing that prevents us adding in the vertex and coloring it the remaining color. ■

We use this to help prove the 5 color theorem.

Theorem 1.9.3 Every graph is 5 colorable.

Proof. By induction. Select a vertex v of degree 5, and color the remaining graph. By our induction hypothesis, the remainder of the graph is 5 colorable. We may assume that the neighbors w_i of v are colored $1, 2, 3, 4, 5$, and let's index them by these colors cyclically around v.

- Suppose there is not an i, j colored path between v_i and v_j . Then v_i and v_j belong to separate connected components of the subgraph colored i, j, therefore in one component we can switch their colors, This reduces the colors in the neighborhood of v to 4 colors, so we can color the graph with 5 colors.
- The above must be the case, as a 13 path cannot intersect a 24 path, but the existence such paths would create a topological $K_{3,3}$ within our graph, contradicting planarity.

■

There is of course the infamous 4 color theorem:

Theorem 1.9.4 Every planar graph is 4 colorable.

This theorem was historically broken by a computer, and we will not go into the proof of it here. However, it is a good time to introduce an equally hard theorem which is equivalent to it:

Definition 1.9.5 A *snark* is a 2-edge-connected graph with every vertex degree 3 and *edge coloring* requiring 4 colors.

It is obvious that every edge coloring requires at least 3, and no more than 4 colors.

Theorem 1.9.6 Every Snark contains the Peterson graph as a minor.

Proof of Four color theorem. We want to show that non existence of a 4 coloring on G will give non existence of a 3 coloring on this new graph. So given a 3 coloring of the graph H will give me a 4 coloring of G. A 3 coloring of the edges of a triangulation gives us a 4 coloring of the graph. So, this gives us necessarily a snark on dual graph. This means that the dual is non-planar, as every snark contains a Peterson graph and therefore a K_5 as a minor. ∎

Now I want to briefly want to explore coloring graphs on surfaces, which is (surprisingly) completely solved.

Theorem 1.9.7 Let Σ be an oriented surface which is not a sphere of genus g. Then no more than $\left\lfloor \frac{7+\sqrt{1+48g}}{2} \right\rfloor$ colors are required to color any graph on Σ. Even more surprisingly, that many colors are required in general.

Proof. First, we have Euler's inequality for graph, which is

$$|V| - |E| + |F| \geq 2 - 2g$$

which tells us that

$$|E| \leq 3|V| + 6(g-1)$$

Let δ be the minimal degree of a vertex of g. Then

$$\delta|V|/2 \leq |E| \leq 3|V| + 6(g-1)$$

which tells us that

$$(\delta - 6)|V| \leq 12g - 12$$

because $|V| \geq \delta + 1$ we have

$$(\delta - 6)(\delta + 1) \leq 12g - 12$$

which simplifies to $\delta \leq \frac{7+\sqrt{1+48g}}{2}$ giving us a minimal vertex. By the same induction argument as before, we have that every graph is at least this colorable.
The surprising result is that this is result is enough! Here is an example of K_7 drawn on the torus. ∎

It's not that hard to get to this case by construction– in fact, its on your homework.

Exercise 1.9.8 Let $K(g)$ be the size of the largest K_n that can be embedded in a surface of genus g. Find functions $K_-(g) \leq K(g) \leq K_+(g)$ so that $K_-(g) - K_+(g) < 2\sqrt{g}$.

1.10 February 13

1.10.1 Surfaces

We are going to start moving toward the developing surfaces before we move onto the higher dimensional cases.

Definition 1.10.1 A *triangulated surface* is a graph G with a prescribed set of 3-cycles with the following properties:

- Each edge belongs to exactly 2 triangles
- There is associated dual graph given by triangles. Take a vertex and look at all of the associated faces in the dual graph. This forms a cycle.

We denote the set $\Sigma = (V, E, F)$ to represent the triangulated surface.

We can go between triangulated surfaces by subdividing edges.

Definition 1.10.2 Let Σ be a triangulated surface, and e an edge in Σ.
The subdivision of Σ at e is denoted $\Sigma \div e$, and is the surface with an extra vertex in the middle of e, and extra triangles containing the new vertex.
The subdivision of Σ at v with neighbors $N_1 \cup N_2 = N(v)$ is the graph with an additional vertex w, and an edge vw, and edges from v to N_1 and w to N_2. We then add appropriate faces. This is denoted $\Sigma \div (N_1 \cup N_2)$
We say a surface is *orientable* if there exists a cyclic ordering on each triangle so that if T_1, T_2 are two triangles that share an edge, they have opposite ordering of edges.

The problem that we are interested is when do two triangulated surfaces have a common subdivision. The theorem that we want to prove is the following:

Theorem 1.10.3 The type of the surface is determined it's Euler characteristic, and orientability.

Claim 1.10.4 Any topological surface with Euler characteristic 2 is a planar graph.

Missing lots of pictures. We would like to show that a triangulation gives us something like a "simple basis" for the cycle space. This simply follows from an argument on the number of cycles. We know the dimension of the cycle space is given by $E - V + 1$. The Euler characteristic tells us that the dimension of the cycle space is $-\chi(\Sigma) + F + 1 = F - 1$. This means that the set of faces, will be a basis if they are linearly independent, less a face.
So, given some cycle, we need to show that it is written as the sum of these other cycles. We will show a stronger result about surfaces.
First, instead of working with triangulated surfaces, we can work with "rectangulated surfaces", because we can subdivide the faces to give us triangles, and combine triangles to give us quadrilaterals. In a rectangulation, I claim that there is a way to put a flow across the edges so that we see the following flows at every single face:

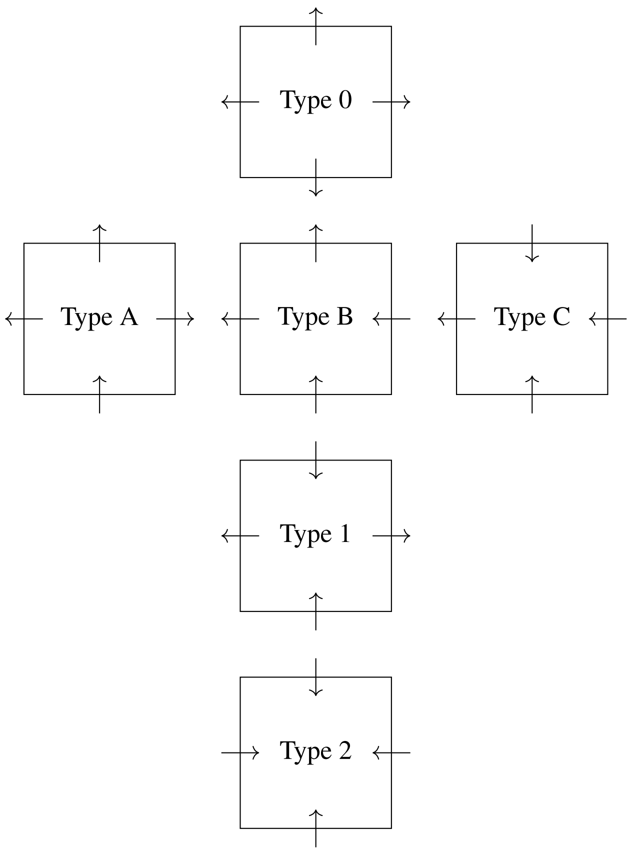

Start by adding a face of type 0. Then a subdivision argument shows that you can always add faces f to the old faces H so that if $v \in f \cap H$, we have that there is an edge $vw \in f \cap H$. In other words, you only join the new face to the old faces at edges. This induces an arrow on some of the edges of f. Make all the other arrows point outward.
This is an algorithm that clearly produces a labeling of all of the faces.

Claim 1.10.5 If the Euler Characteristic is 2, there exists a labeling that does not use faces of type 1.

Notice that adding faces on each step of type A, B, C does not change the Euler characteristic. Adding faces of type 1 lowers the Euler characteristic by 1, and it the only way to do so. You only add a face of type 0 once, while adding faces of type 2 raises the Euler characteristic by one.
Based on this, you know that there is at least one face of type 2 for every face of type 1 you have drawn. Take the last face of type 1 you added, and start making a path along the edge arrows. Whenever you encounter a face of type ABC, travel directly across it. This will give you to a face of type 2. Now reverse every arrow that you walked across. This converts all the type 2 to type C, and converts type 1 to type C. The type ABC faces are permuted.
So, by the fact that the Euler characteristic is 2, we know that there is an arrow labeling of edges that does not use type 1. Now add faces to your diagram in an order induced by the arrow labeling. By adding faces of type A first, then type C, then type B, you at every step add a face which is linearly independent of the remaining faces. Therefore you have $f-1$ linearly independent cycles. This tells us that face space is a basis, and that you have a planar graph. ■

So, where do we go from here? We show that every surface can be cut into a surface of lower degree in a controlled manner by using this kind of decomposition.

1.11 February 18

We want to show that we can cut apart a surface in a controlled manner. Right now we have a labelling every surface in such a way that we have a minimal number of critical points is achieved. Using this Label, we want to show that you can simplify a surface in some way.

Definition 1.11.1 Given 2 surfaces Σ_1 and Σ_2, with cycles of same length $C_i \subset \Sigma_i$, with $\Sigma_i \setminus C_i = \Sigma_i^1 \sqcup \Sigma_i^2$ two disconnected components surfaces with boundary with $\chi(\Sigma_i^1) = 1$, we define the *connected sum* $\Sigma_1 \#_{C_1,C_2} \Sigma_2$ to be the surface $(\Sigma_1^2 \sqcup \Sigma_2^2)/(C_1 \sim C_2)$.

Claim 1.11.2 Connected Sum has the following properties:

- (Dependence on Type 1) The type of $(\Sigma_1 \# \Sigma_2)$ is independent of the cycle chosen.
- (Dependence on Type 2) If $\Sigma_1 \sim \Sigma_1'$ then $(\Sigma_1 \# \Sigma_2) \sim (\Sigma_1' \# \Sigma_2)$
- (Associativity) $(\Sigma_1 \# \Sigma_2) \# \Sigma_3 = \Sigma_1 \# (\Sigma_2 \# \Sigma_3)$
- (Commutativity) $\Sigma_1 \# \Sigma_2 = \Sigma_2 \# \Sigma_1$
- (Identity) If $\chi(\Sigma) = 2$, then $\Sigma \# \Sigma' = \Sigma'$
- (Grading) Define the *Barred Euler Characteristic* $\bar{\chi}(\Sigma) = \chi(\Sigma) - 2$. Then $\bar{\chi}(\Sigma_1 \# \Sigma_2) = \bar{\chi}(\Sigma_1) + \bar{\chi}(\Sigma_2)$. [1]

In particular, the first 5 statements tell us that homeomorphism types of surfaces form an *graded monoid*.

Theorem 1.11.3 (Classification of Surfaces) The Sphere, Torus and Projective Plane generate the monoid of surfaces. In fact this monoid is $\langle T, P \mid PT = PPP, PT = TP \rangle$.

[1] This is not a real thing, but it will be useful for this proof.

Proof. By the work we did in classifying all surfaces of Euler characteristic 2, we know that there is a flow on the surface that uses a minimal number of rectangles of type 1. Let f be the last rectangle type 1. Let H be all of the faces that were added after that rectangle. Because we can reverse the flow, we know that H is a planar graph with boundary. We ask how this boundary glues into the remainder of the surface. Does $H \cup f$ have one cycle in the boundary, or does it have 2 cycles in the boundary?

- If it only has one cycle boundary component, then $H \cup f$ and $G \setminus \{H \cup f\}$ is a decomposition of the surface which both have cycle boundary components. We have removed a projective plane from the space, and thereby decreased the Euler characteristic.
- If it has 2 boundary components, let f' be the next rectangle of type 1. We know that it has flows to each of the two boundary components. Take the faces in the path of those flows and call that H'. Then $H \cup f \cup H' \cup f'$ is connected with boundary which is circle. We have therefore removed an Barred Euler Characteristic -2 piece from our manifold, and have reduced it.

By repeating the above process, we can write every surfaces as the connected sum of spheres, projective planes, and barred Euler characteristic -2 pieces. We simply need to classify these.

1.12 February 20

The barred Euler characteristic -2 pieces are simply the Torus and the Klein Bottle. If you have Euler characteristic -2, then you are composed of two parts– a flow from the bottom of the surface and a flow from the top of the surface. There are 2 cases.

1. Both the top and bottom parts have 2 boundary components.
   (a) You glue these together and the orientations of the cylinders are compatible. In this case, you have a torus.
   (b) You glue these together, and the orientations are non=compatible. In this case, you have a Klein bottle.
2. Both the top and bottom parts have 1 boundary component. In this case, you glue together a Klein bottle.

We need to justify that the Klein Bottle is equal to two 2 projective planes. A good picture can be found here:

`http://people.maths.ox.ac.uk/hitchin/hitchinnotes/Geometry_of_surfaces/Chapter_1`

■

1.12.1 Odds and Ends- Discrete Curvature

In this section we prove a small result of Discrete curvature:

Definition 1.12.1 Let Σ have a linear embedding into $\mathbb{R}^n$, so that every edge is a line segment, and every face is contained in some 2-dimensional subspace. Then at each vertex v, define the *curvature*

$$\kappa(v) = 2\pi \sum_{wvu} \angle wvu$$

as 2π less the measure of each angle that contains v in radians.

Example 1.12.2 In the case of a standard embedding of the cube, $\kappa(v) = 2\pi - 3 \cdot (\pi/4) = \pi/4$.

By taking a sum of the curvature over each vertex, we get the *total curvature*, $\kappa(\Sigma)$.

Theorem 1.12.3 Let Σ be a linearly embedded surface. Then $\kappa(\Sigma) = 2\pi\chi(\Sigma)$.

Proof. Write out the sum $2\pi(V - E + F)$, and think of this as a distribution of weights onto the faces of the surface.

- $2\pi V$ gets distributed by adding an angle worth of weight to each of the faces containing a vertex. This leaves a surplus of $\kappa(v)$ of weight at every vertex.
- $-2\pi E$ gets distributed by adding a weight of $-\pi$ to each face containing a specified edge.
- $2\pi F$ gets distributed by adding a weight of 2π to each face.

Now at each face, we have a weight of $-(n-2)\pi + \sum \angle v$. But $(n-2)\pi$ is the sum of the interior angles of a face with n sides, so this contribution is zero. We are left with the surplus of $\kappa(v)$ weight to distribute. So the sum of the curvatures is equal to 2π times the Euler characteristic. ■

1.12.2 Odds and Ends- Sperner's Lemma

We digress for a moment to talk about the 2 dimensional case of Sperner's lemma, which discusses colorings of triangulations. A triangulation is a network of vertices and edges, where every face of this network is a triangle.

Theorem 1.12.4 (Sperner's Lemma) Let T be a triangulation of $x + y + z = 1$ with $x, y, z \geq 0$ with labeled vertices t_{xyz}. Suppose each vertex is colored either red, green or blue in such a way that

- If $x = 0$, then t_{xyz} is not colored red
- If $y = 0$, then t_{xyz} is not colored blue
- If $z = 0$, then t_{xyz} is not colored green

Then there exists a small triangle in the triangulation, Δ that has the color of all three vertices different.

Proof. We will call the small triangles Δ rooms, which we will usually identify with their bottom left corner. We add a "door" to the edge in the array T if the vertices that bound the edge are blue and red. Then all rooms with just 1 door have all different colors for vertices. Every other room has either no doors or 2 doors. By construction, one side of T contains only blue and red vertices. If a door separates the interior and the exterior of T, then it must lie on this blue-green side.

