Topological Graph Theory

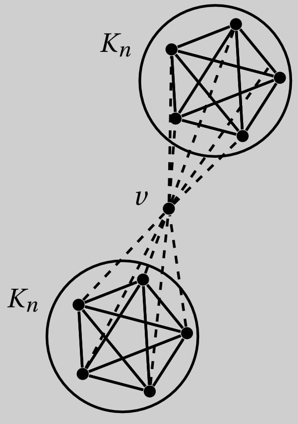

The Dumbbell Graph

A graph is a set of vertices which are connected by edges. This is a simple example of a topological space. Two topological characteristics of graphs are the number of connected components, and the number of cycles they contain.

Definitions

Theorems and Lemmas

Examples

Basic Definitions

A *graph* is a mathematical representation of network. We will explore some basic combinatorial properties of graphs, such as paths and cycles. Here, you can also find the definitions of some commonly used graphs like trees and complete graphs.

Graph theory is a branch of mathematics that naturally arises when you study a group of objects and a relationship that exists between pairs of objects. For example, you might be studying a group at a party, and the friendships that exist between pairs of attendees. Even this simple setting presents itself with some interesting questions, like how tight-knit this group is, or which folks at the party are the least related. Since this type of structure shows up so frequently in mathematics, we will set up an abstract mathematical object called a graph which contains the mathematical data to describe these types of relationships.

1 Definition *Graph*

A *graph* G is a collection of *vertices* $V(G)$ and a set of *edges*

$$E(G) \subset V \times V.$$

The edge is required to be unordered in the sense that whenever $(v, w) \in E(G)$, the pair $(w, v) \notin E(G)$.

In the case where the graph G is clear, we will suppress the label and simply denote the vertices by V and edges by E. While this definition gives us the mathematical rigor necessary to begin our explanation of graphs, it is useful to have some good examples in mind to ground our discussion.

2 Example *Friends in the Order*

Let us look at the motiviating discussion of friend groups at a party, and see how this fits into the definition of a graph. Each person in the population would represent a vertex, and the edges between vertices correspond to the pairs of attendees who are friends.

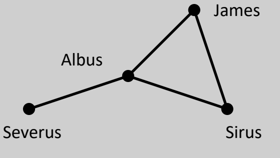

A good way to understand a graph is to draw a diagram of it, by placing a dot for each of the vertices, and drawing a line segment or curve between two vertices if that pair is inside of the edge set E. The edges of these diagrams are allowed to cross each other, and the edges need not be drawn straight.

Even after introducing a couple of graphs, we can already see some interesting topological properties. For instance, the complete graph on 5 vertices can only be drawn with edges that cross, whereas every graph that represents a platonic

3 Some Common Graphs

It is good to have some examples of graphs in mind before we go around discussing the theory of graphs.

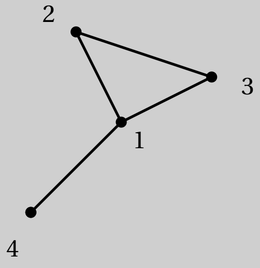

*Graph:*One way to construct a graph is to build it by hand. For example, we can give a graph four vertices and 4 edges by specifying

$$V = \{1, 2, 3, 4\} \quad E = \{12, 23, 13, 14\}.$$

While mathematically precise, this presentation is not very intuitive, and so when we want to specify a graph in this text, we will usually just draw a picture of that graph, and assume that it has some explicit (although unwritten) labeling of the vertices and edges.

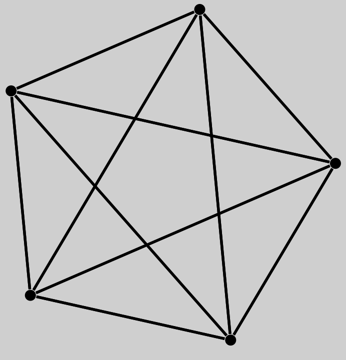

*Complete Graph:*One especially important family of graphs are the *complete graphs*, which are saturated with edges. Let $n \in \mathbb{N}$ be a natural number. The *complete graph on n vertices*, denoted as K_n, is the graph with n vertices and an edge between every pair of vertices. Since every edge corresponds to a choice of 2 vertices, and in the complete graph we've chosen every pair, it follows that

$$|E_{K_n}| = \binom{n}{2}.$$

One reason we call this graph complete is because every graph on fewer than n vertices is a subgraph of K_n.

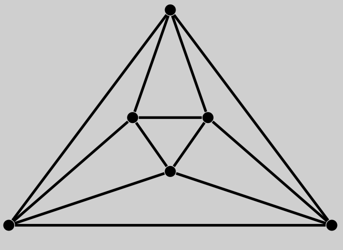

*Octahedron:*Graphs naturally arise from other branches of mathematics. For example, every polyhedra in $\mathbb{R}^3$ determines a graph by its edges and vertices. In the diagram on the left, we see the graph corresponding to the octahedron. Later, we will classify the platonic solids by understanding the combinatorics of their corresponding graphs. Notice that a polyhedron has more data than just its underlying graph, as it also knows what combinations of edges and vertices make up faces of the polyhedron.

solid can be given by a planar drawing without edges crossing. Before we get to describing topological properties of graph, let's first describe some of the combinatorial data attached to a graph.

In this text, the vertices of a graph will be labelled with lowercase letters u, v, w, x, y. For edge, we will either use the letters e, f; or we will refer to an edge with endpoints u and v by the pair uv. With this notation, we say that uv is an edge if either $(u, v) \in E$ or $(v, u) \in E$. The simplest piece of data that we can ask about a vertex in a graph is the number of neighbors it has.

4 Definition *Neighborhood*

Fix some vertex v. Then the subset of vertices which are connected to v by an edge is called the *neighborhood* of v, and will be denoted

$$N(v) := \{u \in V \mid (u, v) \text{ or } (v, u) \in E\}$$

The number of edges connected to v is called the *degree* of v and is denoted

$$\deg(v) := |\{u \in N(v)\}|.$$

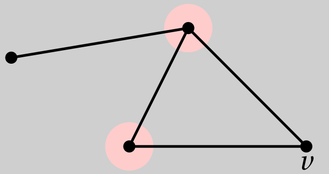

$\deg(v) = 2$

We can already capture a lot of data about a graph simply by knowing the degrees of the vertices in it. For example, if $|V(G)| = n$, and every vertex has degree $|n-1|$, then it must be the case that $G = K_n$. If instead every vertex has degree 2, then it must be the case that G is a collection of cycles. Both of these give us example of *regular graphs*, which are graphs whose vertices all have the same degree. A natural family of regular graphs that show up are the Platonic solids.

The degree of a vertex provides *local* information about the graph. This means that knowing the degree of a particular vertex $v \in G$ doesn't tell us a lot about the graph as a whole — we only learn about a small portion of the graph around v. This is contrasted with quantities like $|E(G)|$ and $|V(G)|$, which tell us *global* information about our graph. In many ways, topology is the study of how local information can be meaningfully assembled into global data.

5 Claim *Average vertex degree*

Let $d(G)$ be the average degree of the vertices of a graph. The total number of edges in G is:

$$2|E| = d(G)|V|.$$

Here, we are using averaging to take local information – the degree – and obtaining some information about the entire graph. The rough idea of proof is that every edge contributes $+1$ to the degree of each of its ends, so that the sum of all the degrees of the vertices will be twice the number of edges. We therefore obtain the equality

$$2|E| = \sum_{v \in V} \deg(V) = d(G)|V|.$$

But let's try to make this argument a little more mathematically watertight.

Proof: From the definition of average degree

$$d(G)|V| = \sum_{v \in V} \deg(v) = \sum_{v \in V} \sum_{u \in V} \delta_E(u,v) + \delta(v,u)$$

where $\delta_E(u,v) = 1$ if $(u,v) \in E$, and 0 otherwise. The reason both orders are necessary is because our definition of edge used ordered pairs (u,v).

$$= \sum_{(u,v) \in V \times V} \delta_E(u,v) + \delta_E(v,u)$$

As $\delta(u,v)$ only takes a value of 1 on sets (u,v) that are in E, this term counts the number of edges

$$=2|E|$$

A good example keep in mind for this lemma are the complete graphs. These n-vertices of these graphs each have $(n-1)$ neighbors, so by Claim 5 the number of edges in a complete graph is $\frac{n(n-1)}{2}$. We also get this strange, but surprisingly useful, corollary.

Corollary 6
An Even number of odd vertices

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

The proof is left to Exercise P1. One general mathematical principle is that a large, complicated object (like a graph) can be understood by asking questions about its substructures. For instance, by asking how many edges or how many vertices are in a graph, we can begin to get an understanding of its global structure. In the setting of graphs a particularly useful substructure to study are subsets of the edges and vertices which themselves make *subgraphs* of G. Here are two especially important types of subgraphs which we will look at throughout the course.

Definition 7
Paths and Cycles

Let G be a graph. A *path* in G is a sequence of distinct vertices $\{v_i\}_{i=0}^n$ 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. We say that a path *starts at* v_0 *and ends at* v_k. A *cycle* in G is a path P such that the first and last vertex of P share a common edge (which is not already in the path P)

Paths and cycles are useful subgraphs to study as they can probe *global* data about the graph. If you know that a graph has a path between any two points, or that a graph contains no cycles, then you've learned something interesting about the global topological information of that graph.

8 Example *Paths and Cycles*

In the drawn example, we have a path 4 highlighted in red, and a cycle of length 3 highlighted in blue. The longest path that one can draw would have length 5, and the largest cycle has length 6.

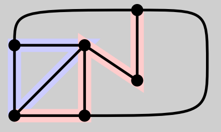

Since every edge constitutes a very short path, it will necessarily be the case that a non-trivial graph has some paths in it. However, it need not be the case that a graph have any cycles.

9 Example

A *tree* is a graph with a unique path between any 2 points. This is a global property of a graph. If we fix a vertex v in the tree (called a root), then the set of paths that start at v are in bijection with the vertices in the tree. In Exercise (P 3), several different interesting properties of a graph are shown to be equivalent to being a tree.

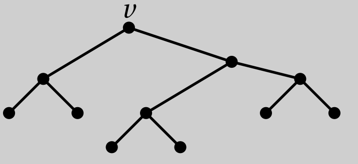

The length of the longest path in a graph G is telling us something about the global structure of the graph– you need to look at the entire graph to find the longest path. As before, we look at how we can assemble local information – in this case, the degrees of the vertices – to provide some information on this invariant.

10 Claim

Let δ be the minimal degree of the vertices in G. Then G has a path of length δ.

Proof: Start with any path P. Let $\{v_i\}$ be the vertices of the path. Let's look at initial vertex of the path , v_0. Suppose that we can find a vertex $w \in N(v)$, which is not already contained within the path P. Take this vertex w, and append it onto P to build a new path wP which is one vertex longer.

We can continue this process unless we've grown our initial path to a path P' which can no longer be extended. This will happen if every point in the neighborhood of the end of the path is contained within the path. For this path P', we have $N(v_0) \subset V(P')$. Since $\delta \leq |N(v_0)| < |V(P)|$, we conclude that $\delta \leq |E(P)|$.

2

Connectivity

A graph is connected if you can travel between any two vertices along the edges of the graph. We also look at several different quantitative measurements of connectivity.

Topology is the study of mathematical structures which have a notion of how their parts are connected to eachother. A graph is an example of a topological structure, as each vertex "knows" which neighbors it is connected to. By travelling from vertex to vertex along edges, we can explore questions about the global connectivity of a graph.

Definition 11
Vertex Connectedness

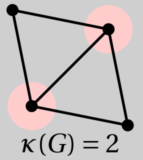

A graph G is called *connected* if for any two vertices v and w, there exists a path from v to w. A disconnecting set $U \subset V$ is a set of vertices with the property that the graph $G \setminus U$ is disconnected. The *connectivity of* G, denoted $\kappa(G)$, is the size of the smallest disconnecting set of G.

The connectivity of a graph is an important measurement for applications. For example, if we are building a power grid for a utility, then whether or not the power can be delivered to the entire network from a single node is determined by the connectedness of the underlying graph. The connectivity of the network is a slightly stronger measure, which tells us the maximal number of nodes of the network can fail and still have the network remain connected. We have a property of a graph which interpolates between the connectivity number and the connected property. We say that a graph is k-connected if no vertex set of size k disconnects the graph G. Equivalently, a graph is k-connected if its connectivity is at least k.

Example 12
Bottlenecks and Connectivity

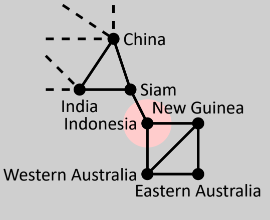

If $H \subset G$ is a subgraph, one can measure the size of the smallest set it takes to disconnect H from the remainder of G. This will always be at least the connectivity of the entire graph. In the game of *Risk*, the continent of Australia is especially prized because of it's low vertex connectivity to the remainder of the graph.

Connectivity has a strange relationship with topology. Whether or not a graph is connected is a *topological* property of the graph. However, the connectivity of a graph is *not* a topological property. An example that demostrates that connectivity is not a topological property are the graphs of the triangle and the square, which are topologically indistinguishable. However, the triangle is three connected, and the square is only 2 connected! We will later see in Station 3 exactly why connectivity is not a topological property. You can observe that even making small modifications to the graph (like making the edge of a graph into a path,) can greatly change to the connectivity of a graph.

It is possible to get an estimate on the connectedness of an average component of the graph by knowing the average degree of the vertices of a graph. In short, if you have a lot of edges, then we expect the graph have a connected component with high connectivity (see Example 14.)

There is no particular reason why we choose to use deletion of vertices to define the connectivity of a graph. We could have instead used edges to get a measure of connectivity, and define the *edge connectivity* as the minimal number of edges that you must remove to disconnect the graph. Somewhat surprisingly, the vertex connectivity and the edge connectivity are usually not related to each other.

14 Example

Dumbbell Graph

Consider the dumbbell graph which is created by taking two K_n and mutually connecting them to a new vertex v. In order to disconnect this graph by removing edges, you need to remove at least n edges. However, to vertex disconnect the graph, it suffices to take out the middle vertex v.

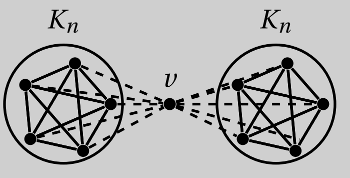

The discrepancy between the edge and vertex connectivity is due to the fact that a concentrating edges onto a single vertex gives it low vertex connectivity, but has little impact on edge connectivity.

Spectral Graph Theory

A more quantitative measure of connectivity examines the average distance squared between two vertices in the graph. We now give a nice algebraic characterization of this metric. Let $V = \{v_1, \dots, v_n\}$ be the vertices of a graph G. The *adjacency matrix of* G is the $n \times n$ matrix A whose ij entry is 1 if $v_i v_j$ is an edge, and 0 otherwise. Define the *degree matrix* D to be the diagonal matrix whose ith diagonal entry is $\deg(v_i)$. Finally, we define the *Laplacian* of the graph G to be the matrix

$$L = D - A.$$

The eigenvalues of L bound the connectivity of G, and in applications gives a more nuanced definition for connectivity.

As demonstrated by the dumbell graph (Example 14), the vertex connectivity tells us where the bottlenecks are in our graph. Another way of counting the bottlenecks in a graph would be to ask how many independent paths there are betweeen two vertices, as the presence of a bottleneck in the graph will force this number to be small.

14 Mader's Theorem

Every graph with average vertex degree $d(g)$ greater than $4k$ contains a k-connected subgraph.

The idea of this proof is to induct on the number of vertices, where the vertex we remove during the inductive step is one of low degree. We start by massaging $d(G) \geq 4k$ to obtain the inequalities

$$|V| \geq 2k-1 \qquad\qquad |E| \geq (2k-3)(|V|-k+1)+1$$

For the first inequality, notice that if the average degree is $4k$, there must be a vertex with degree at least $4k$; therefore it has $4k$ neighbors. Therefore the number of vertices is at least $4k$. For the second inequality, we apply Claim ❺,

$$|E| = d(G)|V| \geq 4k|V| \geq (2k-3)(|V|-k+1)+1$$

With these two new inequalities, we run an inductive argument on the number of vertices in the graph. In the base case, G is a graph on $2k-1$ vertices and therefore is complete, for which the claim follows trivially.

For the inductive step, we split into two cases based on whether there exists a vertex of low-degree. If v has low degree, then $G \setminus v$ will have *larger* average degree, and by induction we can still run our argument.

Case 1: Suppose that there exists a vertex of degree less than $2k-3$. Then even after removing this vertex, our inductive hypothesis holds as we've decreased the number of vertices by 1 but only removed at most $(2k-3)$ edges.

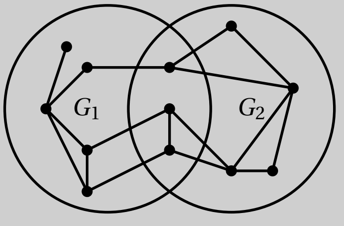

Case 2: Suppose that the inequality is not sharp, and we cannot find a vertex of degree $2k-3$ or less. Let's assume for contradiction that G is not k-connected. Then there are two subgraphs $G_1, G_2 \subset G$ such that $G_1 \cap G_2$ has fewer than k vertices. Every vertex in $G_1 \setminus G_2$ has neighbors only in G_1. Since the minimal degree of each vertex is $2k-2$, we have that G_1 has at least $2k-1$ vertices. Similarly, G_2 has $2k-1$ vertices. So, both the graphs G_i satisfy the first inequality for our induction hypothesis.

As $E = E_1 \cup E_2$, we get $|E| \leq |E_1| + |E_2|$. We additionally know that $|V_1| + |V_2| \leq |V| + k$. Combining these inequalities gives

$$(2k-3)(|V_1|-k+1)+1) + (2k-3)(|V_2|) - k+1) \leq (2k-3)(|V|-k+1)+1 \leq |E| \leq |E_1| + |E_2|$$

from which we conclude that one of the G_i satisfy the induction hypothesis, and therefore contains a k-connected subgraph.

A graph is *k-path connected* if any two vertices v and w can be joined by at least k disjoint paths.

15 Definition

Average vertex degree

You can check that the dumbell graph (Example **14**) is only 1-path connected, as there is only one way to get from the left portion of the graph to the right portion. Menger's theorem (Theorem **17**) tells us that this is generally the case: the path connectivity of a graph is equal to its vertex connectivity.

Menger's theorem seems a bit strange on first reading because of a contrast in the definition of path connectedness and vertex connectedness. The path connectivity looks to *maximize* a set of independent paths, while the vertex connectivity is tries to *minimize* the size of a disconnecting set.

These types of statements are common in combinatorics related to optimization. Max-Min properties extend to many combinatorial objects beyond graphs, and the Max-Min principle has several statements which are all relevant to optimization of networks. Some equivalent theorems to Menger's theorem include the Max-flow Min-cut theorem, König's theorem, Dilworth's theorem, and Hall's theorem.

Max-Min

One application of Menger's theorem is to classify the structure of graphs with connectivity 2.

Suppose that $\kappa(G) \geq 2$. Let v, w be two vertices in G. Create a new graph, $G \cup_{v,w} P$ by attaching a path of length at to G, whose endpoints are v and w. Then $\kappa(G \cup_{v,w} P) \geq 2$.
If additionally we require that the path have at least 3 vertices, then $\kappa(G \cup_{v,w} P) = 2$.

16 Lemma

Adding paths preserves $\kappa(G) \geq 2$

Proof: Suppose for contradiction that $\kappa(G \cup_{v,w} P) = 1$. Then there is a disconnecting vertex u so that $G \cup_{v,w} P \setminus \{u\}$ is disconnected. It must be the case that $u \in P$, as otherwise $G \setminus \{u\}$ would be disconnected, contradicting that the 2-connectivity of G. If P has only two vertices, $\{v, w\}$ we are done (as every candidate vertex for u needs to lie outside of G, but $V(G \cup_{w,v} P) = G(P)$. Therefore the length of P must be at least 2. The removal of u from P separates the path into two 2 components. Each of those components is connected to G by their ends, and therefore $G \cup_{w,v} P \setminus \{u\}$ is still connected. u fails to be a disconnecting vertex for $G \cup_{w,v} P$, contradicting our assumption.

It remains to show that if the length of P is at least 2, that there exists a disconnecting set for $G \cup_{w,v} P$ of size 2. Observe that $G \cup_{w,v} P \setminus \{w, v\}$ is disconnected if there are any vertices in P besides w or v.

17 Menger's Theorem

The path connectivity of a graph is equal to the vertex connectivity of the graph.

We first show that the path connectedness is at most the vertex connectedness. Suppose that $v_1, \ldots, v_k$ form a disconnecting set which separates G into two components G_1 and G_2. Pick a vertex $u_1 \in G_1$, and $u_2 \in G_2$. There can be at most k disjoint paths between u_1 and u_2, as each path must use one of the k points in the intersection. This shows that the path connectivity is less than the vertex connectivity.