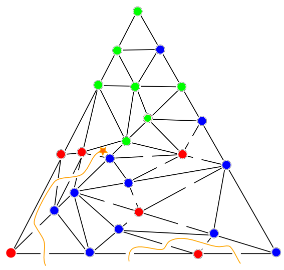

Exercise 1.12.5 Prove that there an odd number of doors that separate the interior of T to the exterior of T.

For each door on the boundary of T, construct a path that travels through T by going through doors, and only going through a door once. This process uniquely determines a path for each initially selected door. This path either terminates in T, or it at some point leaves T through another door. As each path that can both enter and leave T uses 2 boundary doors, there must be a path that is unable to leave T, as the number of doors on the boundary of T is odd. This implies that a path gets "stuck" in T, i.e. it moves into a room with just 1 door. This implies the existence of a red-green-blue room. ■

Using Sperner's Lemma, we prove Brower's fixed point theorem.

Theorem 1.12.6 Every continuous function $f : D^2 \to D^2$ fixes some point.

Proof. We first prove that there is no function $g : D^2 \to \partial D$ that acts as the identity on ∂D. Take a triangulation of D^2, and divide ∂D into three components: Red, green, blue. Then a map $g : D^2 \to \partial D$ gives a coloring to any triangulation of D^2. By Sperner's lemma, there must be a RGB triangle. This means that the vertices of the triangle are mapped to the Red, Green and Blue components respectively. This is impossible, as the map is suppose to be continuous (and therefore uniformly continuous), but an arbitrarily small triangle is mapped to a set of large size.
With this, it is easy to prove that every function $f : D^2 \to D^2$ fixes a point. We prove by contradiction. Suppose that f is a function that has no fixed points. Then define a new function, $H(x,t) : [0,1] \times D^2 \to D^2$ by the following method:

- Suppose that you want to know $H(x,t)$. Draw the points x and $f(x)$.
- Draw a line from $f(x)$ through x. This hits a unique point on the boundary of the circle.
- Define $H(x,t)$ to be that point which is $100 \cdot t\%$ along the line that you have just drawn. When $t = 0$, this should be the identity. When $t = 1$, $H(x,t)$ always gives a point on the boundary circle.

Notice that the second step in this definition crucially uses the fact that $f(x) \neq x$ for all x. Additionally, $h(x)$ is a continuous function. Finally, notice that $H(x,t)$ is a deformation retract to the boundary of the disk. However, we know that there is no deformation retract of the disk to the circle; this is a contradiction! ■

Chapter 2

Simplicial Complexes

2.1 February 23

2.1.1 Basic Definitions

We begin with some basic definitions.

Definition 2.1.1 A *finite abstract simplicial complex* is a set S, and a finite set of subsets $\Delta \subset \mathcal{P}(S)$ which is *downward closed* in the sense that if $X \in \Delta$ and $Y \subset X$, then $Y \in \Delta$. Be sure to include the empty element!

We usually denote a simplicial complex Δ and forego the set S, as we can choose some universal set (such as $\mathbb{N}$) for all of our finite simplicial complexes.

Example 2.1.2 Here are a few abstract simplicial complexes.

- A *graph* is a simplicial complex with the property that $|X| \leq 2$ for $X \in \Delta$. The elements of size 2 are called edges, and the elements of size 1 are called vertices.
- A n-simplex is the simplicial complex with a special largest element $X \in \Delta$ of size $n+1$ with the property that for all $Y \in \Delta$, $Y \subset X$.
- Given any simplicial complex, we have a *geometric realization* of the complex. Choose S with finite size, and consider the affine space $(\mathbb{R}^+)^{|S|}$ with coordinates $\{e_i\}_{i \in S}$. Then to the simplicial complex Δ, we associate the topological space

  $$|\Delta| := \bigcup_{X \in \Delta} \mathrm{Conv}(\{e_i\}_{i \in X})$$

  . This is a topological space. We will avoid working with the geometric realization, because it is not very combinatorial.
- Let S represent the set of all possible edges on n vertices. Let Δ_k be the simplicial complex which includes edge sets which are k colorable. This gives a simplicial complex.
- The above example extends to a lot of different properties of graph. One could take planarity, surface embeddability, etc... It turns out that studying the topological properties of these spaces gives us data on algorithms that try to detect these above properties.
- Given some subset V of a metric space M, δ, define the L *Vietoris-Rips complex* as the set of subsets of V which have diameter less than L. The structure of the Vietoris-Rips complex is an active field of research in applied mathematics, image sensing, and communications.

So, we are in fact looking at something interesting here. I want to develop another language for talking about simplicial complexes.

Definition 2.1.3 A *partially ordered set*, or *poset* for short, is a set P with an order relation $\leq$ satisfying the properties of transitivity, reflexivity, and antisymmetry.

Notice that to every abstract simplicial complex, we can associate a poset, by saying that $Y \leq X$ in $P(\Delta)$ if $Y \subset X$. In this case, we *forget* the empty face. What properties do all of these Posets have?

Definition 2.1.4 A *simplicial poset* is a poset with a minimal element $\hat{0}$ with the property that for all $x \in P$, the interval $[0, x]$ is isomorphic as a poset to some Boolean poset B_n.

It is clear that not every poset is a simplicial poset.

Claim 2.1.5 Simplicial posets are the same as abstract simplicial complexes.

To every poset, we can associate a *order complex* to it, which consists of all chains in the poset. We denote this simplicial complex as $\Delta(P)$.
If P is a poset with a minimal element, let $\bar{P}$ be the poset missing that minimal element.

Exercise 2.1.6 What is the structure of $\Delta(\overline{P(\Delta)})$?

2.2 2/25

2.2.1 Posets

Definition 2.2.1 Let Δ be an abstract simplicial complex on the set S, and $X \in \Delta$ some simplex. Define the simplex $\mathrm{sd}_X(\Delta)$, the subdivision of Δ at x to be the simplicial complex on the $S \cup \{v_X\}$, with simplices of the form

- Simplices Y from Δ which have the property that $X \not\subset Y$.
- Simplices of the form $Y \cup \{v_X\}$, where $Y \in \Delta$, and $Y \cup X \in \Delta$, but Y does not contain X as a subset.

In class, we gave a couple of examples of subdivision. Notice that this definition of subdivision agrees with the definition of subdivision in the case where Δ describes a graph.

Claim 2.2.2 The simplicial complex $\Delta(\overline{P(\Delta)}$ is a subdivision of Δ.

Proof. The algorithm is to start taking subdivisions in order of simplices by their dimension. One needs to show that this is independent of order of subdivision. Notice that this algorithm associates to every simplex of Δ a vertex of $\Delta(\overline{P(\Delta)}$.
Now, given some chain $X_1 < X_2 < \cdots < X_k$ in $\Delta(\overline{P(\Delta)})$, I need to find some simplex which represents this, and the way that I need to find this simplex needs to preserve the ordering of chains. When I do this type of subdivision, I get a vertex for every simplex in Δ. We would like to show that for every chain C in X_1 of length $|X_1|$, there is a unique set $Y_C \in \Delta(\overline{P(\Delta)}$ with $|Y_C| = |X_1|$. We create this set Y inductively as follows:

- We say a chain D is saturated if there is no chain D' with same endpoints such that $D \subset D'$. So in this context, the chain C is saturated.
- $\mathrm{sd}_{X_1} X_1$ has maximal simplices corresponding to saturated chains of length 2. These correspond to the $|X_1| - 2$ dimensional simplices in X_1.
- By induction, each $|X_1| - 2$ dimensional simplex X_2 has the property that all chains of simplices of length $|X_2|$ in X_2 are associated uniquely with a $|X_2| - 1$ simplex in $\Delta(\overline{P(X_2)}$.

- Show that $\Delta(\overline{P(X_2)}) \subset \Delta(\overline{P(X_1)})$.
- Show that every $|X_2|-1$ simplex in $\Delta(\overline{P(X_2)})$ is associated to a unique $|X_1|-1$ simplex in $\Delta(\overline{P(X_1)})$
- Therefore, every chain of length $|X_1|$ can be associated to a unique $|X_1|-1$ simplex.

Once we have the identification of the maximal chains, the smaller chains are identified by subsimplex. ■

A more in depth proof of this can be found in Combinatorial Algebraic Topology by Kozlov, Proposition 2.23.

Definition 2.2.3 The Euler characteristic is the sum $\chi(\Delta) = \sum_{X \in \Delta} -1^{|X|}$.

While this sum makes sense for topological reasons, is there a combinatorial reason that we might study it. Indeed, before mathematicians made the connection of topology to posets, combinatorially had already come up with a function which returns the Euler characteristic in certain cases.

Definition 2.2.4 Given any poset P, define the incidence algebra $I(P)$ to be the functions $f : \{\text{intervals of P}\} \to k$, where k is some field. The operation on the incidence algebra is convolution, which states that given some element p in the incidence algebra,

$$f * g([x,z]) = \sum_{x \leq y \leq z} f(x,y)g(y,z)$$

The reason that this is called convolution is that in the case where your poset looks like n, you have the $\sum_{n \in \mathbb{N}} f(a,n)g(b,n)$, which looks a bit like the convolution sum. These things show up naturally in the study of analytic number theory. Usually in this case, one looks at the poset of numbers under the relation of divisibility. Here are a few examples of interesting functions that exist in every single poset:

Example 2.2.5

- The *delta function* is the function that returns 1 if the interval is trivial, and zero otherwise.
- The ζ function is the function that returns one on every interval. You should think of convolution against ζ as being the same as integration.
- It turns out that the zeta function is invertible! The *Möbius function* $\mu(a,b)$ is defined by the following relation:
$$\mu(x,z) = \begin{cases} 1 & x = z \\ -\sum_{x \leq y < z} \mu(x,y) & x < y \\ 0 & \text{otherwise} \end{cases}$$
Let's do an example of a Mobius function. Then $\mu(a,b) = \mu(b/a)$. On the poset of natural numbers ordered by divisibility, the Möbius function is equal to -1^k, where k is the number of distinct prime divisors, and 0 if n contains a square.

Claim 2.2.6 The delta function is the identity.

Claim 2.2.7 The Mobius function is the inverse of the ζ function.

Proof. We need to show that the $\mu * \zeta(x,z) = 1$ if $x = z$, and 0 otherwise. We prove by induction on the size of the interval.
We need to show that $\sum_{x \leq z \leq y} \mu(x,z) = 1$ is $x = y$, and 0 otherwise. But this follows immediately from our induction, as this is equal to $\mu(z,z)$. ■

2.3 February 27

So where are these names coming from? In the case of number theory, to every function $f : \mathbb{N} \to \mathbb{R}$, we can associate a new function given by series, called it's *Dirichlet L-function* defined by

$$L_f(s) = \sum_{n=1}^{\infty} \frac{f(s)}{n^s}$$

. For instance, to the $\zeta : \mathbb{N} \to \mathbb{R}$ function, we get the infamous Riemann ζ function. Now, multiplication of Dirichlet L-functions gives us the Dirichlet L function associated to the convolution product we had given before. In this case, the statement $\mu * \zeta = 0$ is the same as saying that $1/L_\zeta = L_\mu$. So, studying properties of the Möbius function gives us data on the Riemann Zeta function.
Another place where the Möbius function shows up is the study of *inclusion-exclusion*. The general question is given a collection of objects X, and a set of properties (specified subsets) P on X, it is easy to answer "how many objects satisfy some set of properties $S \subset P$," but hard to answer " how many objects satisfy exactly the properties $S \subset P$ and no others."

- Let $f(S)$ be this first function (Satisfies properties S)
- Let $g(S)$ be the second function (Satisfies only properties S).

Notice that going from g to f is very easy. The sets that satisfy the property S, and possibly more properties, are the sets that satisfy exactly any property T which surrounds S. Think of g as providing a disjoint union count.

$$f(S) = \sum_{S \subset T \subset P} g(T).$$

One can "upgrade" the functions f and g to be elements incidence algebra on Properties ordered by inclusion:

$$f(S,T) := \{\text{Number of elements which satisfy property } T\text{, which also satisfy property } S\}$$

$$g(S,T) := \{\text{Number of elements which satisfy property } T\text{, satisfying exactly property } S\}$$

Then we can recover our original functions by $f(S,X) = f(S)$ and $g(S,X) = g(S)$.
Now applying the theory of incidence algebras to simplify these expressions:

$$f(S,T) = \sum_{S \subset U \subset T} \zeta(S,U) g(U,T)$$

which tells us that we should be able to obtain g from U by Möbius inversion. In this case, the Möbius function is

$$\mu(S,T) = \begin{cases} (-1)^{|T|-|S|} & S \subset T \\ 0 & \text{otherwise} \end{cases}$$

So, the Möbius function is really a Natural thing to study! How does it relate to topology for us?

Claim 2.3.1 If P is a poset with maximal element 1 and minimal element 0, then the value of the Möbius function on $[0,1]$ is $\sum_c -1^{|c|}$, where c are the chains of the poset going from 0 to 1.

Proof. By induction. Let 1 be the maximal element. Then

$$\mu(0,1) = -\sum_{0 \leq x \leq 1} \mu(0,x)$$

But the term on the right counts ever chain with endpoint x which are less than 1. Of course, every such chain gives us a chain with endpoint 1 of length one longer. By induction hypothesis, we get the desired sum. ∎

Let $P(\Delta)$ be the poset of some complex. It comes with a minimal element, and let's append an extra maximal element 1, and call this poset $\widehat{(}P(\Delta))$. Now, the function $\mu(0,1)$ on this poset counts the number of chains. However, the alternating sum of chains in this poset by length gives us the alternating sum of chains by length in the original poset times -1, which gives us the Euler characteristic.

2.4 Evasive Properties

We now take a look at combinatorial and computational called evasiveness. Here is the historical context for the problem:

Definition 2.4.1 Let P be a subset of graphs on n vertices. The property P graph game is a 2 player game. Player two selects a graph G. Player one is allowed to ask "is the edge e_{ij} in G?"? If Player one can determine if G satisfies property P in fewer than $\binom{n}{2}$ moves, then player one wins. We say that a graph property is evasive if there is a graph G which has no winning strategy for player one.

For example, the property of being a complete graph is evasive. Here is a non-evasive property for graphs:

Definition 2.4.2 A graph is *scorpion* if there exists 3 special vertices, the sting, tail and body with the property that the sting is only connected to the tail, and the tail is only connected to the sting and body. The body is connected to every vertex but the sting.

The algorithm first looks for the special vertices. Notice that once we find any of the special vertices, we can identify if a graph is a scorpion in linear time. Take any vertex x.

- If x has degree 0, then we are not a scorpion
- If x has degree 1, then it is either a sting, or its neighbor is the body.
- If x has degree 2, then it is either a tail, or one of its neighbors is the body. This can be checked in linear time.
- If x has degree 3 to $n-3$, then it is not very special.
- If x has degree $n-2$, then it is the body.
- If x has degree $n-1$, then we are not a scorpion.

So, suppose that x has degree 3 to $n-3$. Let B be the neighborhood of x. Let S the complement of B. We know the sting and tail are in S, and the body is in B.

- Pick a vertex b_0 of B. Now go through S looking for stings. Whenever find a vertex connected to b_0, delete it from S. When you find a vertex not connected to b_0, delete b_0 from B.
- Pick a new vertex b_1 from the remaining B, and proceed.

This algorithm terminates when we have removed all the vertices from S or from B. If we remove all the vertices from S, there is no sting, and we are not a scorpion. If we remove the last vertex from B, the last tested vertex in S must have been the sting (as otherwise we did not delete the body from B at some step...) Test this to be the sting.
This algorithm only requires you to check far fewer than $(n-k)(k)/2$ edges, where k is the size of $|B|$.