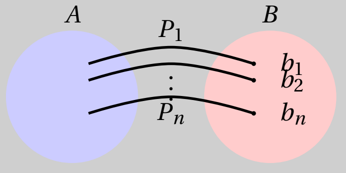

Setup: To show that the path connectedness is at least the vertex-connectedness, we will prove a stronger statement. Suppose that G is k-connected, and let A and B be disjoint subgraphs of G. We will show that whenever we have a collection of fewer than k disjoint paths from A to B, we can find a larger connection of disjoint paths from A to B. Let's set up some notation for this.

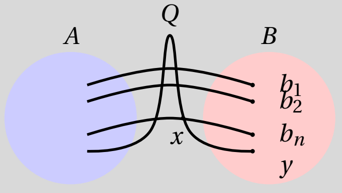

Claim (*Inductive statement for Menger's Thm.*). Suppose that A and B are subgraphs of G each containing at least k vertices. Let $b_1, \ldots, b_n$ be vertices of B, with $n < k$. Let $\mathscr{P}_n = \{P_1, \ldots, P_n\}$ be a collection of disjoint paths G, which

- only intersect A at their left endpoints,
- only intersect B at their right endpoints, which are the specified vertices b_i.

Then there exists a point $y \in B$ and collection $\mathscr{P}_{n+1}$ of $n+1$ disjoint paths in $G \setminus (A \cup B))$ which satisfy the above conditions, with right endpoints $b_1, \ldots, b_n, y$.

We will prove this by inducting on the size of $G \setminus B$. When B is all of $G \setminus A$, then each path P_i consists of a single edge, and k-connectedness gaurentees that A cannot be separated from B by fewer than the removal of k-vertices. For our inductive step, let us assume that the claim holds for every subgraph B' containing B. Let's look at our current graph and subgraph B. Now, we will randomly construct a new path Q from A with a random endpoint in B. If this path is disjoint from the $\mathscr{P}_n$, then we are done. Now, we use the inductive hypothesis by expanding the subgraph B to include a point from the path Q. Let x be the final point where the path Q intersects the collection $\mathscr{P}_n$, and without loss of generality we will assume that x lies on the path P_n.

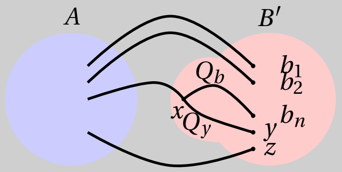

Denote by Q_y the segment of Q going from x to y, and Q_b the segment of Q going from x to b_n. Now enlarge B by including the path Q_b and Q_y,

$$B' := B \cup Q_b \cup Q_y.$$

The end points $b_1, \ldots, b_{n-1}, x$ satisfy the conditions of the claim. By our inductive hypothesis, there exist disjoint paths $\mathscr{P}'_{n+1}$ going from A to end points $b_1, \ldots, b_{n-1}, x, z$ in B'.

At this point we've found a subgraph B' containing B, and we would like to reduce down to B. We break into different cases based on the location of the point z.

Case 1: In the easy case, z belongs to our original set B. In this case, replace y by z in the original step. The paths $P'_1, \ldots P'_{n-1}$ and P'_{n+1} with endpoints $b_1, \ldots, b_{n-1}, z$ are disjoint. To create a final path with endpoint on b_n, we take the concatenation of the path P'_n with endpoint x, and the path Q_b. Since $Q_b \subset B'$, it is disjoint from all of the P'_i we've constructed so far. This gives us the collection $\mathscr{P}_{n+1}$

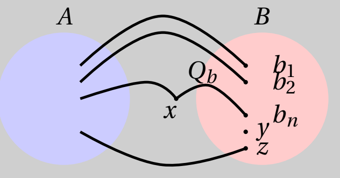

more difficult case is when z only belongs to the enlargement B'. Then $z \in B' \setminus B = Q_b \sqcup Q_y$. By construction, the paths Q_b and Q_y are disjoint, so either $z \in Q_b$ or $z \in Q_y$.

Case 2a Suppose that $z \in Q_y$. Then consider the paths

- P''_n, which is P'_n concatenated with Q_b, and
- P''_{n+1}, which is P'_{n+1} concatenated with the portion of Q_y lying after z.

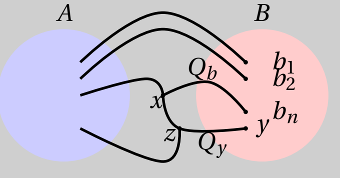

These two paths are disjoint from the $P'_1, \ldots, P'_{n-1}$, and are additionally disjoint for eachother. These paths are seen to have interior vertices which are disjoint from B, and by construction have left endpoints in A. The $\{P'_1, \ldots, P'_{n-1}, P''_n, P''_{n+1}\}$ satisfy the conditions of the claim.

Case 2b Alternatively, it may be the case that z lies on the path Q_b. T

- P''_n, which is P'_{n+1} concatenated with Q_b, and
- P''_{n+1}, which is P'_n concatenated with the portion of Q_y lying after z.

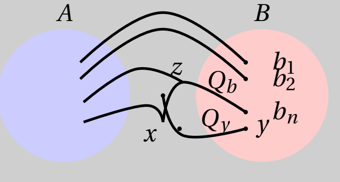

As in the previous case, the paths $\{P'_1, \ldots, P'_{n-1}, P''_n, P''_{n+1}\}$ satisfy the conditions of the claim.

This gives us a method to build larger 2-connected graphs from smaller 2-connected graphs. This can be strenghthend to a characterization of 2-connected graphs.

Theorem 18 *Characterization of 2-connected graphs*

Let G be a 2-connected graph. Then either

- There is a 2-connected graph H so that G may be obtained by attaching a path of H and
$$G = H \cup_{v,w} P.$$
- G is a cycle.

Proof: Suppose that G is not a cycle. Since G is 2-connected, G contains at least a cycle. This means that G contains a 2-connected subgraph, and so we can find a maximal 2-connected proper subgraph H. Here, maximal means that there does not exist another proper 2-connected subgraph H' containing H. We would like to show that G is obtained by attaching a path onto the graph H.

Look at a vertex $v \in G \setminus H$, and a vertex $w \in H$. As G is 2-connected, Menger's theorem ensures that there exists disjoint paths P_1, P_2 from v to w. This gives two paths from w to that are contained in H. $H \cup P$ is a 2-connected subgraph of G. It cannot be the case that $H \cup P$ is a proper subgraph, because H was assumed to be the maximal proper 2-connected subgraph of G. It must then be the case that $G = H \cup P$, concluding the proof.

This gives us way to build up all 2-connected graphs.

Corollary 19

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

We will have to develop a few more tools before we're able to get a classification result for 3 connected graphs, and will return to this in Station 5.

Minors and Topological Minors

Subdivision and contraction are two dual operations which modify a graph at an edge. We'll look at how these opertaions preserve and modify the topological and combinatorial properties of a graph.

3

The main goal of topology is to understand what kinds of topological spaces are out there, and classify them. This means not only understanding how to construct topological spaces, but also when two different spaces are equivalent. So far, we've been using graphs as the building blocks for topological spaces.

20 Example

When we draw a graph, we do not care how we draw the edges of the graph – the edges are allowed to be bent, or straight or cross. Only the adjacency relations between those edges and vertices matter to us. For example, both the pentagon and the star drawn represent the same graph, just drawn differently in the plane. They also represent the same topological space: both are a combinatorial representation of the circle, S^1.

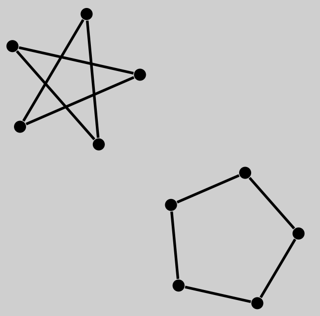

Before we understand when two graphs represent the same topological space, we should understand when two different graphs are equivalent.

21 Definition *Graph Isomorphism*

We say that two graphs G and H are *graph isomorphic* if there exists a bijection $\phi : V(G) \to V(H)$ so that whenever vw is an edge of $E(G)$, then $\phi(v)\phi(w)$ is an edge of $E(H)$.

Isomorphism is an equivalence relation on graphs. Figuring out if two graphs are graph isomorphic isn't an easy task – proving that there does not exist a map between two graphs can be very tricky. Usually, we prove the non-existence of an isomorphism between two graphs by assigning *invariants or properties* to our graphs which only depend on the isomorphism class of a graph. This allows us to distinguish non-isomorphic graphs.

A simple invariant would be to count the number of edges in a graph. clearly two graphs cannot be graph isomorphic if they have a different number of edges or vertices. However, there are plenty of graphs which have the same number of edges and vertices and are not isomorphic. For example, all trees on n-vertices have $n-1$ edges, and all appear the same when only looking at the number of edges and number of vertices.

The properties of the graph that we have studied so far (connectivity of a graph, maximal vertex degree) all can be used to distinguish graphs, and when we develop new properties and invariants throughout this course, I encourage you to construct pairs graphs which can be distinguished by the invariant, and pairs of graph which are non-isomorphic but are not distinguished by the property being discussed. One could weaken the notion of isomorphism to homomorphism, which drops the requirement of Definition 21 that the map of vertices be in bijection. A typical example of a graph homomorphism comes from subgraphs, where one can express the subgraph inclusion $H \subset G$ via an injective graph homomorphism from H to G. Asking whether or not there exists a graph homomorphism between two graphs is not only interesting as a relation between graphs, but can be used to construct invariants between graphs. For example, the graph homomorphisms from $G \to K_n$ can be used to understand the vertex colorings of the graph G.

Graph Homomorphisms

0.3.1 Subdivision and Topological Minors

Graph isomorphism does not describe the whole entire story of topological equivalence. To a topologist, there is no difference between a path and an edge– they both describe “things that look like lines.” In the same light, all cycles, no matter their size, topologically describe the circle. While all graph equivalences give us topological equivalences, the notion of topological equivalence is weaker and a bit more nuanced. One approach to understanding topological equivalence of graphs is to develop some operations which modify a graph, but keep the topology the same.

Definition 22 *Subdivision*

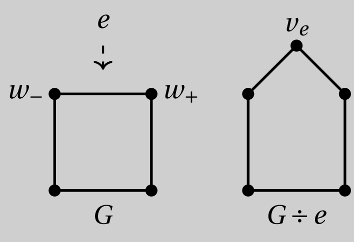

Let G be a graph, and let e be an edge in G, with endpoints w_- and w_+. The *subdivision* of G at e is the graph $G \div e$, whose vertex set has an additional vertex

$$V(G \div e) = V(G) \cup \{v_e\}$$

and whose edge set connects this new vertex connecting to the ends of e:

$$E(G \div e) = V(G) \setminus \{e\} \cup \{v_e w_-, v_e w_+\}.$$

If H is obtained from G by a sequence 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.

While the definition of subdivision may look intimidating, it is a visually intuitive process. One starts with a graph, selects an edge, and then adds a new vertex into the middle of that edge.

Every cycle is a subdivision of the basic cycle on 3 vertices. Subdividing any edge of a cycle produces the next larger cycle, so one may say that whenever $m < n$, C_m is a topological minor of C_n. The set of graphs whose topology matches that of the circle can be described using subdivisions as TC_3.

23 Example

K_5 Subdivision

Graphs contain many more graphs as topological minors than subgraphs. Notice that every subgraph is a topological minor, but it is usually not the case that a topological minor is a subgraph. The graph drawn on the right is a subdivision of K_5, obtained by taking every edge and subdividing it. Since this graph is a subdivision of K_5, it contains every graph on five vertices as a topological minor. However, it does not even contain a triangle as a subgraph. In fact, every cycle in this graph has even length, which is a pretty uncommon property for a graph to have.

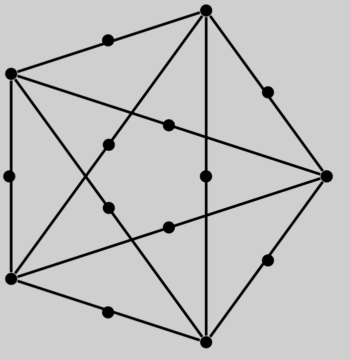

From a combinatorial perspective, subdivision is a bit confusing. It doesn't preserve many graph invariants: for example, whenever one subdivides a graph, the vertex connectivity of the graph is reduced down to at most two. While things like the maximal degree of a vertex are preserved, almost every other combinatorial properties (like edge or vertex connectivity, coloring) is thrown out of the window. Furthermore, subdivision only serves to make a graph more complicated by adding in additional vertices and edge.

However, subdivision satisfies all the condition that we value as topologists: when you subdivide an edge, you are allowed to stretch paths out, but not allowed to break them or create new ones. We will therefore say that the *topological* properties of a graph are those which are preserved under subdivision. If H is a subdivision of G, we will write $G \Rightarrow H$.

This doesn't form an equivalence relation on graphs, as the relation is rarely reflexive: if $G \Rightarrow H$ then $H \Rightarrow G$, it must be the case that $G = H$. However, there is a general trick for construcing equivalence relations out of operations which do not satisfy the reflexive axiom.

Claim 24
Homeomorphic

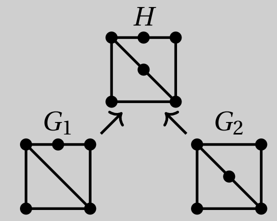

Graphs G_1 and G_2 are *homeomorphic* or *topologically equivalent* if there exists a *common subdivision* H so that

$$G_1 \Rightarrow H \text{ and } G_2 \Rightarrow H.$$

We then write $G_1 \simeq G_2$. Homeomorphism is an equivalence relation.

As a general rule, subdivision preserves the topological properties of a graph, but loses a lot of the combinatorial data. In particular, the operation of subdivision takes a graph G, and produces a slightly sparser, more complicated graph $G \div e$. This means that the operation of subdivision is not particularly useful for proof techniques. In short: subdivision is good for topologists, but bad for graph theory.

0.3.2 Contraction and Minors

A slightly more aggressive type of graph deformation is contraction along and edge, which one can use to invert the operation of subdivision.

Definition 25
Contraction

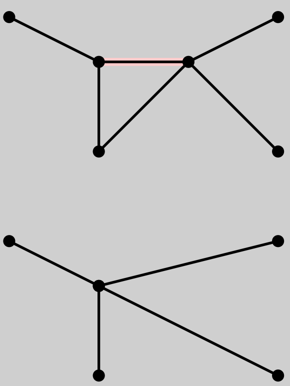

Let G be a graph, and let e an edge in G with ends w_+ and w_-. Then define the *contraction* G/e to be the graph with a new vertex v_e

$$V(G/e) = V(G) \setminus \{v_-, v_+\} \cup \{v_e\}$$

whose edges are given by neighbors of e,

$$E(G/e) = E(G) \cup \{v_e v \mid e \in N(w_-) \cup N(w_+)\}.$$

A graph G is a *minor* of H is G is a subgraph of a contraction of H. When G is a minor of H, we write $G \preceq H$.

The graph G/e is in some sense "simpler" than the graph G in that G/e has fewer edges and vertices than G. Contraction is an opposite operation to subdivision, in that every subdivision can be undone by a contraction. However it is not the case that every contraction can be undone by a subdivision. For instance, the contraction given in the above example is not a contraction which can be undone by subdivision. One also sees this relation between subdivision and contraction when comparing the topological minors of a graph to their minors.

While every topological minor is a minor, it is not the case that every minor is a topological minor. A famous example of a graph which contains many minors, but not very many topological minors, is the *Peterson graph*. The Peterson graph can be contracted to K_5. As a result, the peterson graph contains every graph on 5 or fewer vertices as a topological minor. However, it does not contain K_5 as a topological minor. Compare this to **?**? , which contains K_5 as a topological minor.

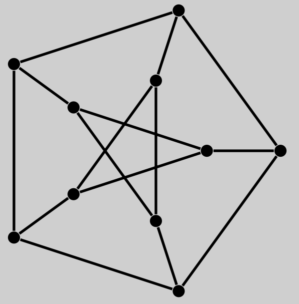

26 Example

Peterson Graph

Subdivision preserves the paths and cycles which exist in a graph, while contracting an edge can possibly destroy the paths and cycles which contain that edge. For this reason, we do not usually think of contraction as an operation which plays well with the topological aspects of graph theory. Consider, for instance, that every connected graph can eventually be contracted down to a point! This means that contraction doesn't give meaningful equivalence relations between graphs.

This does not mean that contraction is less useful for studying graphs. In fact, many of the combinatorial properties of graphs that we care about are not related to the topological type of the graph. Furthermore, because contraction simplifies a graph by lowering the number of edges and vertices, it becomes a valuable tool for inspecting the combinatorial properties of a graph. Contraction will become a more powerful tool for studying the combinatorial properties of a graph. There are even a few topological properties of a graph which are well behaved under contraction. For instance, whether or not a graph can be drawn in the plane without edges crossing (Station 1) is a property which is preserved under the operation of contraction. Additionally, the number of ways that we can color a graph so that no two adjacent vertices have the same color is well behaved under the operation of contraction. We will return to this at Station 2.

Planarity

While subdivision seems to make a graph sparser (and less connected), contraction makes a graph denser. At Example 27 we will see how this relation can be used to compute the reliability of a graph, which is another measure of connectivity. This is an example of graph properties which can be calculated via a recursion relation called *deletion-contraction*. Examples of properties that can be computed by deletion contraction include the Pott-Ising model which describes spins on a lattice, or the number of spanning forests in a graph. The Tutte Polynomial of a graph is an invariant of the graph which universally captures all such properties.

Deletion-Contraction

27 Application: Reliability Polynomial

The *reliability* of G is a function $R_G(p)$, which calculates the probability that G remains connected if we remove each edge of G with probability $(1-p)$.

For instance, a tree does not make a very reliable network– for a tree to fail, all we need is for one of its edges to give out. This means that the reliability of the tree is $p^{|E|}$.
Naïvely, one computes the reliability polynomial by looking at every subgraph of the G, checking if it is connected, and taking the average connectivity of these states. We can use contraction to obtain a nice recurrence relation which computes the reliability.

Proposition (*Recursive computation of R_G*). Let G be a graph. Then the reliability polynomial can be computed by the relation:

$$pR_{G/e}(p) + (1-p)R_{G\setminus e}.$$

Proof: Take any edge e in G. When we start removing edges of G, we start with this edge, and either this edge fails with probability p, or it remains reliable and is not deleted.

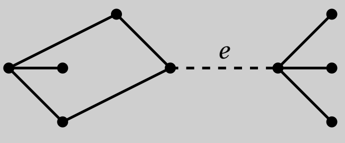

Case 1: The edge e fails, which occurs with probability p. Even though we remove e there is still the possibility that the graph is connected. The probability that the graph remains connected after the possible removal of more edges is given by the reliability of the remaining graph. So, all of the cases where the edge e is deleted contributes contributes $p \cdot R_{G\setminus e}(p)$ to the reliability polynomial.

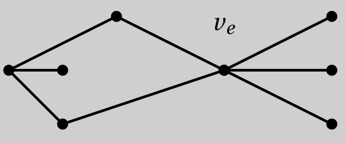

Case 2: With probability $(1-p)$, the edge will not fail. This means that the two vertices at the edge of the graph are gaurenteed to remain connected. One way to represent this configuration is to take the graph G and contract it along the edge e. The contribution from states containing the edge e is $(1-p)\cdot R_{G/e}(p)$.

These two cases are disjoint, and exhaust all the possibilities for how edges of G may fail or remain. Therefore the probability that G is connected after removing each edge with probability $(1-p)$ is

$$p \cdot R_{G/e}(p) + (1-p)\cdot R_{G\setminus e}.$$

As a corollar, we learn something about the structure of $R_G(p)$ by applying induction on the number of edges in G.

Corollary.The Reliability of a graph is a polynomial in the variable p.

The Cycle Space

By representing the vertex set and edge set of a graph with vector spaces, we can recast some of the graph properties we previously explored (like path-connectivity) in the language of linear algebra. We show that the dimension of the cycle space is a topological invariant of a graph.

4

In this section we introduce a way of encoding data that we'll use in the future. The idea is to take some of the combinatorial data of the graph and upgrade it to algebraic data. Here is a question that frames the tools we'll be developing in this section.

28 Example
How to count cycles

What is the proper way to count the number of cycles in a graph G? One way to get a count of cycles is to list them all out: for instance, this graph has 3 cycles. However, it looks like the large cycle can be drawn by combining two smaller interior cycles. We would like to say that this graph has 2 essential cycles, and that the third comes about from the fact that when 2 cycles share an edge, you get a third cycle for free.