Definition 2.4.3 We say that a graph property is *monotone* if G satisfies P, and $H \subset G$, then H satisfies P. A graph property is *trivial* if either every graph or no graph satisfies the property.

2.5 March 2

Conjecture 2.5.1 Every monotone non-trivial graph property is evasive.

So how does this relate to topology? If P is a monotone graph property, it describes a simplicial complex for us on the set edges. We now rephrase the Evasiveness problem as follows:

Definition 2.5.2 Let S be the simplex on vertices S, and let Δ a simplicial complex of S which is known to us. Let $X \subset S$ be some simplex. We play a game asking "is the vertex x in $X?''$ If you can solve whether $X \in \Delta$ in fewer than $|S|$ questions for every X, we say that Δ is *non-evasive*, and we say that it is *evasive* otherwise.

Notice that this is equivalent to solving the graph theoretic problem for monotone graph properties.

Theorem 2.5.3 If $\chi(\Delta) \neq 1$, then K is evasive.

In fact, we will prove that every single non-evasive complex is contractible.

Definition 2.5.4 A *guessing algorithm* for property Δ of a set S is a binary tree which where every vertex corresponds to the question "is vertex v_k in my set." On each path , no question may be repeated, and the depth of the tree is exactly $|S|$, so the leaves are identified with $\mathcal{P}(S)$.

Given a guessing algorithm, a *evader* is an element $X \in \mathcal{P}(S)$ with the property that one must ask all k questions to determine if X is in Δ.
Now, notice that every single guessing algorithm induces a pairing on the vertices of the poset of S with edges from the Hasse Diagram. This restricts to a matching on the edges of the poset of Δ. Now suppose that some cell X is a evader. This means that the matching edge from this pairing pairs X with a cell not in Δ! So, existence of evaders is equivalent to an incomplete pairing on Δ.
Now, suppose that you have an algorithm that is has no evaders. This tells you that you have a pairing on the Hasse Diagram of $P(\Delta)$. This means that each d face is paired with a $d+1$ face, which means that the Euler characteristic is 1. (Recall, the empty set gets paired here too!)
This proves the theorem.
In fact, we can do slightly better than this. We can prove that if K is non-evasive, it is *homotopy equivalent* to a point.

Definition 2.5.5 Let Δ be a simplicial complex, and $X, Y \in \Delta$ simplices with the property that $Y \subset X$. Suppose that there is no other simplex containing Y. The *simplicial collapse* of Y along X is the complex $\Delta \setminus X$.

The idea is that if a k-simplex has a face that belongs to no other simplex, we can collapse it by pushing it in. Here is the basic idea:

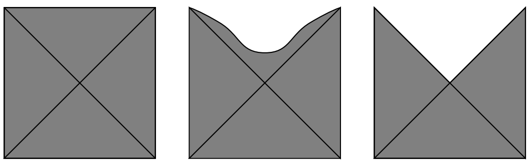

Definition 2.5.6 We say that Δ is contractible if it simplicially collapses to a point. We say that Δ_1, Δ_2 are *simplicially homotopic* if they have a common simplicial collapse.

Notice that a simplicial collapse of X along Y states that the dimension of Y is one less than the dimension of X. So, we can think of a simplicial collapse as dictating some edges on the Hasse diagram. But how can we think of these edges on the Hasse diagram as giving us a simplicial collapse?

Definition 2.5.7 A set of edges $F \subset E$ in the Hasse diagram are called a *acyclic partial matching* if

No two edges in F share a vertex.

If you orient all the edges of E upwards, and all the edges of $F \setminus E$ downward, there are no directed cycles in the poset.

Clearly, every single contraction describes a acyclic matching. One can show that every single acyclic matching gives you some kind of contraction as well, if you would like. We prove the following theorem:

Theorem 2.5.8 Every acyclic complete matching describes a contraction to a point.

Proof. We induct on the size of the simplex. A complete acyclic matching means that every simplex is paired. In particular, a maximal simplex X is paired to Y. We claim that Y does not belong to any other simplex X'.
Suppose for contradiction that Y belongs to X'. Pick any Then we have the following picture in the Hasse Diagram:

$$\begin{tikzcd} X & & X' \arrow[dl] \\ & Y \arrow[ul] & \end{tikzcd}$$

We know that we can keep on travelling " down- up" from here on out, as every maximal simplex has a unique arrow in the wrong direction. This eventually terminates, and gives us a simplex which we can simplicial collapse upon.

■

So, it remains to show that every single guessing algorithm gives us an acyclic matching on the poset! With this, we would be done.

2.6 March 4

We can now prove that every non-evasive complex is contractible.

Proof. We claim that every non-evasive complex has a free face which can be collapsed. Let $X \in \Delta$ be some maximal simplex. Suppose that X has no free face. Since every simplex is paired, X is paired with some simplex Y less than it. This simplex is contained in X_1, another maximal simplex, which is paired with Y_2, and so on. Create a sequence in this fashion. The sequence either terminates and we find a free face, or it winds up back at X_1. Think about how these are laid out on a binary tree. Say that the "no" answers to each question are on the left, and the "yes" answers on the right. Then we have that each X_i, X_{i+1} differ by a single yes-no question. These questions must be in the negative in each case, and so you move only in one direction on the tree as you go along this chain.

■

Today, we move onto simplicial homology, one of the staples of algebraic topology.
When we talked about graph theory, it was pretty useless to ask "how many cycles were in my space." It was, however, useful to talk about things like "the dimension of the cycle space." We want to create something that mimics this in the simplicial case. However, we first need to make sense of what the cycle space should be! We need a little bit of fanciness to get these definitions to work out. Let's start by putting an order on S, the vertex set of Δ some simplicial complex.

Definition 2.6.1 Let R be a ring and Δ a simplicial complex. Define $C_k(\Delta)$ to R-module generated on a basis of the k-simplices of Δ. For each $X \in \Delta$ with $|X| = k + 1$, let e_X be the corresponding basis vector.

What should the "cycles" of this space be? They should be ways that we can glue together the simplices to get something which "has no boundary" in some sense.

Definition 2.6.2 Let Δ be a simplicial complex, and $C_k(\Delta)$ be as before. Define the *differential* $d_k : C_k(\Delta) \to C_{k-1}(\Delta)$ on the basis e_X

$$d_k(e_X) = \sum_{v \in X} -1^{\operatorname{ord}_X x} e_{X \setminus x}$$

where $\operatorname{ord}_X v$ is the position of v in the order of vertices inside of Δ.

Because this minus sign is very annoying, we will usually work with $R = \mathbb{Z}/2\mathbb{Z}$. We will also drop the index on d all of the time.

One can think of d as measuring how much "boundary" a set has. When a set has zero boundary we call it a "cycle."

Definition 2.6.3 Let $v \in C_k(\Delta)$ be an element with $dv = 0$. Then we say that v is a *cycle*. The space of cycles is denoted $Z_k(\Delta)$.

Now what about the image of the map d. d measures how much boundary a set has by actually returning its boundary!

Definition 2.6.4 Let $v \in C_k(\Delta)$ be an element such that $v = dw$ for some $w \in C_{k-1}(\Delta)$. Then we say that w is a *boundary*. The space of all boundaries is denoted $B_k(\Delta)$.

Naturally we should expect that all boundaries are, themselves, cycles. This equates to proving that $d^2 = 0$.

Claim 2.6.5 $d_{k-1}d_k = 0$.

Proof. Let $e_X \in C_k(\Delta)$ be some basis vector. Then we have that

$$\begin{aligned} d_{k-1}d_k(e_X) =& d_{k-1}\left(\sum_{x \in X} -1^{\operatorname{ord}_X x} e_{X \setminus x}\right) \\ =& \sum_{x \neq y \in X} -1^{\operatorname{ord}_{X \setminus x} y} e_{X \setminus \{x,y\}} \end{aligned}$$

Now, notice that summing over $x \neq y$ means that each subsimplex $e_{X \setminus \{x,y\}}$ is counted twice. The order on the signs though is different– in one, x is counted first, then y, and in the other y is removed first, then x. So, the sign convention $-1^{\operatorname{ord}_X}$ will cause these to cancel out. Therefore $d^2 = 0$. ∎

This proof implies that the space of cycles contains the space of boundaries – $B_k \subset C_k$. A question that we might have is – how much does the space of boundaries fail to be the space of cycles?

Example 2.6.6 An example where the space of boundaries is not the space of cycles

2.7 March 6

The structure that we have just set up is called a *chain complex.*

Definition 2.7.1 A *chain complex* over a ring R is a sequence of R-modules C_k along with maps $d_k : C_k \to C_{k-1}$ satisfying $d^2 = 0$.

We are interested in measuring the difference between the dimension of the space of cycles and the space of boundaries.

Definition 2.7.2 Given any chain complex C_k, the kth homology groups are defined by the quotient $H_k := Z_k/B_k$.

The homology groups of a simplicial complex can be thought of measuring the number independent k dimensional holes that are in your space– the cycles which cannot be filled with shapes that give their boundary. In the case where R is a field, each of these homology groups is a vector space, and we can measure it's dimension.

Definition 2.7.3 If R is a field of characteristic 0, then the quantities $\dim(H_k) = b_k$ are called the *Betti Numbers* of the space.

We delay proofs of the following facts until next week. But they are important facts!

- The homology groups of subdivided spaces are isomorphic.
- If Δ simplicially collapses to Δ', they have the same homology groups.

Let's compute a bunch of homology groups now.

Example 2.7.4 The homology of the n-simplex is

$$H^0(\Delta_n) = 1$$

$$H^k(\Delta_n) = 0 \quad k \neq 0$$

- In the case that $n = 0$, this is trivial.
- In every other case, we have a simplicial collapse to the point.

Example 2.7.5 The homology of the sphere, S^n. Here, every k cycle for $k \neq n$ restricts to a cycle in Δ, so we have that for $0 < k < n$, $H^k(S^n) = 0$. What about for n? The only thing that could possibly be in homology is the n cycles of Δ_n. There is only one element here, and it was canceled out by the boundary of the top simplex. So $H^n(S^n) = 1$.

Example 2.7.6 The Torus. Here, we draw the torus with 18 triangles, and a bunch of edges. We use a combination of topological and geometric insight here. We also compute over $\mathbb{F}_2$.

- For vertices, notice that every pair of vertices is connected by a path. Therefore $H^0(T) = R$.
- For faces, notice that the boundary of some set of faces is 0 if and only every edge is counted twice. Therefore, the only element of homology is the one that uses all of the faces. This is the same as saying that $H^2(T) = R$
- The middle computation is a pain, but we can use algebra instead. There are 27 edges. The image of the differential must be 8 dimensional to give the proper dimension of homology in H^0. Therefore, the kernel is 19 dimensional. There are 18 faces. One of them is in the kernel, therefore the image of the differential into $C^1(T)$ is 17 dimensional. This means that the homology is 2 dimensional.

We can do a similar argument over $\mathbb{Z}$. This gives us the same homology groups.

Example 2.7.7 The homology of the Klein Bottle. Notice before that we were not able to distinguish the Klein Bottle and the Torus via the Euler Characteristic. However, we are able to distinguish them using homology. The Klein Bottle is also an example where it matters what you choose to use for coefficients.

2.8 March 9th

Today, we take a break from the topological intuition and explore the powerful tools of homology. Recall the following definition:

Definition 2.8.1 A *chain complex* over a ring R is a sequence of R-modules C_k along with maps $d_k : C_k \to C_{k-1}$ satisfying $d^2 = 0$.

These chain complexes come with a sequence of groups, $H_k(C_\bullet) := \ker d_k / \operatorname{Im} d_{k+1}$, called the *homology groups*.
Chain complexes come with maps between them, just like vector spaces. Of course, these maps come with additional structure.

Definition 2.8.2 A *map of chain complexes* is a sequence of maps $f_k : B_k \to C_k$ which respect the differential in the sense that $d_k^C f_k = f_{k-1} d_k^B$.

This is usually written by saying that the following diagram of maps commutes:

$$\begin{array}{ccccccccc}
\cdots & \xrightarrow{d} & B_k & \xrightarrow{d} & B_{k-1} & \xrightarrow{d} & B_{k-2} & \xrightarrow{d} & \cdots \\
 & & \downarrow{\scriptstyle f_k} & & \downarrow{\scriptstyle f_{k-1}} & & \downarrow{\scriptstyle f_{k-2}} & & \\
\cdots & \xrightarrow{d} & C_k & \xrightarrow{d} & C_{k-1} & \xrightarrow{d} & C_{k-2} & \xrightarrow{d} & \cdots
\end{array}$$

Maps of chain complexes are well behaved with respect to composition.

Claim 2.8.3 Let $f : A_k \to B_k$, and $g : B_k \to C_k$ be maps of chain complexes. Then $gf : A_k \to C_k$ is a map of chain complexes.

Proof. Here, we repress the index k.

$$d^C \circ (gf) = (g \circ d^B) \circ f = (gf) \circ d^A$$

■

Perhaps the most important aspect of maps between chain complexes is that they give us an induced map on homology.

Proposition 2.8.4 Let $f : A_k \to B_k$ be a map of chain complex. Then we have an *induced map* $f_* : H_k(A_\bullet) \to H_k(B_\bullet)$.

Proof. We need to construct such a map. Our candidate definition is that $f_*[a] = [f(a)]$. To show that this is well defined, we need to show that $f(a)$ is actually represents a class of homology of B_k, and that this definition is independent of the choice of representative for homology.

- To show that is a class of homology, notice that $d^B(f(a)) = f(d^A(a)) = f(0) = 0$. We have that $d^A(a) = 0$ because a belongs to the homology of A. Therefore, $f(a)$ is a member of homology of B_k.
- To show that this is independent of class chosen, suppose that $[a'] = [a]$. Then it is the case that $a - a' = dv$ for some v.

$$\begin{aligned}
f_*[a] - f_*[a'] =& [fa - fa'] \\
=& [f(a - a')] \\
=& [fdv] \\
=& [dfv] = 0
\end{aligned}$$

■

As maps between chain complexes are maps between R modules, they have the notions of kernels and images as well.

2.9 March 11

Today, we prove a number of results in homological algebra.

Definition 2.9.1 A *exact sequence* is a sequence of maps and objects $f_i : A_i \to A_{i-1}$ such that $\ker f_i = \operatorname{Im} f_{i+1}$ for every i.

Notice that every exact sequence gives us a chain complex, however the converse is usually false. One can think of the homology groups of a chain complex as giving us a measure of how much a sequence fails to be an exacts sequence.

Theorem 2.9.2 Zig-Zag lemma

Let $A, d_\bullet$, $B, d'_\bullet$ and $C, d''_\bullet$ be chain complexes. Given

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

a short exact sequence, there exists a unique map δ such that the following is a long exact sequence on homology:

$$\cdots \xrightarrow{g_*} H_{n+1}(C) \xrightarrow{\delta} H_n(A) \xrightarrow{f_*} H_n(B) \xrightarrow{g_*} H_n(C) \xrightarrow{\delta} H_{n-1}(A) \xrightarrow{f_*} \cdots$$

Proof. First, let's expand the original diagram:

$$\begin{tikzcd}
 & \vdots \arrow[d, "d"] & \vdots \arrow[d, "d'"] & \vdots \arrow[d, "d''"] & \\
0 \arrow[r] & A_{n+1} \arrow[r, "f_{n+1}"] \arrow[d, "d_{n+1}"] & B_{n+1} \arrow[r, "g_{n+1}"] \arrow[d, "d'_{n+1}"] & C_{n+1} \arrow[r] \arrow[d, "d''_{n+1}"] & 0 \\
0 \arrow[r] & A_n \arrow[r, "f_n"] \arrow[d, "d_n"] & B_n \arrow[r, "g_n"] \arrow[d, "d'_n"] & C_n \arrow[r] \arrow[d, "d''_n"] & 0 \\
0 \arrow[r] & A_{n-1} \arrow[r, "f_{n-1}"] \arrow[d, "d_{n-1}"] & B_{n-1} \arrow[r, "g_{n-1}"] \arrow[d, "d'_{n-1}"] & C_{n-1} \arrow[r] \arrow[d, "d''_{n-1}"] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

We want to construct a function δ from $H_n(C)$ to $H_{n-1}(A)$. Select a $[c] \in H_n(C)$. Then lift it to $c \in C_n$. Take an element b in the preimage of c along g_n. Then as $d''_n(g_n(b)) = d''_n(c) = 0$, we have by commutivity of the diagram that $g_{n-1}(d'_n(b)) = 0$. Therefore, $d_n(b) \in \ker g_{n-1}$. However, $\ker g_{n-1} = \operatorname{Im} f_{n-1}$ by exactness of the rows. There therefore exists unique $a \in A_{n-1}$ such that $f_{n-1}(a) = d_n(b)$. We choose the representative of a in $H_{n-1}(A)$ to be $\delta(c)$.