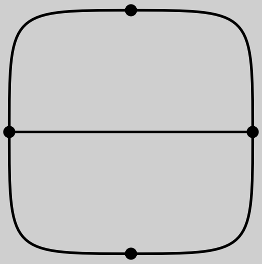

The problem in counting cycles is that we do not have a good notion of what it means to "add" together cycles in a graph. We do not even know how to add together edges. One way we could make sense of taking sums of edges would be to work instead with a vector space $\mathscr{E}$, whose basis is given by the set of edges. In order to this to have any meaning, we'd have to imbue this vector space with some information so that it remembers what the graph G was. This means that we are going to have to abstract our definition of a graph a little bit.

29 Claim
Graphs a 1-complexes

The data of a directed graph equivalent to a set of vertices V, a set of edges E, and two maps, called the left and right boundary maps:

$$\partial_l : E \to V, \partial_r : E \to V.$$

When we work with sets, there is not a way to encode both the left and right endpoint maps into the same map. However, if we upgrade the set E into a vector space, we can encode both of these boundary maps into one function instead.

Sets and Vector Spaces

There are some portions of this course where we'll foray into the world of algebra, and try to build up algebraic tools which capture the geometric of what we are trying to construct. The majority of these tools are summarized in the appendix. In general, these discussions are rather abstract, but are useful to see in parallel with the topologically inspired problems we are solving.

Let's start with a discussion on the limitation of working with sets and function between sets. Suppose we are given two functions $f_1 : S_1 \to T$ and $f_2 : S_2 \to T$. Then in the language of sets, we can assemble these two maps into one map.

$$f_1 \sqcup f_2 : S_1 \sqcup S_2 \to T$$

$$f(x) = \begin{cases} f_1(x) & \text{if } x \in S_1 \\ f_2(x) & \text{if } x \in S_2 \end{cases}$$

It may help to think of the data of the disjoint union and the relevant maps between sets as fitting nicely into the diagram drawn below.

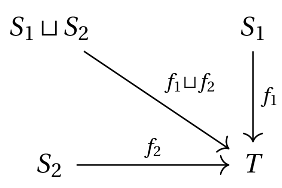

One might say that the disjoint union construct allows us to take maps between sets "combine" them whenever they have the same target. However, if we instead have maps $f_1 : S \to T_1$ and $f_2 : S \to T_2$, there is not a good way to formulate a map from S to $T_1 \sqcup T_2$. The obstacle to the creation of such a map is that it should be multi-valued, but we don't have multi-valued maps between sets.

In order to be able to add together the codomains of maps, we'll need to upgrade our sets to more interesting algebraic objects.

Definition 30
Linearizeation of Sets

Let $S = \{s_1, \ldots, s_k\}$ be a set. Denote by $\mathscr{S}$ the vector space over $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$, with basis labeled by S.

Let $f : S_1 \to S_2$ be a map between sets. When we write $f : \mathscr{S}_1 \to \mathscr{S}_2$, we'll mean the linear map induced from the basis f, which is called the *linearization of* f.

We'll continue to build a dictionary between sets and linear algebra throughout this course. A way to translate from vector spaces back to combinatorics ofsets is to use the relation

$$\dim(\mathscr{S}) = |S|.$$

Let $S = S_1 \sqcup S_2$ be finite sets. Then $\mathscr{S} = \mathscr{S}_1 \oplus \mathscr{S}_2$.

31 Claim

Proof: Recall that $\mathscr{T}_1 \oplus \mathscr{T}_2$ is the set of vectors written as pairs (t_1, t_2), where $t_1 \in \mathscr{T}_1$ and $t_2 \in \mathscr{T}_2$. A basis for $\mathscr{T}_2 \oplus \mathscr{T}_2$ is given by the disjoint union of the basis for $\mathscr{T}_1$ and $\mathscr{T}_2$.

Linear algebra is an extremely flexible piece of mathematical theory. In the land of linear algebra, we can take sums of maps not just along the domain of the map, but additionally along the target. This is an advantage over maps between sets.

Let $f_1 : \mathscr{S} \to \mathscr{T}_1$ and $f_2 : \mathscr{S} \to \mathscr{T}_2$ be two linear maps. Define the sum over the target as $f_1 \oplus f_2 : \mathscr{S} \to \mathscr{T}_1 \oplus T_2$ by

$$(f_1 \oplus f_2)(s) = (f_1(s), f_2(s)).$$

32 Definition

Summing the Target

Let $f_1 : \mathscr{S}_1 \to \mathscr{T}$ and $f_2 : \mathscr{S}_2 \to \mathscr{T}$ be two linear maps. Define the sum over the domain as $f_1 \oplus f_2 : \mathscr{T}_1 \oplus \mathscr{T}_2 \to \mathscr{S}$ to be the map which sends

$$(f_1 \oplus f_2)(s_1, s_2) = f_1(s_1) + f_2(s_2).$$

33 Definition

Summing the Domain

These definitions generalize the disjoint union of two maps between sets. Whenever $f_1 : S_1 \to T, f_2 : S_2 \to T$ are two maps of sets, then $f_1 \sqcup f_2 : S_1 \sqcup S_2 \to T$ has linearization given by $f_1 \oplus f_2 : \mathscr{S}_1 \oplus \mathscr{S}_2 \to \mathscr{T}$.

A return to Graphs

Let's try to use these tools to combine those maps that we were talking about earlier from graph theory.

Given a graph G, define the

- *edge space* $\mathscr{E}$ to be the $\mathbb{F}_2$ vector space generated on a basis E.

34 Definition

– *vertex space* $\mathcal{V}$ to the $\mathbb{F}_2$ vector space generated on a basis V.

– *edge boundary map* $\partial : \mathcal{E} \to \mathcal{V}$ by

$$\partial := \partial_l \oplus \partial_r.$$

In this new setup, elements of $\mathcal{E}$ correspond to subsets of $E(G)$, and the vector addition on E corresponds to symmetric difference. Since each vector $s \in \mathcal{E}$ corresponds to a subset of $E(G)$, we will frequently say that an edge e is in s if e is contained in the corresponding subset. Likewise, we will say the size of a vector s is the number of edges in the corresponding subset.

Example 35

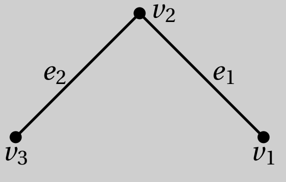

Let's try this new viewpoint by applying it to the graph on the left. The edge space is given by the basis $\{e_1, e_2\}$, and the vertex space has the basis spanned by $\{v_1, v_2, v_3\}$. In this basis, the edge boundary map can be expressed as the matrix:

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} : \mathcal{E} \to \mathcal{V}$$

We have a nice algorithm for writing out the matrix of the *edge differential* if we use the standard basis for $\mathcal{E}$ and $\mathcal{V}$. Each row of the matrix will represent a vertex, and the columns are indexed by the vertices.'' Then whenever an edge in e is has a vertex v as an endpoint, we put a 1 in the corresponding place in the matrix. As a result, the number of ones in each column will be exactly 2 (as each edge has 2 endpoints,) and the number of ones in each row will be the degree of the corresponding vertex.

It will be convenient for us to refer to the elements of $\mathcal{E}$ and $\mathcal{V}$ by simply referring to them by the corresponding vertices and edges, so whenever we write a vertex $v \in \mathcal{V}$, or edge $e \in \mathcal{E}$, we will mean the corresponding basis vector. A good way to verify that we're on the same page for notation is to check that

$$\partial(e_2) = v_3 + v_2$$

for the graph drawn in Example 35.

A reason to upgrade this combinatorial set-up to an algebraic framework is that many graph properties of subsets of edges can be easily stated in terms of properties of the edge differential.

36 Claim *Identifying Paths*

Suppose that $v, w \in V$ are two vertices which belong to the same connected component. Let P be the path which connects v and w. Then $\partial(P) = v - w$. In particular,

$$v - w \in \text{Im}(\partial).$$

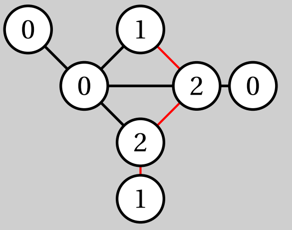

Proof: Let $P \subset E$ be the subset of edges that correspond to a path between v, w. Then $\partial(P)$ is the subset of vertices which belong to P counted with multiplicity based on how many edges they show up in. Since the interior vertices each belong to 2 edges in the path, they get counted with multiplicity 2, which is congruent to 0 mod 2. The boundary vertices, $\{v, w\}$ each get counted once. Therefore, $\partial(P) = \{v, w\} = v - w$.

37 Corollary *Identifying Cycles*

Let $C \subset E$ be a subset of edges that form a cycle, and let $c \in \mathscr{E}(G)$ be the vector corresponding to that cycle. Then $\partial(c) = 0$.

It is not the case that the set of cycles forms a subspace of $\mathscr{E}$, as the sum of two cycles will usually be just the disjoint union of two cycles. However, we can still look at the smallest subspace of $\mathscr{E}$ which contains every cycle.

38 Definition *Cycle Space*

Define $\mathscr{C}$, the *cycle space*, to be the subspace spanned by all the edge sets giving cycles in $\mathscr{E}$.

39 Claim

The cycle space can be related to the boundary by

$$\mathscr{C} = \ker \partial.$$

Proof: The forward direction follows from Claim 36. To show that that every element of the $\ker \partial$ is a sum of cycles, we will use a greedy algorithm which decomposes an element $s \in \ker \partial$ into a sum of cycles.

We induct on the number of edges in $s \in \ker \partial$. Pick any edge e_1 in s. Since $s \in \ker \partial$, there must be another edge $e_2 \in s$ with left endpoint equal to e_1's right endpoint. This process constructs a path, and by finiteness of the graph this path must eventually cross itself. This gives us a cycle $c_1 \in s$, so we have

$$s = s' + c_1$$

where s' has fewer edges than s, and $s' \in \ker \partial$. By our induction hypothesis, we know that s' is a sum of cycles; therefore we can find a decomposition of s into a sum of cycles.

We now can associate to a graph G a subspace $\mathscr{C} \subset \mathscr{E}$ which describes the set of cycles, and this subspace is easy to compute (as it is simply the kernel of a linear map.) We'll use this subspace to give a definition of the number of independent cycles in our graph.

Definition 40
Betti Numbers

Let G be a graph. The *zero Betti number*, denoted by $b_0(G)$, is the number of connected components of G. The *first Betti number*, denoted by $b_1(G) := \dim(\ker(\partial)) = \dim \mathscr{C}$, is the cycle number of G.

These numbers are a topological invariant of our graph. One would hope that the number of cycles or connected components does not vary as we subdivide an edges of our graph.

Theorem 41

Suppose that $H = G \div xy$. Then $b_1(G) = b_1(H)$.

Proof: The remainder of this section is spent on this proof. The goal will be to show that the spaces $\mathscr{C}(H) = \ker(\partial^H)$ and $\mathscr{C}(G) = \ker(\partial^G)$ are isomorphic. We will use our topological intuition to construct maps of vector spaces between $\mathscr{E}(G) \leftrightarrow \mathscr{E}(H)$ and $\mathscr{V}(G) \leftrightarrow \mathscr{V}(H)$. The proof can be broken into a topological portion, and an algebra portion. On the topological part, we need to create maps between vector spaces encoding our intuition on what happens to edges and vertices in the process of subdivision.

$$\begin{array}{ccc} \mathscr{E}(G) & \xrightarrow{\partial^G} & \mathscr{V}(G) \\ \downarrow i_E & & \downarrow i_V \\ \mathscr{E}(H) & \xrightarrow{\partial^H} & \mathscr{V}(H) \end{array} \qquad \begin{array}{ccc} \mathscr{E}(G) & \xrightarrow{\partial^G} & \mathscr{V}(G) \\ \pi_E \uparrow & & \pi_V \uparrow \\ \mathscr{E}(H) & \xrightarrow{\partial^H} & \mathscr{V}(H) \end{array}$$

The second thing we'll need to do is some linear algebra to show that these topologically inspired maps give rise to isomorphisms between the cycle spaces.

42 Claim
Definitions of Maps

Define the following maps between edge and vertex spaces.

- As $V(H) = V(G) \cup \{v_{xy}\}$, there is a natural inclusion $i_V : V(G) \hookrightarrow V(H)$. The map $i_V : \mathscr{V}(G) \hookrightarrow \mathscr{V}(H)$ is the linearization of that map.
- We have that $E(H) = E(G) \cup \{xv_{xy}, v_{xy}y\} \setminus xy$. Define the map $i_V : \mathscr{E}(G) \to \mathscr{E}(H)$ on the basis of $\mathscr{E}(G)$ by

$$i_E(e) = \begin{cases} e & \text{if } e \neq xy \\ xv_{xy} + v_{xy}y & \text{if } e = xy \end{cases}$$

- We slightly modify the map π_E. In the basis $E(H)$ for $\mathscr{E}(H)$, we define the projection map by

$$\pi_E(e) = \begin{cases} e & \text{if } e \neq xv_{xy}, v_{xy}y \\ xy & \text{if } e = xv_{xy} \\ 0 & \text{if } e = v_{xy}y \end{cases}$$

- Define the $\pi_V : \mathscr{V}(H) \to \mathscr{V}(G)$ by defining the values on the basis

$$\pi_V(v) = \begin{cases} v & \text{if } v \neq v_{xy} \\ y & \text{if } v = v_{xy} \end{cases}$$

These maps between edge spaces and vertex spaces seems a tad bit arbitrary, but are are inspired bo the figure to the right. In short, the inclusion map $i_V : \mathscr{E}(G) \to \mathscr{E}(G \div xy)$ works by seinding the edge to the sum of two edges. When undoing a subdivision, we make a choice in the choice of contraction, which gives us the definition for the map π_E and π_V. We chose the contraction which squashes the edge $v_{xy}y$.

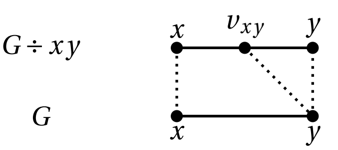

the topology is all fine and good, to show that these have some meaning for cycle spaces we'll show that these map are *chain map*, which means that it preserves the boundary operation. See Station 4 for the full definition.

43 Claim

The following squares commute:

$$\begin{array}{ccc} \mathscr{E}(G) & \xrightarrow{\partial^G} & \mathscr{V}(G) \\ \downarrow i_E & & \downarrow i_V \\ \mathscr{E}(H) & \xrightarrow{\partial^H} & \mathscr{V}(H) \end{array} \qquad \begin{array}{ccc} \mathscr{E}(G) & \xrightarrow{\partial^G} & \mathscr{V}(G) \\ \pi_E \uparrow & & \pi_V \uparrow \\ \mathscr{E}(H) & \xrightarrow{\partial^H} & \mathscr{V}(H) \end{array} .$$

We will first check that this is true for the square given by the inclusion.

$$\begin{aligned}\partial^H(i_E)(e) &= \begin{cases} \partial^H(e) & \text{if } e \neq xy \\ d_H(xv_{xy} + v_{xy}y) & \text{if } e = xy \end{cases} \\ &= \begin{cases} \partial^H(e) & \text{if } e \neq xy \\ x + v_{xy} + y + v_{xy} & \text{if } e = xy \end{cases} \\ &= \begin{cases} i_V\partial^G(e) & \text{if } e \neq xy \\ i_V\partial^G(xy) & \text{if } e = xy \end{cases} \\ &= i_V\partial^G(e)\end{aligned}$$

For the other square,

$$\begin{aligned}\partial^G(\pi_E)(e) &= \begin{cases} \partial^G(e) & \text{if } e \neq xv_{xy} + v_{xy}y \\ \partial^G(xy) & \text{if } e = xv_{xy} \\ \partial^G(0) & \text{if } e = v_{xy}y \end{cases} \\ &= \begin{cases} \pi_V\partial^H(e) & \text{if } e \neq xv_{xy}, v_{xy}y \\ x + y & \text{if } e = xv_{xy} \\ 0 & \text{if } e = v_{xy}y \end{cases} \\ &= \begin{cases} \pi_V\partial^H(e) & \text{if } e \neq xv_{xy}, v_{xy}y \\ \pi_V(x + v_{xy}) & \text{if } e = xv_{xy} \\ & \pi_V(v_{xy} + y)\text{if } e = v_{xy}y \end{cases} \\ &= \pi_V\partial^H(e)\end{aligned}$$

Whenever we have a chain map, we get an induced map between the cycle spaces.

Lemma 44 *Induced Map on Cycle Spaces*

Suppose we have a diagram of maps as given:

$$\begin{array}{ccc} \mathscr{E}(G) & \xrightarrow{\partial^G} & \mathscr{V}(G) \\ \downarrow f_E & & \downarrow f_V \\ \mathscr{E}(H) & \xrightarrow{\partial^H} & \mathscr{V}(H) \end{array}$$

Then the restriction of $(f_V)|_{\ker(\partial^G)} \subset \ker\partial^H$.

Proof: Suppose that $c \in \ker\partial^G$. Then we want to show that $f_E(c) \in \ker\partial^H$. This means that we need to compute $\partial^H f_E(c)$. By commutativity of the diagram,

$$\partial^H f_E(c) = f_V\partial^G(c) = f_V(0) = 0.$$

These results are listed in greater detail in **?**? While we have now produced maps between the cycle spaces of G and H, we still need to show that these maps are isomorphisms.

Claim 45. The restrictions of the map $(i_E)|_{\ker\partial^G} : \ker\partial^G \to \ker\partial^H$ and $(\pi_E)|_{\ker\partial^H} : \ker\partial^H \to \ker\partial^G$ are inverses.

This proof is mostly a computation by hand. Again, we need to check two directions. Let's start with $\pi_E \circ i_E = \mathrm{id}_{\ker\partial^G}$. We have on a basis that

$$\pi_E \circ i_E(e) = \begin{cases} e & \text{if } e \neq xy \\ \pi_E(xv_{xy} + v_{xy}y) & \text{if } e = xy \end{cases}$$
$$= e$$

A similar proof works in the other direction.

$$i_E \circ \pi_E(e) = \begin{cases} e & \text{if } e \neq xy \\ \pi_E(xv_{xy} + v_{xy}y) & \text{if } e = xy \end{cases}$$

Switching basis slightly

$$= \begin{cases} i_E(e) & \text{if } e \neq xv_{xy}, v_{xy}y \\ i_E(xy) & \text{if } xv_{xy} + v_{xy}y \\ i_E(0) & \text{if } e = v_{xy}y \end{cases}$$
$$= \begin{cases} e & \text{if } e \neq xv_{xy}, v_{xy}y \\ xv_{xy} + v_{xy}y & \text{if } e = xv_{xy} + v_{xy}y \\ 0 & \text{if } e = v_{xy}y \end{cases}$$

We've now reduced to three different cases. In the first case, we have an isomorphism on the nose. In the second case, we also have an isomorphism.

Let's see if the third case even occurs. Let c be any element in the kernel of ∂^H. Then if c contains the edge xv_{xy}, it must also contain the edge $v_{xy}y$, as these are the only two edges which contain the vertex v_{xy}. As a result, we can throw the third case out, as it is not contained in the cycle space. A key takeaway is that the map $i_E \circ \pi_E$ is *not* an isomorphism on the edge spaces, but it is an isomorphism on the cycle spaces. ∎

The same proof can be used to show that b_0 is a topological invariant of our space.

Inclusion-Exclusion

One cannot use inclusion/exclusion to compute connected components. However, there is another way.

Let G be a graph, and let $A, B \subset G$ be two subgraphs which *cover* G in the sense that $G = A \cup B$. There are several quantities of G which can be computed via an inclusion/exclusion principle. For instance, both the number of edges and vertices in G can be computed by

$$|V(G)| = |V(A)| + |V(B)| - |V(A \cap B)| \qquad |E(G)| = |E(A)| + |E(B)| - |E(A \cap B)|.$$

This is not surprising, as both the number of edges and number of vertices are really set theoretic properties of the graph.

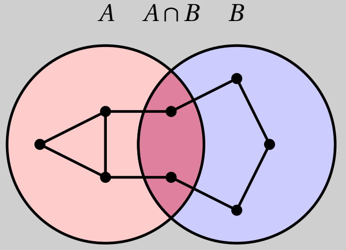

However, the number of connected components does *not* satisfy the inclusion-exclusion property. In the example on the left, $b_0(A) = b_0(B) = 1$, and $b_0(A \cap B) = 2$, but

$$b_0(G) \neq b_0(A) + b_0(B) - b_0(A \cap B) = 0$$

This gives us a striking example that the global property of connectedness does not break up well accross multiple components! The problem in this siuation is that A and B are connected across two different components of their common intersection.

There is a useful example to keep in mind where the number of connected components can be computed via inclusion-exclusion.