Now we have to show that this a is independent of choice of b. Suppose instead we choose b' in the preimage of c along g_n. Let $d'_n(b')$ lift instead to a'. As b' is a lift of c, we have that $g_n(b) = g_n(b') = c$. Then $b - b' \in \ker(g_n)$. By the exactness of the rows, there exists a unique $\bar{a} \in A_n$ such that $f_n(\bar{a}) = b - b'$. Then the lifts of $d'_n(b)$ and $d'_n(b')$ in A_{n-1} differ by $d_n(\bar{a})$. Therefore, $a - a' = d_n(\bar{a}) \in \operatorname{Im} d_n$. Therefore, a and a' represent the same class in the homology $H_{n-1}(A)$. The function δ is now shown to be well defined on homology and is the unique class

$$\delta(c) = f_{n-1}^{-1} d'_n g_n^{-1}(c)$$

■

Theorem 2.9.3 5-lemma

Given that the rows of the following diagrams are exact, and that the dashed maps are isomorphisms:

$$\begin{tikzcd} A \arrow[r] \arrow[d, dashed] & B \arrow[r] \arrow[d, dashed] & C \arrow[r] \arrow[d] & D \arrow[r] \arrow[d, dashed] & E \arrow[d, dashed] \\ A' \arrow[r] & B' \arrow[r] & C' \arrow[r] & D' \arrow[r] & E' \end{tikzcd}$$

Then C is isomorphic to C'.

Proof. We omit proof in these notes. ∎

Notice that the five lemma tells us that if we have a short exact sequence, and the two outside maps are isomorphisms, then the inside map is an isomorphism. On homology, it tells us if any 2 maps are quasi-isomorphisms, then the third map is a quasi-isomorphism.

2.10 March 13

Definition 2.10.1 Let A, d and B, d' be chain complexes. Let $f : A \to B$ and $g : A \to B$ be chain maps. Then we say that f is **chain homotopic** to g if there exists a series of maps (called a **chain homotopy**) $h^i : A^i \to B^{i+1}$ such that

$$f - g = d'h^i + h^{i-1}d$$

We write that $f \sim g$.

Here's a diagram that helps visualize the maps involved in a chain homotopy.

$$\begin{tikzcd} \cdots \arrow[r, "d"] & A_{i+1} \arrow[r, "d"] \arrow[d, dashed, shift right, "g"'] \arrow[d, dashed, shift left, "f"] \arrow[dl, "h_{i+1}"'] & A_i \arrow[r, "d"] \arrow[d, dashed, shift right, "g"'] \arrow[d, dashed, shift left, "f"] \arrow[dl, "h_i"'] & A_{i-1} \arrow[r, "d"] \arrow[d, dashed, shift right, "g"'] \arrow[d, dashed, shift left, "f"] \arrow[dl, "h_{i-1}"'] & \cdots \arrow[dl, "h_{i-2}"'] \\ \cdots \arrow[r, "d'"'] & B_{i+1} \arrow[r, "d'"'] & B_i \arrow[r, "d'"'] & B_{i-1} \arrow[r, "d'"'] & \cdots \end{tikzcd}$$

It is easy to check that chain homotopy is an equivalence relationship on the set of morphism. This next claim shows the usefulness of chain homotopic:

Lemma 2.10.2 Let $f : A \to B$ and $g : B \to A$ be chain maps. Suppose that $g \circ f \sim 1_A$ and $g \circ f \sim 1_B$. Then f and g are quasi-isomorphisms.

Proof. Let's start with a diagram.

$$\begin{tikzcd} \cdots \arrow[r, "d'"] & A_{i+1} \arrow[r, "d"] \arrow[d, dashed, "f"] \arrow[dd, bend left] \arrow[ddl, dashed, "h"'] & A_i \arrow[r, "d"] \arrow[d, dashed, "f"] \arrow[dd, bend left] \arrow[ddl, dashed, "h"'] & A_{i-1} \arrow[r, "d"] \arrow[d, dashed, "f"] \arrow[dd, bend left] \arrow[ddl, dashed, "h"'] & \cdots \\ \cdots \arrow[r, "d'"] & B_{i+1} \arrow[r, "d'"] \arrow[d, dashed, "g"] & B_i \arrow[r, "d'"] \arrow[d, dashed, "g"] & B_{i-1} \arrow[r, "d'"] \arrow[d, dashed, "g"] & \cdots \\ \cdots \arrow[r, "d'"] & A_{i+1} \arrow[r, "d"] & A_i \arrow[r, "d"] & A_{i-1} \arrow[r, "d"] & \cdots \end{tikzcd}$$

The homotopy to the identity map gives us that there exists h such that $g \circ f - 1_A = dh^i + h^{i-1}d$. Suppose that $v \in H_i A$. Then v is in the kernel of d, so $h^{i-1}d(v) = h^{i-1}(0) = 0$. We have that therefore $g \circ f - 1_A \in \text{ Im }(d)$, which is to say that on homology $g \circ f = 1_A$, as we mod out by $\text{ Im }(d)$ when we take homology.
Of course, a similar proof shows that $f \circ g = 1_A$ ∎

Therefore, to show that a morphism is a quasi-isomorphism, one can just show that it is homotopic to the identity. As homotopic form an equivalence relationship on the set of morphism, we can quotient out by homotopic to arrive a new category, the homotopy category of chain complexes$C$. In this category, $f = g$ if $f \sim g$; in particular, maps who are isomorphisms in the homotopy category are quasi-isomorphisms in the original chain category. We call this category $\mathcal{K}(C)$.
Why do we work with homotopy? It turns out that if Δ' is a simplicial collapse of Δ, then there is a map from $C(\Delta)$ to $C(\Delta')$ which is homotopic to the identity.

Theorem 2.10.3 Suppose that Δ' is a simplicial collapse of Δ. Then Δ and Δ' have the same homology.

Proof. We have a natural inclusion i of Δ' into Δ, which has a right inverse π by projection. We need to show that the compositions πi and $i\pi$ are homotopic to the identity.
The composition πi is clearly homotopic to the identity, as it is the identity on δ.
Suppose that this is the collapse of X where Y is the free face. Then define $h(X) = 0$, $h(Y) = X$, and $h(dY) = dX - Y$, and 0 everywhere else. This is a homotopy! ■

With this, we can at least check that all subdivisions of D^k, the k complex, have the same homology. Now, we bring everything back to the topology. On your homework, you showed that simplicial maps induced chain maps. We now use this, and the tricks from homology that we have already uncovered, to make classify lots of different topological spaces.

Theorem 2.10.4 *Mayer-Vietoris Sequence* Suppose Δ is a simplicial complex, and Δ_1, Δ_2 are two subcomplexes with $\Delta = \Delta_1 \cup \Delta_2$. Then there is a long exact sequence of homology:

$$\cdots \to H_k(\Delta_1 \cap \Delta_2) \to H_k(\Delta_1) \oplus H_k(\Delta_2) \to H_k(\Delta) \to H_{k-1}(\Delta_1 \cap \Delta_2) \to \cdots$$

Proof. Apply the Zig-Zag lemma to the exact sequence

$$0 \to C_k(\Delta_1 \cap \Delta_2) \to C_k(\Delta_1) \oplus C_k(\Delta_2) \to C_k(\Delta) \to 0$$

where the first arrow is by inclusion into both components, and the second arrow is signed inclusion into the union. ■

This shows us really that we can do a lot of computations!

Example 2.10.5 The Sphere, again. Notice that the sphere S^n decomposes into U, V which are discs and $U \cap V = S^{n-1}$. By applying the Mayer Vietoris Sequence to this, we can compute the homology of the sphere.

2.11 March 16

Example 2.11.1 There is no continuous map $f : D^n \to \partial D^n$ which acts by identity on the boundary. If so, the identity map of the sphere would factor through the disk. But the sphere has nontrivial homology in the $n - 1$ place, so this cannot happen.

Example 2.11.2 Let's compute the $\mathbb{Z}/2$ homology of $\mathbb{RP}^n$. We claim that it is $H_k(\mathbb{RP}^n) = \mathbb{Z}/2$ for every $n \le k$. We represent $\mathbb{RP}^n$ as a simplicial complex by taking the octahedron S^n and quotienting out by the relation $X \sim -X$. We can cover $\mathbb{RP}^n$ by the quotient of the upper hemisphere, and the quotient of the equator, which is $\mathbb{RP}^{n-1}$. (In reality, take a slightly smaller version of the upper hemisphere, and a slightly fattened $\mathbb{RP}^{n-1}$ equator.) Their intersection is the boundary sphere S^{n-1} of the upper hemisphere.

This gives us a long exact sequence

$$\begin{tikzcd}
H_n(S^{n-1}) \arrow[r] & H_n(U) \oplus H_n(\mathbb{RP}^{n-1}) \arrow[r] & H_n(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
H_{n-1}(S^{n-1}) \arrow[r] & \cdots \arrow[r] & H_2(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
H_1(S^{n-1}) \arrow[r] & H_1(U) \oplus H_1(\mathbb{RP}^{n-1}) \arrow[r] & H_1(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
H_0(S^{n-1}) \arrow[r] & H_0(U) \oplus H_0(\mathbb{RP}^{n-1}) \arrow[r] & H_0(\mathbb{RP}^n) \arrow[r] & 0
\end{tikzcd}$$

We have some known replacements that we can make from our knowledge of homology of disks (U) and the induction hypothesis.

$$\begin{tikzcd}
0 \arrow[r] & 0 \oplus 0 \arrow[r] & H_n(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
\mathbb{Z}/2 \arrow[r] & \cdots \arrow[r] & H_2(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
0 \arrow[r] & 0 \oplus \mathbb{Z}/2 \arrow[r] & H_1(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
\mathbb{Z}/2 \arrow[r, "i_0 \oplus j_0"] & \mathbb{Z}/2 \oplus \mathbb{Z}/2 \arrow[r] & H_0(\mathbb{RP}^n) \arrow[r] & 0
\end{tikzcd}$$

Since $\mathbb{RP}^n$ is connected, the right map is an isomorphism. The inclusion i_0 is an isomorphism. The map j_0 is multiplication by 2, because the equator is wrapped around 2 times when it is included in the upper hemisphere. Since we are working in $\mathbb{Z}/2$ coefficients, this is 0.

$$\begin{tikzcd}
0 \arrow[r] & 0 \oplus 0 \arrow[r] & H_n(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
\mathbb{Z}/2 \arrow[r] & \cdots \arrow[r] & H_2(\mathbb{RP}^n) \arrow[dll, "\delta"] \\
0 \arrow[r] & 0 \oplus \mathbb{Z}/2 \arrow[r, "\pi_1"] & H_1(\mathbb{RP}^n) \arrow[dll, "\delta_1"] \\
\mathbb{Z}/2 \arrow[r, "1 \oplus 0"] & \mathbb{Z}/2 \oplus \mathbb{Z}/2 \arrow[r, "(0,1)"] & \mathbb{Z}/2 \arrow[r] & 0
\end{tikzcd}$$

Now, the bottom right map $(1 \oplus 0)$ is injective. This means that it has no kernel. Therefore, the map δ_1 has trivial image, which means that the map π_1 is surjective. We also have that π_1 is injective, by the long exact sequence. So this means that $H_1(\mathbb{RP}^n) \simeq H_1(\mathbb{RP}^{n-1}) = \mathbb{Z}/2$.
This argument continues all the way up the diagram, until the top row. On the top row, exactness clearly gives us $H_n(\mathbb{RP}^n) = H_{n-1}(S^{n-1})$.

Theorem 2.11.3
Simplicial Homolgy

The simplicial homology of a simplicial complex is independent subdivision.

Proof. Suppose that we want to show that a complex Δ and a subdivision Δ' have the same simplicial homology. We will prove by induction on the size Δ. The base case is for subdivisions of simplices, which we proved using simplicial collapse and homotopy.
Take any simplex $X \in \Delta$. Let $s(X)$ denote the set of $|X|$-simplices in Δ which come from X in the

subdivision. Then we have a map

$$f : C_\bullet(\Delta) \to C_\bullet(\Delta')$$
$$e_X \mapsto e_{s(X)}$$

which is a chain map. We would like to show that f is a quasi-isomorphism.
Now for the induction step. Pick any maximal dimension simplex Y in Δ. Let $\bar{Y}$ be the subcomplex generated by Y in Δ. By induction, we know that $\Delta \setminus Y$ has the same homology as $\Delta \setminus \mathrm{sd}(Y)$. We have the Mayer-Vietoris decompositions:

$$\begin{tikzcd}
0 \arrow[r] & C(\bar{Y} \cap \Delta) \arrow[r] \arrow[d, "f|_{\bar{Y} \cap \Delta}"] & C(Y) \oplus C(\Delta \setminus Y) \arrow[r] \arrow[d, "f|_Y \oplus f|_{C \setminus Y}"] & C(\Delta) \arrow[r] & 0 \\
0 \arrow[r] & C(\overline{s(Y)} \cap \Delta') \arrow[r] & C(s(Y)) \oplus C(\Delta' \setminus C(s(Y)) \arrow[r] & C(\Delta') \arrow[r] & 0
\end{tikzcd}$$

Now the first map and second maps are quasi-isomorphisms by the induction hypothesis. By the 5 lemma, the last map must also be a quasi-isomorphism. ■

Definition 2.11.4 Let $f : S^n \to S^n$ be a simplicial map. Since $f_* : H_n(S^n) \to H_n(S^n)$ must be of the form $x \mapsto nx$, we say that n is the *degree* of the map f.

There is in fact a way to generalize this to manifolds in general, but we have to define what a manifold is. Studying the degree of maps will have interesting implications for us topologically. Topology will give tell us things about combinatorics. Here are two problems that we are interested in exploring:

Example 2.11.5 Suppose that you have a necklace, with $k \cdot n$ beads on it. The beads come in t different colors, and each color occurs $k \cdot a_i$ times, where $1 \leq i \leq t$. k thieves would like to split the necklace fairly, which means that each thief gets the same number of beads of each color. They would like to split the necklace using the minimal number of cuts possible. Clearly, we can construct an example that uses $(k-1)t$ cuts. What is the worst and best case scenario for the number of cuts that they will need.

Example 2.11.6 A *Kneser graph* $KG_{n,k}$ is the graph whose vertices correspond to k elements subsets of an n element set. Two vertices are connected if their corresponding subsets are disjoint. What is the chromatic color of this graph?

2.12 March 18

Both of these proofs will depend on a famous theorem, the Borsuk-Ulam Theorem.

Theorem 2.12.1 (Borsuk-Ulam) Suppose a map $f : S^n \to S^n$ has the property that $f(x) = f(-x)$ for every x. Then the degree of f is odd.