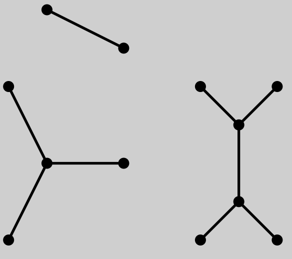

If the graph H is a *forest*, meaning that it is a disjoint union of trees, then the number of connected components can be computed via the formula

$$b_0(H) = |V(H)| + |E(H)|.$$

Since both $V(H)$ and $E(H)$ can be computed with inclusion-exclusion, we see that $b_0(H)$ can be computed in terms of the subgraphs of H whenever H is a forest.

Both of the previous examples point to role that cycles play in computing the number of connected components of G in terms of its subgraphs. One may also try to compute the cycle number $b_1(G)$ by inclusion exclusion, and once again sees that this fails to satisfy the inclusion/exclusion principle. However, there is a simple example to keep in mind when this succeeds. When $b_0(G) = 1$, then the cycle number can be be computed by

$$b_1(G) = |E(H)| - |V(H)| + 1$$

From this, we can make the following inclusion/exclusion computation.

Claim. Suppose that $b_0(A) = b_0(B) = b_0(G) = 1$. Then

$$b_1(G) = b_1(A) + b_1(B) - b_1(A \cap B) + b_0(A \cap B) - 1.$$

This formula seems a bit ad-hoc, but can be understood as partitioning the cycles of G into those cycles which are completely outside of B, completely outside of A, completely contained in the intersection, or passing between A and B through a distinct connected component of the intersection.

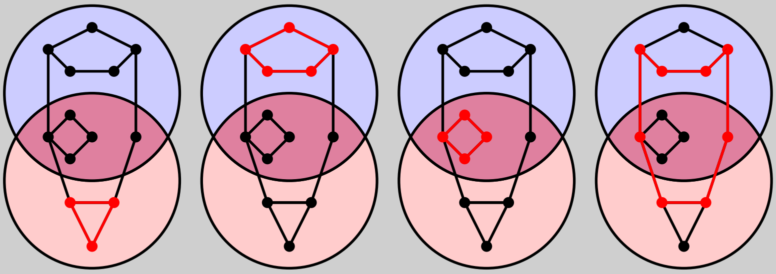

One interpretation of this is the following: if the conneected components of $A \cap B$ causes us to *overcount* $b_0(G)$, then this overcount is realized in the *undercount* of $b_1(G)$. There is a delicate balencing that occurs here, and problematically the numbers do not have enough structure for us to remember all of the balencing that occurs. Fortunately the numbers b_0 and b_1 are only shadows of vector spaces,

$$b_0(G) = \mathcal{V}(G)/\operatorname{Im}(\partial) \qquad b_1(G) = \ker(\partial)$$

and we can reconstruct the decomposition of cycles of G by using maps of vector spaces.

Claim. There exists a map $\delta : \ker(\partial_G) \to \mathcal{V}(A \cap B)/\operatorname{Im}(\partial_{A\cap B})$.

Proof: Let $c \in \ker(\partial_G)$ be a cycle of G. Let $c|_A \in \mathcal{E}(A)$ be the restriction of the cycle to the subgraph A. Note that $c|_A$ will no longer be a cycle. This truncated cycle will now have boundary $\partial_A(c_A)$ at every point where the cycle c crossed over into B, so $\partial(c_A) \in \mathcal{V}(A \cap B)$. We define $\delta(c) := [\partial(c_A)] \in \mathcal{V}(G)/\operatorname{Im}(\partial_G)$.

This can be extended (see **??**) to show that the failure of b_0 to satisfy the inclusion/exclusion principle is exactly equal to the failure of b_1 to satisfy the inclusion/exclusion principle so that:

$$\begin{aligned} 0 = &(b_0(A \cup B) - (b_0(A)) + b_0(B) + b_0(A \cap B)) \\ &- (b_1(A \cup B) - (b_1(A) + b_1(B)) + b_1(A \cap B)) \end{aligned}$$

5 Some Applications to 3-connected Graphs

We prove Tutte's Lemma which states that every 3-connected graph G contains an edge e so that G/e is still 3-connected. While this is an innocuous looking lemma, it allows us to make induction arguments within the subset of graphs that are 3-connected. Using this lemma, we prove that the cycle space of a 3-connected graph has a particularly nice basis.

If we remember that $\mathscr{E}$ and $\mathscr{V}$ come from the graph G they come with additional algebraic structure. The most important structure is the preferred basis (given by the vertices and edges.) While $\mathscr{C}$ does not come with a preferred basis, it still has a preferred generating set: the cycles. However, there is no reason that the set of cycles should be linearly independent.

Example 47 *Faces as Basis*

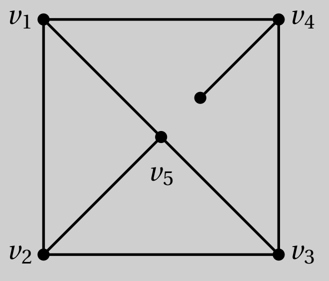

One thing which is interesting to note: when looking at graphs, almost everybody comes up with the same set of cycles for a basis for the cycle space. For instance, in this graph, you would probably say that the three cycles

$$\{v_1 v_2 v_5, v_2 v_3 v_4, v_1 v_5 v_3 v_4\}$$

give a basis for the cycle space. The reason for this is that we are naturally inclined to use the faces from planar drawings as a basis for the cycle space (and it does, indeed, form a basis!)

We could eliminate a few generators by restricting to a smaller set of cycles. 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 48 The induced cycles of G generate the space $\mathscr{C}$.

This set is still a redundant set of cycles, but without further conditions on the graph there is not a clear candidate of generating cycle.

If a graph G is 3-connected, then there are enough cycles that we can be slightly more picky when choosing a basis for the cycle space. The additional connectivity allows us to pick cycles which do not separate the graph into different components.

A cycle C is *non-separating* if $G \setminus C$ is connected.

49 Definition

A good example of non-separating cycles comes from a graph drawn in the plane. A planar graph is called *polyhedral* if it is 3-connected. If a planar graph is polyhedral, then every face gives an example of non-separting cycle. In fact, these are the only non-separating cycles of a planar graph, as every other cycle divides the graph into an interior and exterior region. We will return to this discussion in **?**? . There is a generalization of this characterization which does not need the planarity requirement.

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

50 Theorem
Tutte

Before proving this theorem, it is useful to look at an non-example.

First, let's look at a 2-connected graph where this is not true. In this example, the cycle space is 3 dimensional, so one needs to pick 3 independent cycles to get a basis. One can check by hand that every such collection of 3 cycles will necessarily include a generator which is separating.

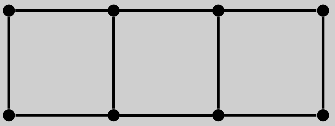

51 Example
Tutte Counterexample

The proof of this theorem relies crucially on a lemma of Tutte that allows us to make induction-type arguments with 3-connected graphs.

Proof: We will use **?**? to prove this by induction on the number of edges, and follow the exposition in [Die00]. Let e be an edge of G so that G/e is 3-connected. By our induction hypothesis, we assume that the theorem holds for G/e.

52 Tutte's Lemma

If G is 3-connected with more than 4 vertices, then there exists an edge $e \in G$ so that G/e is still 3 connected.

Suppose for contradiction there is no such edge. Define the function $I : E \times V \to \mathbb{N}$ by

$$I(uv, w) = \text{Size of the smallest component of } G \setminus \{u, v, w\}$$

Given our hypotheses we will prove that $I(uv, w)$ has no minimal value which is clearly impossible!

First we show that given our hypothesis, there is a pair (uv, w) so that $G \setminus \{u, v, w\}$ is disconnected. Pick any edge uv. By our hypothesis G/uv is 2-connected. Since G/uv is 2-connected but G is not, it must be the case v_{uv} is part of a separating set of G/uv. Let w be any vertex such that $(G/uv) \setminus \{v_{uv}, w\}$ is disconnected. Then u, v, w separate G.

Now we show that I has no minimal value. Pick any uv, w so that $G \setminus \{u, v, w\}$ is disconnected. We will show that $I(uv, w)$ is not minimal.

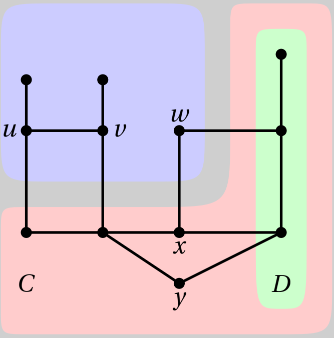

Let C be the smallest component of $G \setminus \{u, v, w\}$. Because all three vertices are necessary to separate the graph, w must have a neighbor in the component C. Call this vertex x. Notice that the neighbors of x are entirely contained in C. By our assumption, G/wx is 2-connected. Therefore, there exists another vertex y so that $G \setminus w, x, y$ is disconnected. We will show that this has a connected component which is smaller than C.

Let D be a component of $G \setminus \{w, y, x\}$ which does not contain uv. x has a neighbor in D, otherwise $G \setminus \{w, y\}$ would be disconnected. Therefore $D \cap C$ is nonempty. Furthermore, every vertex of D is contained in C, because D is disjoint from the connected components of $G \setminus \{u, v, w\}$ containing u and v, and additionally C is a connected subset.

Since D does not contain x, we have that D is a proper subset of C. Therefore D is smaller size than C. So $I(wx, y) < I(uv, w)$.

Corollary (*Characterization of* $\kappa(G) \geq 3$). Every 3-connected graph has a K_4 minor.

The reverse direction holds as well (see Exercise P 11.)

We prove the theorem by relating the vector spaces $\mathscr{C}(G)$ and $\mathscr{C}(G/e)$. Consider the map

$$\pi_E : \mathscr{E}(G) \to \mathscr{E}(G/e)$$
$$\pi_V : \mathscr{V}(G) \to \mathscr{V}(G/e)$$

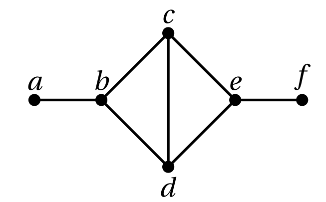

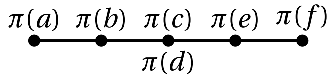

constructed in Claim 43. We know that there is then a map on the spaces of cycles $\pi : \mathscr{C}(G) \to \mathscr{C}(G/e)$. In contrast to the example we considered earlier, it is possible that this map has a kernel given by the triangles which are collapsed to an edge under the contraction. We call these cycles the *fundamental triangles*, and they span the kernel of the map. We can construct chain maps

$$\imath_E : \mathscr{E}(G/e) \to \mathscr{E}(G)$$
$$\imath_V : \mathscr{V}(G/e) \to \mathscr{V}(G)$$

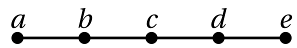

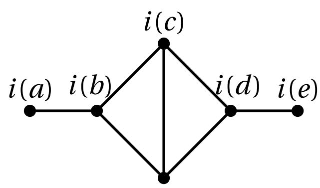

giving rise to a chain map

$$i : \mathscr{C}(G/e) \to \mathscr{C}(G)$$

which is a right inverse to the map $\pi : \mathscr{C}(G) \to \mathscr{C}(G/e)$. Note that there are many different choices that we could have made in constructing the inverse map i.