Proof. We use a special exact sequence related to double covers with coefficients in $\mathbb{Z}/2$. Consider the exact sequence :

$$0 \longrightarrow C_k(\mathbb{RP}^n) \xrightarrow{\tau} C_k(S^n) \xrightarrow{\pi_*} C_k(\mathbb{RP}^n) \longrightarrow 0$$

where π_* is the map induced by the projection, and τ is a lifting map defined by $\tau_{e_X} = \sum_{Y \in \pi^{-1}(X)} e_Y$. Because we work in $\mathbb{Z}/2$ coefficients this is an exact sequence.
Being odd will be the same as saying that f is an isomorphism on the $Z/2$ homology $H_n(S^n)$. Because the map f commutes with the antipode map, it factors to $\tilde{f} : \mathbb{RP}^n \to \mathbb{RP}^n$. This means that we have 2 exact

sequences:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_k(\mathbb{RP}^n) & \xrightarrow{\tau} & C_k(S^n) & \xrightarrow{\pi_*} & C_k(\mathbb{RP}^n) & \longrightarrow & 0 \\ & & \downarrow{\tilde{f}} & & \downarrow{f} & & \downarrow{\tilde{f}} & & \\ 0 & \longrightarrow & C_k(\mathbb{RP}^n) & \xrightarrow{\tau} & C_k(S^n) & \xrightarrow{\pi_*} & C_k(\mathbb{RP}^n) & \longrightarrow & 0 \end{array}$$

This gives us a pair of exact sequences on homology.

$$\begin{array}{ccccccccccc} H_n(\mathbb{RP}^n) & \longrightarrow & H_n(S^n) & \longrightarrow & H_n(\mathbb{RP}^n) & \longrightarrow & H_{n-1}(\mathbb{RP}^n) & \longrightarrow & H_{n-1}(S^n) & \longrightarrow & \cdots \\ \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \\ H_n(\mathbb{RP}^n) & \longrightarrow & H_n(S^n) & \longrightarrow & H_n(\mathbb{RP}^n) & \longrightarrow & H_{n-1}(\mathbb{RP}^n) & \longrightarrow & H_{n-1}(S^n) & \longrightarrow & \cdots \end{array}$$

$$\begin{array}{ccccccccccccc} \cdots H_k(S^n) & \longrightarrow & H_k(\mathbb{RP}^n) & \longrightarrow & H_{k-1}(\mathbb{RP}^n) & \longrightarrow & H_{k-1}(S^n) & \longrightarrow & H_{k-1}(\mathbb{RP}^n) & \longrightarrow & H_{k-2}(\mathbb{RP}^n) & \longrightarrow & H_{k-2}(S^n) \\ \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} \\ \cdots H_k(S^n) & \longrightarrow & H_k(\mathbb{RP}^n) & \longrightarrow & H_{k-1}(\mathbb{RP}^n) & \longrightarrow & H_{k-1}(S^n) & \longrightarrow & H_{k-1}(\mathbb{RP}^n) & \longrightarrow & H_{k-2}(\mathbb{RP}^n) & \longrightarrow & H_{k-2}(S^n) \end{array}$$

$$\begin{array}{ccccccccccc} \cdots H_1(S^n) & \longrightarrow & H_1(\mathbb{RP}^n) & \longrightarrow & H_0(\mathbb{RP}^n) & \longrightarrow & H_0(S^n) & \longrightarrow & H_0(\mathbb{RP}^n) & \longrightarrow & 0 \\ \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \downarrow{f_*} & & \\ \cdots H_1(S^n) & \longrightarrow & H_1(\mathbb{RP}^n) & \longrightarrow & H_0(\mathbb{RP}^n) & \longrightarrow & H_0(S^n) & \longrightarrow & H_0(\mathbb{RP}^n) & \longrightarrow & 0 \end{array}$$

We would like to show that all of the downward arrows are isomorphisms.

Ok, so on the final row we know that the downward arrow is an isomorphism, because a point is sent to a point. As we know the homology of a sphere is 0, we have that most of these diagrams break into 4 row blocks:

$$\begin{array}{ccccccc} 0 & \to & H_k(\mathbb{RP}^n) & \to & H_{k-1}(\mathbb{RP}^n) & \to & 0 \\ & & \downarrow{f_*} & & \downarrow{f_*} & & \\ 0 & \to & H_k(\mathbb{RP}^n) & \to & H_{k-1}(\mathbb{RP}^n) & \to & 0 \end{array}$$

So knowing that the $f_* : H_{k-1}(\mathbb{RP}^n) \to H_{k-1}(\mathbb{RP}^n)$ is an isomorphism tells us $f_* : H_k(\mathbb{RP}^n) \to H_k(\mathbb{RP}^n)$ is an isomorphism. This tells us by the 5-lemma that the map $f : H_n(S^n) \to H_n(S^n)$ is an isomorphism. ■

Corollary 2.12.2 For every map $f : S^n \to \mathbb{R}^n$, there exists x such that $f(x) = f(-x)$

Proof. We prove by contradiction. Consider the map $\frac{f(x)-f(-x)}{|f(x)-f(-x)|}$ which goes from the sphere to the equator. This is an antipodal map, and therefore must be odd. However, any map from the sphere to it's equator is trivial on top homology, and therefore even. This is a contradiction! ■

Suppose that we have sets[1] $U_1, \dots U_n \subset \mathbb{R}^n$. Let H be a hyperplane in $\mathbb{R}^n$. Then H divides $\mathbb{R}^n$ into two half spaces, and cuts every set U_i into $U_i^{\pm}$. Suppose additionally we have a volume function μ that assigns to subsets of $\mathbb{R}^n$ a volume.

Theorem 2.12.3 (The Ham Sandwich Theorem) Let $U_1, \dots U_n \subset \mathbb{R}^n$. Then there exists a hyperplane H such that $\mu(U_i^{H+}) = \mu(U_i^{H-})$.

[1] measurable, but not really important

Proof. We define a function $S^{n-1} \to \mathbb{R}^{n-1}$ as follows. For every point $x \in S^{n-1}$, there is a vector $v \in R^n$ which points in the direction of x. There is also a 1-dimensional family of hyperplanes which lies perpendicular to the vector v.

Claim 2.12.4 Let v be a direction. Then there exists a hyperplane H_v perpendicular to v which bisects the first body, that is $U_1^{H_v+} = U_1^{H_v-}$.

This is because these hyperplanes are parameterized by translation along v, and the intermediate value theorem guarantees such a value.
Given this claim, we create a function S^{n-1} to $\mathbb{R}^{n-1}$ as follows. Because for every $v \in S^{n-1}$, create the map

$$\begin{aligned} f : S^{n-1} &\to \mathbb{R}^{n-1} \\ v &\mapsto (\mu(U_2^{H_v+}), \mu(U_3^{H_v+}), \dots, \mu(U_n^{H_v+})) \end{aligned}$$

By the Borsuk Ulam theorem, there must be a point where $f(v) = f(-v)$. This tells us

$$\mu(U_i^{H_v+}) = \mu(U_i^{H_{-v}+}) = \mu(U_i^{H_v-})$$

for every i, which means that every single object is bisect. Boom. ∎

2.12.1 Necklaces

A necklace with n beads on k different colors is a string $a_1, a_2, a_3, \dots a_n$ each $a_i \in \{1, \dots, n\}$ is one of the k different colors. The necklace splitting the problem is the following:

Question 2.12.5 Suppose that we have a necklace with an even number of beads of each color. Suppose we would like to split this necklace between the two thieves. The thieves want an equal number of beads of each color. How many cuts do we need to add to the necklace?

Theorem 2.12.6 No more than k cuts are required to cut the necklace.

Lemma 2.12.7 The **moment curve** in $\mathbb{R}^k$ is the curve

$$\begin{aligned} \gamma : \mathbb{R} &\to \mathbb{R}^k \\ x &\mapsto (x, x^2, x^3, \dots, x^k) \end{aligned}$$

Every hyperplane H intersects γ in at most k places.

Proof. We "build" the necklace in $\mathbb{R}^k$ as follows. For each color, we have a color set

$$U_k = \{\gamma(x) \mid x \in [i, i+1) \text{ with } i \in a_k\}$$

By Borsuk-Ulam, there exists a hyperplane which cuts this necklace in half. The cut only intersects the necklace γ at k places. If you cut a segment not at an endpoint, there is another endpoint of the same color which hasn't been cut at the endpoint. Change the cut so it only intersects γ at integer values. This is a valid cutting of the necklace. ∎

2.12.2 Graphs

Definition 2.12.8 The **Kneser graph** $KG_{n,k}$ is the graph whose vertices are k-element subsets of $\{a_1, a_2, \dots, a_n\}$ where we put an edge between 2 vertices if their corresponding sets are disjoint.

Definition 2.12.9

A **coloring** of a graph G is an assignment of colors $1, 2, \dots m$ to the vertices of a G such that no edge has monochromatic endpoints. The smallest such m such that $1, 2, \dots, m$ colors G is called the **chromatic number** of G.

Theorem 2.12.10

The chromatic number of $KG_{n,2n+k}$ is $k+2$.

To prove that $k+2$ colors suffices, order the elements of the $2n+k$ element sets. For each i in this set, let K_i be the collection of all n element subsets whose least element is i.
Then $K_1, K_2, \dots, K_{k+1}$ are all disjoint. Finally, throw in the lucky set $K_{k+2} \cup \dots \cup K_{n+k+1}$. If two sets are disjoint, they lie in different K_i. This gives an upper bound for the chromatic number of $KG_{n,2n+k}$.
To prove the lower bound we need the following lemma:

Lemma 2.12.11

If S^n is covered by $n+1$ sets, each of which is either closed or open, then one of the sets contains a pair of antipodes.

Proof. Take U_i to be the open sets covering S^n. Define $S^n \to \mathbb{R}^n$ by

$$f_i(x) = \inf_{y \in F_i} d(x, -x).$$

By Borsuk-Ulam, we have that $f(x) = f(-x)$ for some $x \in S^n$. If $f_i(x) = 0$ for some x and all i, then we have that $x, -x \in U_i$. Otherwise, we have that none of the coordinates are zero, which means that $x, -x$ are not contained in any of the sets U_i for $i \leq n$. This means that $x, -x$ are contained in U_{n+1}. ■

Proof of Kneser Conjecture. We prove the following equivalent statement to the Kneser Conjecture: If the n-element subsets of a $(2n+k)$ element set are partitioned into $k+1$ classes, then one of the classes must contain a pair of disjoint subsets. Put the $2n+k$ points on S^{k+1} randomly. You can do this in such a way so that no $k+2$ points lie on a great k sphere. Suppose for contradiction there were coloring classes $A_1, \dots A_{k+1}$ of the subsets of the $2n+k$ points. For each of these $i \in 1, \dots, k+1$, let U_i be all the points $x \in S^{k+1}$ such that the open hemisphere centered at x contains an n-subset of A_i.
Take the set V to be all the points where are not contained in one of the U_i. Now it is clear that $V, U_1, \dots U_{k+1}$ are $k+2$ sets that cover S^{k+1}. By the lemma, there exists x such that $x, -x$ are contained in one of these covering sets. We have 2 cases to try and build a contradiction off of.

- If F contains x, then $H(x)$ and $H(-x)$ contains fewer than n points from our $2n+k$ randomly placed points. This means that $k+2$ points are on the boundary of $H(x)$, which means that $k+2$ points lie on a great circle. Bad!
- If $x, -x$ are contained in a U_i. We know that A_i contains a set of points in $H(x)$, and that A_i contains a set of points in $H(-x)$. But the points in $H(x)$ and $H(-x)$ are disjoint! This is a contradiction to the structure of A_i.

■

2.13 March 20

This lecture is a survey of the topics that we'll be covering in the last month after the break:

Definition 2.13.1

A *knot* is an injective piecewise linear function $i : \mathbb{R}/\mathbb{Z} \to \mathbb{R}^3$.

We will be interested in studying knots up to equivalence of continuous deformations. Because we are not usually able to study these functions abstractly, we realize them with diagrams representing their planar projections. Such a diagram is a curve in the plane, where every self-intersection is designated as an over crossing or an under crossing. The main question we will ask is: how do we determine if two planar representations give the same knot? To do so, we will develop a number of invariants that classify knots.

1. The first invariants that we will study are knot polynomials. Knot polynomials are something like the graph polynomials that we defined at the very start of the course, such as the chromatic polynomial and the connectivity polynomial. In fact, we should think of these polynomials as prescribing some kind of inclusion/ exclusion relation in knots. If you are interested in knot polynomials, there are 2 possible routes that you could take for a large project:
   - Knot Polynomials have some deeper topological or geometric definition. Look into this! In particular, the Alexander polynomial is define by taking some covering space of $\mathbb{R}^3 \setminus K$. In this way the Alexander polynomial is related to the fundamental group. More to come on that later! But you could look at the Alexander Polynomial as a project idea.
   - There is a way to think of many polynomials (including the ones we defined for knots) as coming from a kind of filtered homology theory. In other words, there is a way to construct a chain complex $C^\bullet(K)$ which has Euler characteristic equal to these knot polynomials. This is super weird! Furthermore, the homology of this chain complex is an invariant of the knot. The first example of a knot homology invariant is Khovanov homology, and you could look at this as a project.
2. Another invariant that we will study are braids. Braids are a bit like knots, but they are better behaved algebraically. A braid is a bunch of strings that go from one plane to another. They form a group under composition. We will study this group, and understanding the relations in this group will give us an algebraic way to categorize knots. Unfortunately, figuring out if two braids represent the same knot is just as difficult as figuring out if two knots are the same. If you are interested in the braid group, there is a nice representation of the braid group as a certain set of automorphism of the free group– one can think about the braid group as the permutation group which "doesn't forget" previous permutations made. This is the Artin Presentation.
3. We will also look at the fundamental group of the knot complement. We will give one way to compute this group, but it'll be pretty rough at first. If you were interested in learning about the fundamental group of knots, you could do a project on the interplay between the braid group and the knot group.
4. Finally, we will look at the theory of Seifert surfaces, which are bounding surfaces for knots. It's pretty easy to show that such a surface exists, but the topological properties of such surfaces are interesting. In particular, the minimal genus of these surfaces is an invariant that we would like to study.
5. Time permitting, we'll look at things like slice genus, or knot cobordisms.

Chapter 3

Knots

All of us have some idea of what a "knot " should be: something involving tangled string. The key feature of a knot is that it cannot be untangled- if knots could be easily untangled then they would be rather useless in their practical applications. In math-world, we have a slightly similar definition of a knot. We can't talk about open pieces of sting- strings with two ends- because all knots formed with a string with ends can be untangled with sufficient effort. However, we can talk about knots formed on loops of strings: these knots may not be untangled. I will propose some definitions of a knots, and why we do not use these definitions.

- A knot is the image of a circle in space. We can represent this as a function, $f : S^1 \to \mathbb{R}^3$. However, this does not work, as the function $0 : S^1 \mapsto 0$ should not represent a knot.
- A knot is the image of a circle under an injective function, $F : S^1 \hookrightarrow \mathbb{R}^3$.

However, this doesn't work as one can come up with an injective function whose knot has nasty image. For instance, we can consider a knot with an infinite number of crossings. This type of knot has all kinds of interesting analytic properties, but unfortunately does not describe a practical knots. We could also have embeddings that have cuts in them. We solve this problem by asking our knots to be continuous. We will restrain ourselves to knots that do have these desired properties. With that in mind, we use this definition of what a knot is:

Definition 3.0.2 A *knot* is an injective, continuous function $K : S^1 \hookrightarrow \mathbb{R}^3$ whose image is a finite collection of line segments. We will frequently refer to the image set as the knot.

Definition 3.0.3 A *link* is an injective, continuous function $L :\hookrightarrow \mathbb{R}^3$ whose image is a finite collection of line segments. We will frequently refer to the image set as the link.

What does this mean? It means that we formally think of a knot (or link) as a collection of line segments. However, we will frequently represent knots with curved lines, with the full understanding that if we want to, we could approximate these curves accurately with a finite number of line segments.

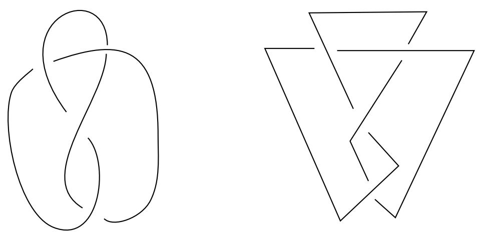

Figure 3.1: An example of a Knot and it's rectification

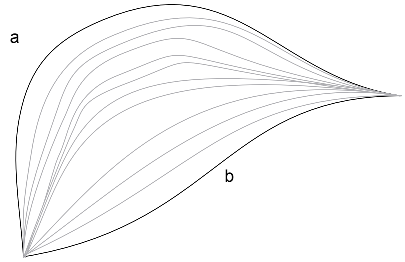

Figure 3.2: A homotopy between curves a and b

We want to talk about deformations of knots as well, and make sure that these deformations reflect our understanding of how real knots work. To do this, we need to define some topological ideas.

Definition 3.0.4 Let $f, g : X \to Y$ be two continuous functions. A *homotopy between f and g* is a continuous function $H : X \times [0,1] \to Y$ which satisfies the following properties:

$$H(x,0) = f(x)$$
$$H(x,1) = g(x)$$

A homotopy is a continuous deformation of one function into another. Visually, we think of a homotopy as a continuous deformation of the images of f and g into each other. This is wonderful, except for a few small details. We are not sure that this transformation preserves the wonderful properties of knots that we want. Homotopy is too weak a condition to preserve our knot structure. We present two improvements here.

Definition 3.0.5 Let $f, gX \hookrightarrow Y$ be two continuous functions. A *isotopy* between f and g is a continuous injection $H : X \times [0,1] \hookrightarrow Y$ which satisfies the following properties:

$$H(x,0) = f(x)$$
$$H(x,1) = g(x)$$

Even this isotopy is not good enough for us. If we consider a overhand knot being pulled tight, this constitutes an isotopy between our knot and the unknot. We need to do a little better. We want our knot transformations to mean “no piece of string gets too close to another piece of string." We start making this definition formal by describing triangle moves.

Definition 3.0.6 A Δ-move on a knot K (pronounced either delta-move or triangle move) is a replacement of a line segment ac in the knot with two line segments ab and bc in such a way that the triangle bounded by abc does not intersect the knot K. The new set of edges given by ab and bc constitutes a new knot.

If two knots differ by a sequence of Δ-moves, we say that they are equivalent via an ambient isotopy. We may also so say the knots are equivalent, the same knot, deformable, etc. We will have lots of names for these kinds of things, their obvious relationship is there. The name equivalence rightfully suggests that ambient isotopy places an equivalence relation on the set of knots.

Since we can’t actually visualize knots well in 3-dimensions, we frequently look at the their two dimensional projections. In order to make these pictures understandable, we require that the projection only have double intersections.

Definition 3.0.7 Let $K : S^1 \to \mathbb{R}^3$ be a knot. A projection $\pi : \mathbb{R}^3 \to \mathbb{R}^2$ is a *knot projection of K* if there are no points $x, y, z \in S^1$ such that $\pi(K(x)) = \pi(K(y)) = \pi(K(z))$. When we draw the picture of the knot, we can assign over and under crossings based on the projection. Such a picture is called a *knot diagram*.

With some analysis, one can show that such a projection always exists if K is a knot. It is interesting to note that given a knot diagram of K, one can always construct a knot K_1 which is ambient isotopic to K. This means that somehow, knot diagrams are an accurate way to represent different knots. We can now formulate the basic question of knot theory. How can we tell if two knot diagrams represent equivalent knots. There are a few ways to figure this out. We have some moves to transform diagrams into each other. We call theses the Reidemeister moves.

Theorem 3.0.8 All Δ transformations can be described by either Δ moves in the plane which do not effect crossings, or one of these three *Reidemeister moves*.

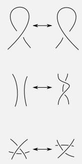

Proof. The idea is to break up your triangle move into lots of smaller triangle moves. This turns into a lot of casework, but let's get the idea. Suppose that we take E and replace it with A, B two edges via a triangle move. Look at the projection of this triangle move into the plane. This is a triangle– so subdivide this into smaller triangles so that at every step, the edges E, A, B only encounter at most 1 edge each.
If E encounters an edge, and A B do not, then this is the first Reidemeister move.
If E does not encounter an edge, and one of A and B do, then this is the first Reidemeister move.
If E does not encounter an edge, and both A and B encounter the same edge, then this is the second Reidemeister move.
If E does encounter an edge, and one of A and B do, this is just an ambient move.
If E does encounter an edge, and A and B both encounter different edges, then this is the third Reidemeister move. ■

This will make our lives much easier, as we will only have to deal with knot diagrams from here on out. When we refer to knots, we will call them K. Likewise, as knots are determined by their knot diagrams, we will call the knot diagram K as well. Here is an example of two knot diagrams that are equivalent to one another via a sequence of R-moves

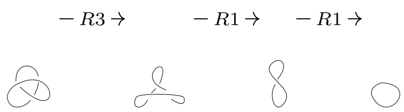

We unfortunately don't have such a good way to tell knot diagrams apart. However, we have a few good tools to get us started.

Definition 3.0.9 A *knot invariant* is a set X and a function $f : K \to X$, where the value of the $f(K)$ is independent of the diagram of K.

The idea is that we can come up with a quality (or invariant) that is associated to each knot, and if two knots return different values under the same invariant, then they cannot be related to each other via ambient isotopy. Let's see an example:

3.0.1 3-colorability and Fox p-colorings

We call the individual segments in a knot diagram *arcs* and any place where arcs meet to be crossings. We note that each crossing is the meeting of 3 arcs- one arc that goes over, and two arcs that meet "under". For the 3-coloring invariant, we choose 3 colors, say Red, Green. We label the arcs of our knot with the colors with the following rules: At each crossing the 3 arcs that meet must all be of the same color, or of 3 different colors.

A knot is called 3-colorable if the above conditions are met, and at least 2 colors are used.

How do we show that 3-colorability is an invariant of knots? We can check that if a knot is 3-colorable, then any Reidemeister move of that knot is still 3 colorable. One we check this, it is evident that 3 colorability is an invariant property of knots. We can use this property to show that the unknot is not equivalent to the trefoil.

Theorem 3.0.10 The Unknot is not deformable to the trefoil.

Proof. We first have to show that the 3-coloring invariant, is in fact an invariant. Suppose a knot is three colorable. Then we would like to show that when we apply a Reidemeister move to part of the knot, the knot remains 3 colorable. Fortunately for us, Reidemeister moves alter a knot locally and only effect one small section of the knot. ∎

Unfortunately, 3 colorability is not a very powerful invariant. 3-colorability maps knots into a two element set- either "3-colorable" or "not 3-colorable". This is pretty terrible, as there are surely more than 2 different knots. In other words, this knot invariant can tell us only if knots are different, not if they are the same.

We can generalize this idea of the *Fox p-coloring*. Let p be a prime number, and consider the field $\mathbb{Z}/p\mathbb{Z}$. We label the arcs of our knot diagram $0, 1, 2, 3, \dots p-1$ according to the following rule: At each crossing, let x be the value of the arc going over, and y and z the arcs that went under each crossing. We require

$$2x - y - z \equiv 0 \mod p$$

If there exists a coloring that uses at least 2 different numbers and satisfies the above condition, the knot is called n-colorable. This is equivalent to a linear algebra condition when you work over a finite field.

3.1 Jones Polynomial

In this lecture, we will construct a invariant returns a polynomial for each link. We start by talking about smoothings of crossings.

Definition 3.1.1 Let c_i be a crossing in a diagram L. We say that we have given c_i the *0-smoothing* or *1-smoothing* if we replace c_i with one of the smoothings of the crossing below.

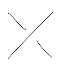

Figure 3.3: Crossing

Figure 3.4: 0-smoothing

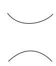

Figure 3.5: 1-smoothing

Definition 3.1.2 Let L be a link diagram. let $\{c_1, c_2, \dots c_n\}$ be the set of crossings for L. Give each crossing c_i the e_i smoothing, where $e_1 \in \{0, 1\}$. Then we call this diagram a *smoothing of L* and write $L_{e_1 e_2 \dots e_n}$.

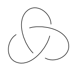

Figure 3.6: K

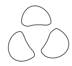

Figure 3.7: K_{000}

An example of a trefoil and it's 000 smoothing is given in Figure **??**.

The idea behind the Jones Polynomial is simple. We will take a large link, and decompose it into smaller links. By keeping track of our decomposition, we will generate an invariant of a link. Let's start by assigning a polynomial to a simple link. The polynomial for the unlink is simply

$$J(U) = q^{-1} + q^1$$

Now the goal is to create a way to break down a link into a bunch of unlinks while simultaneously keeping track of these decompositions. We have already "computed" the Jones Polynomial for every link with 0 crossings. Suppose that we have computed the Jones polynomial for every link diagram with $n-1$ crossings. Let K be a link with n crossings and pick some crossing c_n in the link. Let K_0 and K_1, be the two smoothings of c_1. Then K_0 and K_1 are both links with $n-1$ crossings, and therefore have associated to them polynomials $J(K_0)$ and $J(K_1)$. We then define the Jones polynomial of

$$J(K) = J(K_0) - qJ(K_1)$$

We need one more rule to compute the Jones Polynomial: if K_1 and K_2 are links, and $K_3 = K_1 \sqcup K_2$ is the disjoint link made of K_1 and K_2, then

$$J(K_3) = J(K_1)J(K_2)$$

We therefore have inductively defined a polynomial for every link. Now we want to show that this polynomial is an invariant of the link-i.e., if two link diagrams vary only by a Reidemeister move, then they produce the same polynomial (up to a change of sign or shift in degree). Let's check it for the first Reidemeister move.

Theorem 3.1.3 The Jones Polynomial is an invariant of the link, up to a shift in q degree and sign.

Proof. The idea is to show that the three Reidemeister moves only change the Jones polynomial of a link by a degree of q or by a sign change. We will show the invariance for each of the R-moves.

1. We show invariance under the first R-move.

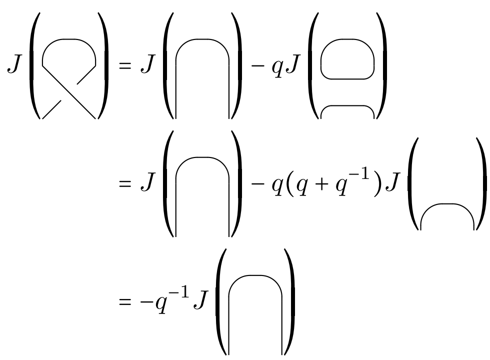

2. We show invariance under the second R-move.

With the relationships 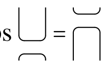 and

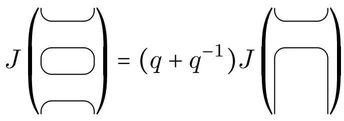

we perform the computation of the second Reidemeister move:

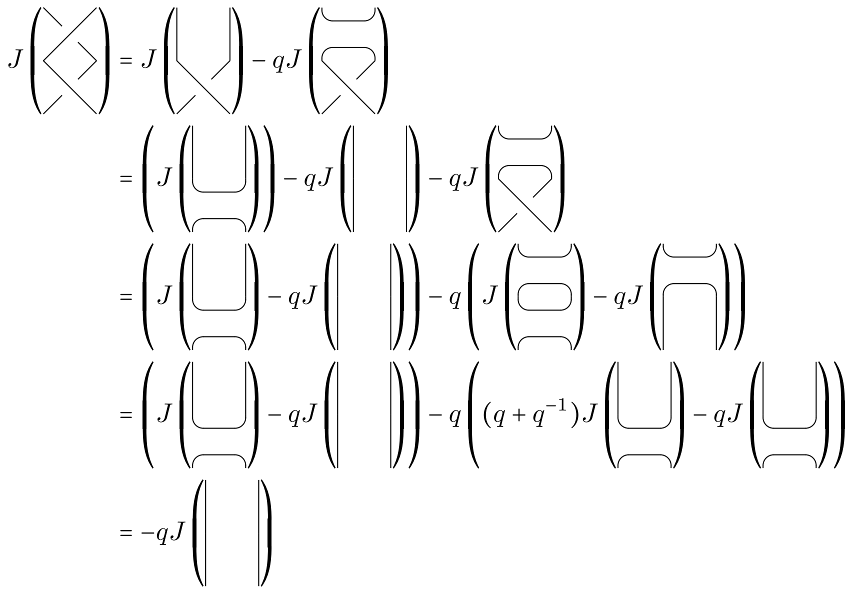

3. We show invariance under the third R-move. This one follows from the second Reidemeister move.

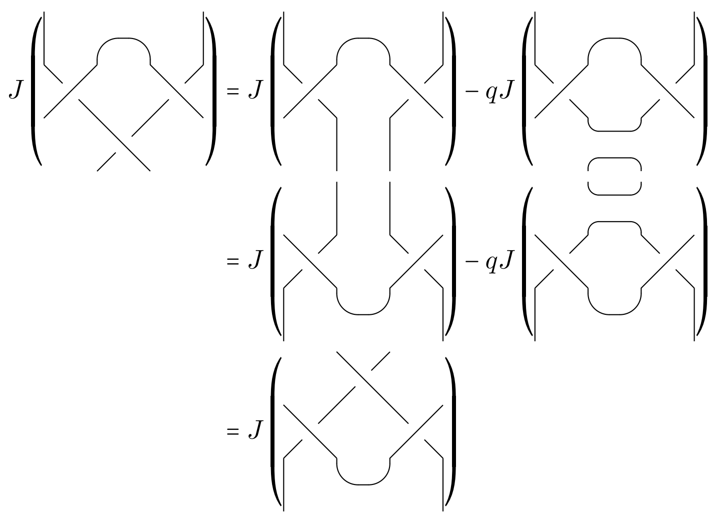

■

We now have shown that the Jone's polynomial is an invariant of the knot up to a sign and a multiple of q. There is a way to correct this shifting. If we assign an orientation to the knot, we can now label crossings in two different ways, as in the below figure

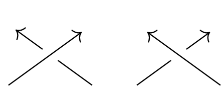

Let n_+ and n_- count the number of + and − crossings respectively. We see that if we multiply our J polynomial by a factor of $(-1)^{n_-}q^{n_+-2n_-}$ that this new polynomial now a complete invariant of the knot.
If we try and compute the Jone's polynomial of a knot by continually taking 0 and 1 resolutions, we see that we reach every possible smoothing of the knot. This means that if our original knot diagram has n crossings, there are 2^n different resolutions of the knot. We see that for each 1 resolution the knot contains, it picks up factor of $-q$. This gives us a simpler way to compute the jones polynomial. The Jones polynomial is given by

$$\sum_{i\in\{0,1\}^n}(-q)^{|i|}J(K_i)$$

where K_i means the resolution where the crossings are given by the string i, and $|i|$ counts the number of times 1 appears in the string i. Let's do a sample computation. We start by creating a cube of resolutions.

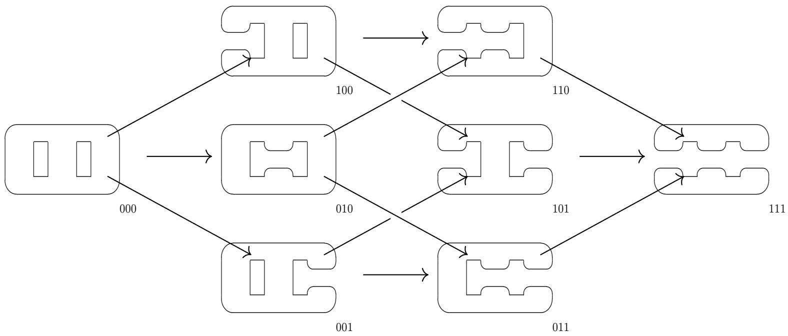

Each vertex on the cube is a smoothing of the knot, and the smoothings are ordered from left to right in columns corresponding to the number of 0 smoothings chosen. We draw an arrow between two vertices if they differ by just a single place in their smoothing. We now create a cube that lists each resolution's Jone's polynomial times $(-q)^i$, where i is the column the polynomial lies in.

$$\begin{tikzcd}
 & q(q+q^{-1})^2 \arrow[r] \arrow[dr] & q^2(q+q^{-1})^1 \arrow[dr] & \\
(q+q^{-1})^3 \arrow[ur] \arrow[r] \arrow[dr] & q(q+q^{-1})^2 \arrow[ur] \arrow[dr] & q^2(q+q^{-1})^1 \arrow[r] & q^3(q+q^{-1})^2 \\
 & q(q+q^{-1})^2 \arrow[ur] \arrow[r] & q^2(q+q^{-1})^1 \arrow[ur] &
\end{tikzcd}$$

We now can compute the unnormalized Jones polynomial: it is just the alternating sum of the columns of this cube. That is, the unnormalized Jones polynomial of the trefoil is $(q+q^{-1})^3 - 3q(q+q^{-1})^2 + 3q^2(q+q^{-1})^3 - q^3(q+q^{-1})^2 = q^{-2} + 1 + q^2 - q^6$.
One last thing to notice about the Jones polynomials– it is multiplicative with respect to knot connect sum.

Definition 3.1.4 Given two knots oriented knots K_1 and K_2, define their *connect sum* to be the knot $K_1\#K_2$ where we delete a small segment of K_1 and K_2, and glue them by their boundaries in an orientation preserving way.

Claim 3.1.5 Up to a multiple of q and sign, $J(K_1\#K_2) = J(K_1)J(K_2)$.

Proof. The cube of resolutions that we've drawn out is a product of the cube of resolutions for K_1 and the cube of resolutions for K_2. ■

3.1.1 Seifert Surfaces

But do these polynomials mean anything? One way that we can see these polynomials take form geometrically if via Seifert Surfaces.

Definition 3.1.6 Let L be a link. A *Seifert Surface* for L is a connected, compact, oriented, embedded surface whose boundary is L.

Now, a given link L may have many different Seifert surfaces; one can see that given a Seifert surface, you can modify it by adding on a handle and obtain a new Seifert surface. This makes it seem like this may not be such a good tool for studying links.

Theorem 3.1.7 Let L be a link. Then there exists a Seifert surface for L.

Proof. This one is pretty to write out, but pretty easy to draw, so we are going to do that instead! The idea is to assign an orientation to the knot. Apply a resolution to each crossing so that the orientations of the resolution are compatible with the orientation of the knot. This will give you a bunch nested disks. Each crossing now corresponds to a twisted handle that connects the boundary of one disk to the next.
Here is another way to visualize this: given a knot diagram, and a resolution of the crossings in n orientation preserving way, we get a collection of circle. For each circle S_i, label the spots x_{ij} of the resolutions connected to this disk. Now to create a surface, let D_i be a collection of disks. Arrange the disks so that D_i has no variation in the z axis direction, so that the projection $\pi_{xy}(\partial D_i) = S_i$, and that the z value of D_i is higher than D_j whenever $\pi(\partial D_i) \subset \pi(D_j)$.
Then connect the points x_{ij} and x_{kl} with a twisted handle whenever they came from the resolution of the same crossing. This gives us a ∎

This will give you a surface which has first homology generated by each of its boundary components. Now, instead of working in $\mathbb{R}^3$, we are going to work in S^3. Let M be the Seifert surface.

Claim 3.1.8 There is a nondegenerate form $g : H_1(M) \times H_1(S \setminus M) \to \mathbb{Z}$. In particular, the first homology of the Seifert surface its complement are isomorphic.

Proof. In fact, we prove something a little different: given any *handlebody* in S^3, there is a pairing with its complement. In fact, a handlebody and its complement are homeomorphic.
Notice a handlebody is just a 3-ball with some disks which are identified to each other. Given a handlebody, pick disks which fill in the "donut-holes" of the handlebody, which lie in the complement. Cutting the complement along these holes gives us a 3 ball. The number of disks that we need to cut will be equal to the genus of the handlebody, which determines handlebodies up to homeomorphism.
Let α be a generator of homology of the handlebody, and β an element in $H_1(S \setminus M)$. Then the generator α corresponds to some disk, and the (signed) intersection of β and the disk associated to α gives us the symmetric form.

∎

There is an easier way to think of his form. We think of elements of $H_1(M)$ and $H_1(S \setminus M)$ as begin represented by oriented loops. Then the pairing is the *linking number* between the two loops, which is the number of positive crossings, less the number negative crossings in a projection of the two loops.

Claim 3.1.9 The orientation of the Seifert surface gives us two different maps $H^1(M) \to H^1(S^3 \setminus M)$.

Proof. As the Seifert surface is orientable, there is a normal direction at every point of the Seifert surface. One can take any loop representing a class of homology, and push it off the Seifert surface into the complement. This gives us a loop in the complement, which represents an element of homology in $H^1(S^3 \setminus M)$. Of course, if we use the other normal, we get a different map. ∎

Definition 3.1.10 Let K be a knot. Let M be a Seifert surface for the knot, and $i_+ : M \to S/M$ be an inclusion. Then there is a map, called the *Seifert Matrix*

$$\begin{aligned} A : H_1(M) &\to H_1(M) \\ \alpha &\mapsto g(\alpha, i_+(\alpha)) \end{aligned}$$

The determinant of the map $\Delta(K) := \det(tA - A^t)$ (where t is a dummy variable) is called the *Alexander Polynomial*.

Claim 3.1.11 The Alexander polynomial is a knot invariant up to sign and a multiple of t.

Proof. We show that the alexander polynomial is generated by a Skein relation as well. However, this is a Skein relation that involves reversing the crossings. At a given crossing (with orientation), denote the + and − and 0 resolutions of a particular crossing by

We like to prove that the Alexander polynomial as defined here satisfy the Skein relation

$$\Delta(K^+) - \Delta(K^-) = (1-t)\Delta(K^0)$$

In order to prove this, we pick a very special basis for the homology, given by a certain realization of the Seifert Surface.
Draw the Seifert surface as before, with a collection of disks and bridges between them. On every disk. there are bridges that point "outward" to lower disks, and bridges that point inward to "higher"disks. On the disk D_i, let T_{ik} be a labeling of the outward pointing bridges in a cyclic order.
Then we can give a basis for $H_1(\Sigma)$ in the following way: Let v_{ik} be the loop that starts on the disk D_i, goes down the T_{ik} crossing, goes over to the $T_{i(k+1)}$ crossing, then goes up back to the disk D_i.Note here that we *do not* take these indexes to be cyclic– there is no entry v_{in}, where $1 \leq k \leq n$. This is because we can write v_{in} as $\sum_k v_{ik}$.
Now, let's put an ordering on the basis so that v_{ij} comes before v_{mn} if $i \leq m$, and then if $j \leq n$. In this ordered basis, what does the matrix A look like? The entry a_{ij} is nonzero if and only if the loops link– but notice that v_{ij} and v_{mn} link if and only if $i = m$ and $j = n \pm 1$. From this, we know that A splits as a bunch of matrices along the diagonal, one block for every disk D_i, and has entries only on the diagonal, super diagonal, and sub diagonal.
For ease of notation, let $B = A + tA^T$, and let $A^{\pm}$, and A^0 denote the matrices to the $+, -$ and 0 resolution of a knot on the T_{11} crossing. Now, we would like to compute $\det(B^+) - (\det B^-)$. By minor expansion, we have that

$$\det(B^+) - \det(B^-) = b^+_{11}[B^+_{11}] - b^+_{12}[B^+_{12}] - (b^-_{11}[B^-_{11}] - b^-_{11}[B^-_{11}])$$

Now, look at $b^{\pm}_{11}$. One of these entries will be 0, and the other one will necessarily be $1+t$. Why is this? Becuase the crossing T^+_{11} and T^+_{12} either both contribute ± 1 to the linking of v_{11}, or they contribute opposite signs. Similarly, T^-_{11} and T^-11 contribute opposite signs to the crossing, or a $\pm$ 1. We work on the case where $b^-_{11} = 0$, the other case is similar argument.

$$= b^+_{11}[B^+_{11}] - b^+_{12}[B^+_{12}] + b^-_{11}[B^-_{11}]$$

Now, b^-_{12} and b^+_{12} are the same, as v_{11} and v_{12} interact only at the the T_{12} crossing, which is unchanged. Similarly , the matrices $B^{\pm}_{11}$ are the same.

$$\begin{aligned} &= b^+_{11}[B^+_{11}] - b^+_{12}([B^+_{12}] - [B^-_{12}]) \\ &= b^+_{11}[B^+_{11}] \\ &= (1-t)[B^+_{11}] = (1-t)[B^0] \end{aligned}$$

■

3.2 Braids

Braid groups describe one particular way that we can tangle stings together. Intuitively, a braid is a tangle of strings that go from top to bottom, that is they proceed in only one direction. The stings of the braid can't intersect, and they can't loop back on themselves. A braid is like a permutation with some additional information encoded about the order that you took the permutations in, and we will formally see later why this is an accurate description of the braid group.

Definition 3.2.1 A collection of functions $f_i : [0,1] \to \mathbb{R}^3$, $i \in 1, 2, \dots n$ is called a **braid on n strings** and is written as β if the following properties hold:

1. $f_i(x) \in \{x\} \times \mathbb{R}^2$
2. $f_i(x) \neq f_j(x)$ for all $i \neq j$
3. Im f_i is a finite collection of line segments
4. $f_i(0) = (i, 0, 0)$ and $f_i(1) = (\sigma(i), 0, 1)$ where $\sigma_i \in S_n$ a permutation function.

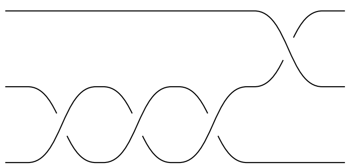

As one can see, this definition says that a braid is much like a knot, in that it can be represented by a finite collection of line segments, and that it does not intersect itself. How a braid is different than a knot (or a link) is that braids "travel" in one direction, that is, if we orient the braid by the orientation $[0,1]$, we see that the image of the braid is oriented against a single axis. We will choose this axis to be the x axis, and hence going along the braid means "going left to right." Like knots, we are interested in braids up to the equivalence of ambient isotopy; when we draw a braid, we really are just picking a representative of the equivalence class of braids on n strings that are isotopic to this one. We will denote the set of all equivalence classes of braids with n strings as $\mathcal{B}_n$. Here are some equivalent braids.

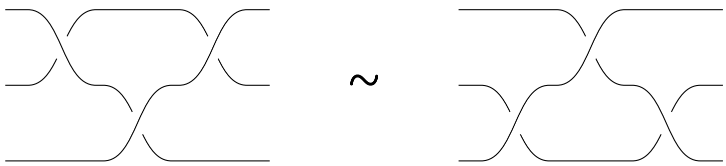

How can we make a group out of braids? The way that designed our braids, we can form a group operation by braid "stacking", that is placing one braid on top of each other.

Definition 3.2.2 Let $\beta_1, \beta_2 \in \mathcal{B}_n$, where $\beta_1 = \{f_i(t)\}$ and $\beta_2 = \{f_i'(t)\}$. Define the new braid $\beta_3 = \{g_i\}$ where

$$g_i = \begin{cases} f_i(2t)/2 & : t \in [0, 1/2) \\ f_\sigma(i)'(2(t-1/2)/2 + 1/2 & : t \in [1/2, 1] \end{cases}$$

In abuse of notation, when we write $f_i/2$ we mean a scaling of the x axis by $1/2$
Then we define a group structure on $\mathcal{B}_n$ by $\beta_1 \circ \beta_2 = \beta_3$ Then call $\mathcal{B}_n$ the **braid group on n strings**.

Of course, we need to check that the axioms for groups are satisfied. The composition of braids is clearly associative: scaling along the x axis is an isotopy of braids, and other than scaling across the x axis, the braids $(\beta_1\beta_2)\beta_3)$ and $\beta_1(\beta_2\beta_3)$ are the same.

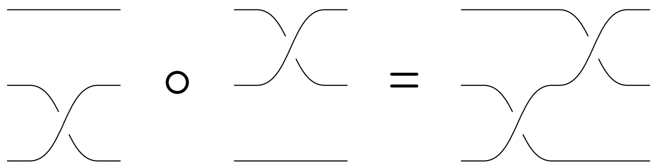

We also want to make sure that there exists an identity element and inverse elements. The identity element is easy to come up with: it is simply a set of horizontal strands that have no twisting.

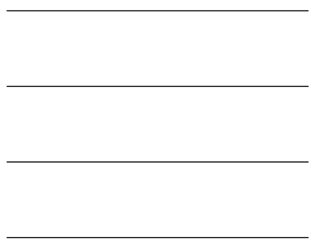

As for inverses, you can have them too. Simply take the braid β you want to inverse, and reflect it along the yz plane, to get a braid β^{-1} This should give a braid that, when composed with it's original, is equivalent to the identity. Therefore, we are justified in calling this object the braid group on n strings. Ideally, we would want a group presentation for this object, so we wouldn't have to draw pictures whenever we wanted to talk about a braid. It's pretty easy to get generators for the braid group: all braids are generated by the sequential crossing of strands that are adjacent to each other. If we want to designate any braid, we need only the **braid word** which corresponds to the sequence of strands being crossed. We write the letter σ_i to mean the crossing of the i and $i+1$ strands in a right over left fashion. Therefore, the braid given in Figure **??** is given by the braid word $\sigma_1\sigma_1\sigma_1\sigma_2^{-1}$, or sometimes written $\sigma_1^3\sigma_2^{-1}$. With these braid words, we see that the stacking of braids corresponds to the concatenation of braid words.

Exercise 3.2.3 Show that the mirror braid of β has braid word given by β^{-1}

Notice that a braid is not uniquely identified by its braid word: in fact, there are many different braid words that represent the same braid. For instance, the braid $\sigma_3\sigma_1$ is the same as the braid $\sigma_1\sigma_3$. However, the relations giving braids are based mostly on the relations for knots.

Theorem 3.2.4 The braid group has a finite presentation as

$$\langle \sigma_1, \dots, \sigma_{n-1} \mid \forall |i-j| \geq 2, \quad \sigma_i\sigma_j = \sigma_j\sigma_i, \sigma_i\sigma_{i+1}\sigma_i^{-1} = \sigma_{i+1}^{-1}\sigma_i\sigma_{i+1}\rangle.$$

Recall that a group presentation is the "free-est" group on specified generators that satisfies the given relations. One way of thinking about this is that a group presentation of G specifies a map $F_n \to G$, where F_n is a the free group. The normal closure of the specified generators gives us the kernel of this map. Notice that every finitely generated group has a presentation by the same generators.

Proof. The generators represent the elementary crossings of the i and $i+1$ strands. These are basically the Reidemeister moves written out as algebraic relations.
Notice, that braid isotopy is almost the same as knot isotopy, except that the first Reidemeister move is forbidden!
The first type of braid relation corresponds to the "ambient isotopy." It says that if you have two crossings, and those crossings are some distance from each other, then you can commute them past each other because they share no stings in common. This means that crossings "commute at distance".
The second Reidemeister move is already encoded into the braid group by default: this is just the inverse relation between generators.
The third Reidemeister move, if you write it down, is the relation that we have written above.
With these relations, we capture the R-moves, and therefore all the additional structure in the braid group. ■

Notice that if I add in the additional relation that $\sigma_i^2 = 1$, then I get the permutation group. Interesting!
Now, we would like to use braids to study knots. Notice that every braid β gives a knot by gluing its ends together in a circle.

Definition 3.2.5 Given a braid β, we say the *braid closure* of β, $\bar{\beta}$, is the link created by adding loops from the upper endpoints of the braid to the lower endpoints of the braid in a non-crossing fashion. If $\bar{\beta} = K$, we say that β is a *braid word* for K.

Given a knot K it is possible to have a lot of different braid words. For instance, both the braid 1 on one string and the braid σ_1 on 2 strings are braid words for the unknot.

Theorem 3.2.6 Suppose that β_1, β_2 are two different braid words for K. Then they are related by the following *Markov moves*:

- Concatenation: replace β with $\alpha^{-1}\beta\alpha$.
- Reidemeister 1: Replace β on n strings with $\beta\sigma_n$ on $n+1$ strings.

Proof. There are 2 types of isotopy which are not taken into account by the braid relations: the first one is the isotopy which does not change crossings in the diagram, but possibly moves the braid outside of the "braided region". This means that part of the braid moves up over the top of the braided region, and then shows up again on the bottom. You can check that this is the same as concatenating the braid with a word.
The other type of isotopy that we have not taken into account is the first Reidemeister move. Suppose we want to do a first Reidemeister move on a little segment. By applying many R_2 moves, we can assume that this Reidemeister move occurs on the nth string. By applying concatenation, we may assume that it occurs at the end of the braid word. Then the Reidemeister move looks like adding a little loop to the end of the braid. However, adding a little loop to the end of the braid is the same as adding an additional string with a crossing, once we blow up the loop in size a bit. This shows why R_1 can be expressed as the replacement of β with $\beta\sigma_n$ on $n+1$ strings. ∎

So now we know that if we have to braid words, then we can algebraically manipulate one into the other. This is really cool, because on the one hand we've reduced the problem of showing that two knots are the same to an algebraic manipulation. On the other hand, you can prove that it is hard to show that two braid words represent the same knot.
But, we do not know if a knot is necessarily representable by a braid.

Theorem 3.2.7 (Alexander) For every knot K there is a braid β with $\bar{\beta} = K$.

Proof. By our assumption about knots, there exists a linearization of the knot image. So, let's pick a linearization of the knot's image so that $K = \bigcup_i L_i$
Now, assign an orientation of the knot. Look at a planar projection (knot diagram), and pick a point in the planar projection. If it is the case that each line L_i has an orientation that agrees with the orientation around the point, then we are done! But this is probably not the case. So, for every line segment L_i which disagrees with the orientation around a point x, use a bunch of R_2 moves to replace it with 2 line segments L_i' and L_i'' so that the left endpoint of L_i' is the left endpoint of L_i, the right endpoint of L_i'' is the right endpoint of L_i, and L_i', L_i'' have a common endpoint which is on the opposite side of x. Now, L_i' and L_i'' have an orientation which agrees with x. Proceed with this process on every single line segment that has a disagreeable orientation. This provides a knot diagram whose orientation is agreeable with x. ∎

So, this means that we can take every knot, and find a knot diagram which represents it. Notice that this gives an extremely easy way to produce a Seifert surface, because all of the crossings get unoriented in the same direction. From this piece of data, we should expect that the braid word gives us a purely algebraic way of computing something like the Alexander polynomial for a knot! In fact, it does, but in order to stray away from matrix manipulations for the rest of the course, we will not be going over this one.

3.3 Fundamental Groups

Definition 3.3.1 Let X be a topological space, and $x_0, x_1 \in X$ be two points. A **path** in X with endpoints x_0 and x_1 is a continuous function $\gamma : [0,1] \to X$ with the property that $\gamma(0) = x_0$ and $\gamma(1) = x_1$. If $x_0 = x_1$, we say that γ is a **loop**, and we call x_0 the **base point** of those loops.

Definition 3.3.2 Let $f_0 : X \to Y$ and $f_1 : X \to Y$ be two different functions. A **homotopy** between f_0 and f_1 is a continuous function $H(x,t) : X \to Y$ with the property that

$$\begin{aligned} H(x,0) =& f_0(x) \\ H(x,1) =& f_1(x). \end{aligned}$$

If there is a homotopy between f_0 and f_1, we say that f_0 is **homotopic** to f_1 and write $f_0 \sim f_1$.

Definition 3.3.3 Let X be a topological space and γ_0 nd γ_1 be two paths with endpoints $x_0, x_1 \in X$. Then a **path homotopy from** γ_0 **to** γ_1 is a homotopy between γ_0 and γ_1 such that

$$\begin{aligned} H(0,t) =& x_0 \\ H(1,t) =& x_1. \end{aligned}$$

We will usually represent a homotopy with this type of square diagram. One axis represents t, the homotopy parameter, and the other one represents x, the loop parameter.

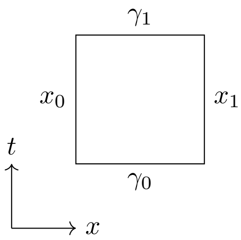

Claim 3.3.4 Path homotopy is an equivalence relation on paths from x_0 to x_1.

Proof. We need to show that path homotopy satisfies the three properties of an equivalence relation.

- (Reflexive) This one is the easiest. We can let $H(x,t) = \gamma(x)$ for all t, and this will give us a homotopy between γ and itself.
- (Symmetric) Suppose that $H(x,t)$ is a homotopy between γ_0 and γ_1. Then $H(x, 1-t)$ is a homotopy between γ_1 and γ_0.
- (Transitive) Suppose that $H_{01}(x,t)$ is a homotopy between γ_0 and γ_1, and $H_{12}(x,t)$ is a homotopy between γ_1 and γ_2. Then define a new homotopy H_{02} by

$$H_{02}(x,t) = \begin{cases} H_{01}(x,2t) & t \in [0,1/2] \\ H_{12}(x,2t-1) & t \in (1/2,1] \end{cases}.$$

  You should check that this gives us an honest continuous homotopy between γ_0 and γ_1.

■

Because of this claim, we can talk about functions up to the equivalence relation of homotopy. If we are interested in writing homotopy classes of a function, we will denote it with square brackets. Notationally, $\gamma_0 \sim \gamma_1$ implies $[\gamma_0] = [\gamma_1]$.

Definition 3.3.5 Let γ_0, γ_1 be two paths with base point x_0. Then we define the **composition loop**

$$\begin{aligned}\gamma_0 \cdot \gamma_1 : [0,1] &\to X \\ x &\mapsto \begin{cases} \gamma_0(2t) & t \in [0, 1/2] \\ \gamma_1(2t-1) & t \in (1/2, 1] \end{cases}\end{aligned}$$

which is again a loop that starts with base point x_0.

Unfortunately, loop composition as written is not an associative operation. That is,

$$(\gamma_0 \cdot \gamma_1) \cdot \gamma_2 \neq \gamma_0 \cdot (\gamma_1 \cdot \gamma_2)$$

Claim 3.3.6 Suppose that $\gamma_0 \sim \gamma_1$ are two path homotopic loops, and β is another loop. Then $\gamma_0 \circ \beta \sim \gamma_1 \circ \beta$. In particular, the composition rule is well defined over equivalence classes of loops up to homotopy.

Proof. We just need to exhibit a homotopy between $\gamma_0 \cdot \beta$ and $\gamma_1 \cdot \beta$. Here is a picture of that homotopy:

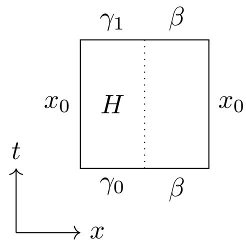

■

Claim 3.3.7 The composition rule is associative over equivalence classes of loops up to homotopy.

Proof. We need to show that $(\gamma_0 \cdot \gamma_1) \cdot \gamma_2) \sim \gamma_0 \cdot (\gamma_1 \cdot \gamma_2)$ via homotopy. The idea is that you can reparameterize the path in a homotopy; the only difference between these to paths is that

- The first goes through γ_0 and γ_1 at 4x speed, and then γ_2 through double speed.
- The second goes through γ_0 at double speed, and γ_1, γ_2 at 4x speed.

Define the following homotopy between $(\gamma_0 \cdot \gamma_1) \cdot \gamma_2)$ and $\gamma_0 \cdot (\gamma_1 \cdot \gamma_2)$.

$$H(x,t) = \begin{cases} \gamma_0(4/(1+t)x) & x \in [0, (t+1)/4) \\ \gamma_1(4(x-(1+t)/4)) & x \in [(t+1)/4, (t+2)/4) \\ \gamma_2(4/(2-t)(x-(t+2)/4)) & x \in [(t+2)/4, 1] \end{cases}$$

This homotopy is given by a really nasty function, so let's instead draw it as a picture so we have an actual idea of what it represents.

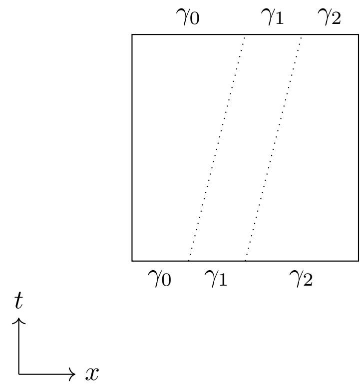

■

Claim 3.3.8 The constant loop 1_{x_0} has the property that

$$[1_{x_0}] \cdot [\gamma] = [\gamma] \cdot [1_{x_0}] = [\gamma]$$

for any loop γ based at x_0.

Proof. We use the fact that the constant loop can be reparameterized to any length that we choose. We define the homotopy as follows:

$$H(x,t) = \begin{cases} x_0 & x \in [0,(1-t)/2) \\ \gamma((2-t)(x-(1-t)/2)) & x \in [(1-t)/2, 1] \end{cases}$$

Again, easier to represent with a picture.

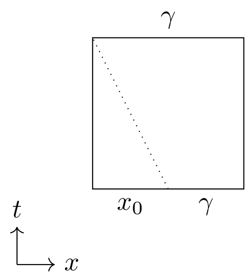

■

Claim 3.3.9 Suppose that $\gamma(t)$ is a loop. Define γ^{-1} to be the loop $\gamma(-t)$. Then

$$[\gamma] \cdot [\gamma^{-1}] = [1_{x_0}]$$

Proof. We need to exhibit a homotopy from the loops γ and γ^{-1} to the constant loop. A formula for this homotopy is

$$H(x,t) = \begin{cases} \gamma(2x) & x \in [0,(1-t)/2) \\ \gamma(t) & x \in [(1-t)/2,(1+t)/2) \\ \gamma(1-2x) & x \in [(1+t)/2, 1] \end{cases}$$

Which is, as usual , easier to visualize as a picture instead of a piecewise function.

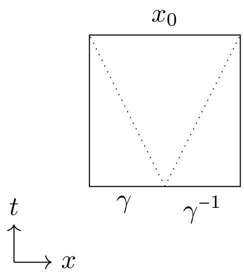

■

Definition 3.3.10 Let X be a topological space, and x_0 a point in X. The **fundamental group** of X is

- $\pi(X, x_0)$ as a set is all equivalence classes of loops at x_0.
- The group multiplication is loop composition.

The three claims that we made above prove that $\pi(X, x_0)$ is a group.

Definition 3.3.11 Let β be a path in X connecting x_0 to x_1. Define the **change of base point map** $F_\beta : \pi(X, x_1) \to \pi(X, x_0)$ as follows. Suppose that γ is a path that has endpoints at x_1. Then let

$$F_\beta(\gamma) := (\beta \cdot \gamma \cdot \beta^{-1})$$

. This map is only dependent on the homotopy class of h.

Claim 3.3.12 The change of base point map is a group homomorphism.

Proof. We need to prove that the map preserves the group composition law. Here is a picture why! ■

Claim 3.3.13 Suppose that X is path connected. Let $x_0, x_1 \in X$ be two points in X. Then $\pi(X, x_0)$ is isomorphic to $\pi(X, x_1)$.

Proof. Since X is connected, it is clear that there is path β with endpoints x_0 and x_1. Check that the change of base point map is in fact an isomorphism of groups by composing it with the reverse change of base point map. ■

Claim 3.3.14 Suppose that X and Y are two different space, and $f : X \to Y$ is a continuous function. Then there is an induced group homomorphism $f_* : \pi(X, x_0) \to \pi(Y, f(x_0))$.

Proof. Using a function $f : X \to Y$, we can take loops in X and get loops in Y as follows. Let $\gamma : [0,1] \to X$ be a loop. Define

$$f_*(\gamma) = f \circ \gamma : [0,1] \to Y.$$

We can check that if $\gamma_0 \sim \gamma_1$, then it is the case that $f_*(\gamma_0) \sim f_*(\gamma_1)$. This is because if $H(x,t)$ is a homotopy between γ_0 and γ_1, then $f(H(x,t))$ is a homotopy between $f_*(\gamma_0)$ and $g_*(\gamma_1)$.
Finally, we need to check that f_* respects the group operation, that is

$$f_*(\gamma_0 \circ \gamma_1) = f_*(\gamma_0) \circ f_*(\gamma_1)$$

■

Claim 3.3.15 Suppose that $f : X \to Y$ and $g : Y \to Z$ are two continuous maps. Then

$$g_* \circ f_* = (f \circ g)_*$$

Theorem 3.3.16 The fundamental group of the circle is $\mathbb{Z}$.

Proof. To come later in this class. But the idea is that the number of times that a loop winds around the origin tells you what number to map every loop to. ■

3.4 Combinatorial Definition of the Fundamental Group

Now that we have an intutition for this thing in the topological sense, let's mork out a combinatorial definition for it. As usual, we work with simplicial complexes.

Definition 3.4.1 Let Δ be a simplicial complex. Pick $x \in \Delta$ to be a basepoint. Let E be the set of 1 simpleices in Δ, and let F be the set of 2 simplieices. Define the group $\pi_1(\Delta, x)$ as strings of vertices $xv_1v_2\ldots v_kx$, where $v_iv_{i+1} \in E$ for every pair. The group operation is given by concatenation, and we mod out the paths by the following two equivalences:

$$v_iv_{i+1}v_i = v_i$$

$$v_iv_jv_k = v_i \text{ whenever } v_iv_jv_k \in F$$

This certainly gives us a group, and in the case when we take the topological realization of Δ, we have that this group si the same as $\pi_1(\Delta, x)$.
Notice that the higher cell structure (for $n \geq 3$) does not affect the fundamental group!

Example 3.4.2 Let C_n be the n cycle. Then $\pi_1(C_n, x) = \mathbb{Z}$. We can prove as follows. Pick a vertex 1 next to x. You now define a morphism $\pi_1(C_n, x) \to \mathbb{Z}$ by sending each string $xv_1 \cdots v_kx$ to the quantity $n_+ - n_- \in \mathbb{Z}$, where n_+ is the number of times $x1$ shows up in the string, and n_- is the number of times that $1x$ shows up. The quantity $n_+ - n_-$ is invariant under the equivalence of strings that we have defined above.