Corollary (*Characterization of* $\kappa(G) \geq 3$). If $c \in \mathscr{C}(G)$ is a cycle, and $c \in \ker(\pi)$, then c is a triangle with one edge e. The set of such triangles generate $\ker(\pi)$. 53

Let's fix some notation to improve readability of the proof. We will call a non-separating induced cycle *basic*. Cycles which can be written as the sum of basic cycles will be called *good*.

The idea of the proof is to start with a circle $c \in \mathscr{C}(G)$, obtain a good cycle $\pi(c) \in \mathscr{C}(G/e)$, then try to lift this back to a new cycle $i \circ \pi(c)$. Because i is only a right inverse there will be a disagreement between c and the lift $i \circ \pi(c) \neq c$. We will show that this discrepancy is a good cycle, which completes the proof.

Every fundamental triangle is basic. 54 Claim

Proof: Let C_3 be a fundamental triangle. If C_3 separates G, then C_3/e separates G/e. But C_3/e only has 2 vertices and G/e is 3 connected. Therefore C_3 is basic.

Claim 55 If $c' \in \mathscr{C}(G/e)$ is basic, then $i(c') \in \mathscr{C}(G)$ is good.

In the best case, we are in a scenario where

$$(G \setminus i(c'))/e = (G/e) \setminus c',$$

which is connected. It follows that $G \setminus i(C')$ is connected (and we may in fact conclude that $i(C')$ is basic).

The more difficult case to handle is when $i(C')$ only passes through one end of the contracted edge e. To handle this case, we will need to label the vertices near the contracted edge $e \in G$ containing the possible lifts of C'.

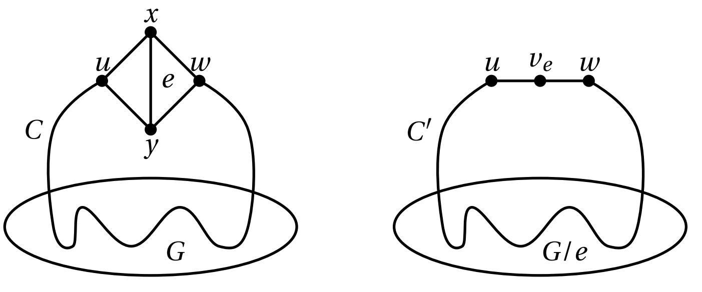

In this case there are 2 potential lifts of the cycle c'– the one that goes through x, and the one that passes through y. We will call these c_x and c_y respectively. Note that if one of c_x or c_y is basic, then the other one is good as they differ by fundamental triangles.

Suppose for contradiction that neither of these are basic. Then upon the removal of c_x or c_y, the vertices x and y are isolated. Therefore the only neighbors of x and y are in the set $\{x, y, u, v\}$. The removal of u and w would separate x, y from G. But G is supposed to be 3-connected, which contradicts our hypothesis.

We now have all the pieces to complete our proof. Start with any cycle $c \in \mathscr{C}(G)$. Because i is an injective map on the cycle space,

$$c - i \circ \pi(c) \in \ker(\pi),$$

and since $\ker(\pi)$ is generated by fundamental triangles, which are basic, we learn that c is basic whenever $i \circ \pi(c)$ is basic. By the claim, this is a basic cycle, proving the theorem.

Exercises P

P1 Exercise

Let G be a graph. We call a vertex *even* if its degree is even, and odd otherwise. Prove that there are an even number of odd vertices.

P2 Exercise

Bridges of Konigsberg

A *Walk* is a sequence of vertices which are pairwise connected by edges, and we are allowed to possibly to repeat vertices. Let G be a connected graph. Show that there is a walk in G that uses every edge exactly 1 time if and only there are at most 2 odd vertices.

P3 Exercise

A connected graph which contains no cycles is called a *tree*. Prove that the following are equivalent:

- G is a tree.
- G is minimally connected, that is $G \setminus e$ is not connected for any removed edge e
- G has no cycles, but adding an edge between any two vertices which do not have an edge already induces a cycle.
- Between every two vertices there exists a unique path.
- G is connected and has $|V| = |E| + 1$.

P4 Exercise

Show that the edge connectivity is necessarily larger than the vertex connectivity.

Exercise (P5)

Show that for every graph G, there exists a graph H so that G is a minor of H and for all $v \in V(H)$, $\deg(v) \leq 3$.

Exercise (P6)

A graph is called *cubic* if every vertex has degree exactly 3. Show that every graph G is a minor of some cubic graph H.

Exercise (P7)

Suppose for all vertices $v \in V(G)$, $\deg v \leq 3$. Show that the following are equivalent:

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

Exercise (P8)

Show that the reliability polynomial of a tree only depends on the number of edges in the tree.

Exercise (P9)

Produce two non-isomorphic graphs G and G' which are not trees and

$$R_G(p) = R_{G'}(p).$$

Exercise (P10)

Show that either $R_G(p) = 0$, or $R_G(p)$ has no zeros in the interval $(0, 1)$.

Show that a graph G is 3-connected if and only if it contains a K_4 minor.

P11 Exercise

Despite the cycle number being defined via the cycle space, we can compute it without using all of this algebra.

- Let e be an edge so that $G \setminus e$ is still connected. Show that $b_1(G)-1 = b_1(G \setminus e)$.
- Suppose that a graph G is connected. Prove that $b_1(G) = |E| - |V| + 1$.
- Generalize the above formula to when $b_0(G) \neq 1$.

P12 Exercise

Recall that $b_0(G)$ is the number of connected components of a graph G. Show that $b_0(G) = \dim \mathcal{V} - \dim \text{ Im } (d)$.

P13 Exercise

Compute $b_1(K_n)$, the dimension of the cycle space of the complete graph on n-vertices. What about $b_1(K_{m,n})$?

P14 Exercise

Embedded Graphs

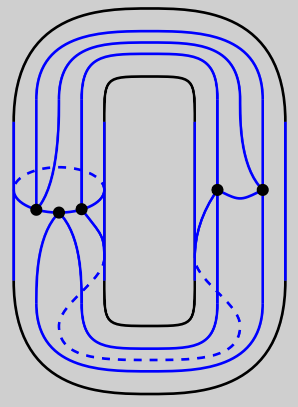

K_5 *drawn on a torus*

A map of the world gives an example of an embedded graph, where the vertices are the countries and edges record adjacencies. We explore the properties of graphs which can be drawn a variety of surfaces.

Definitions

Theorems and Lemmas

Examples

Planarity

A planar graph is a graph that can be drawn in the plane without edges crossings. We look at relations between the number of faces, edges and vertices of a graph. We also associate to every graph a dual graph. We touch on how the faces relate to the cycle space construction from before, and show that certain graphs are not planar.

While a graph is an abstract object, in practice we work with drawings of graphs for our intuition. A planar graph is a graph we can draw without the edges crossing. This means that all of the data of the graph can be easily read off of the picture. We probably won't ever use the actual definition of a planar graph, but it's good at least to write it down once:

1 Definition

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 $f_e(t) = f_{e'}(t')$ and $t, t' \neq 0, 1$, then $e' = e$ and $t = t'$. This means that the paths corresponding to edges do not cross over the interiors.
- If uv is an edge, then $f_{uv}(0) = u$ and $f_{uv}(1) = v$. This means that the paths corresponding to each pair of vertices have endpoints on the appropriate vertices.

We say that a graph is *planar* if it admits a planar representation. Note that it is possible for a graph to have topologically non-isomorphic planar representations.

2 Example

Even a graph as simple as a tree can admit different planar representations. The trees on the right are the same as graphs, but there is no way to distort the first diagram into the second without creating a crossing at some point. To capture the planar information of a graph combinatorially, one can use the data of a rotation system.

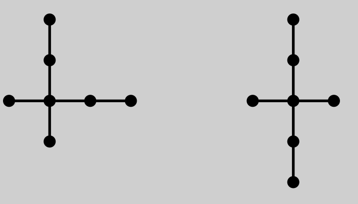

A planar representation of a graph gives us a new piece of data: the *faces* of a graph.

Definition 3

A *Face* of a planar representation of G is a connected component of $\mathbb{R}^2 \setminus G$. The set of such components is denoted $F(G)$.

Notice that this definition gives us a large "outer face" to the graph. To each face, we get a *closed walk* which makes up the boundary of the face. Even though a closed walk is not a cycle, it is an element of the cycle space.

Definition 4

Let G be a planar graph with a fixed embedding, and let $F = \{f_1, \dots, f_k\}$ be the set of faces. Let $\mathscr{F}$ be the $\mathbb{Z}_2$ vector space generated on the set F. Let $\partial_F : \mathscr{F} \to \mathscr{E}$ be the map which sends each face to the subset of edges in it's boundary (counted with multiplicity.)

This is a construction very similar to those that we've employed for our algebraic analysis of graphs, just we've included faces into it.

Claim 5

Consider the following sequence of vector spaces and functions:

$$\mathscr{F} \xrightarrow{\partial_F} \mathscr{E} \xrightarrow{\partial_E} \mathscr{V}$$

This is a *chain complex*, in that $\partial_E \circ \partial_F = 0$

Proof: In order to show that $\partial_E \circ \partial_F = 0$, it suffices to check on a basis of $\mathscr{F}$. Let $f \in \mathscr{F}$ be a face. Then $\partial_F(f)$ is a union of cycles. By **??**, this lies in the kernel of ∂_E. ⊝

We will develop the algebraic theory of chain complexes throughout the course– for more details, see Appendix **??**. We'll bring this complex back throughout this section. Let's look at some basic facts about planar graphs.

A useful geometric construction for planar graphs is the *dual graph* construction. Given a graph G with a planar embeddings, we can construct a multigraph which is dual to it.

Definition 6
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.

7 Example

For example, here is the multigraph given by the matrix

$$\begin{pmatrix} 0 & 2 & 0 & 0 \\ 2 & 0 & 1 & 3 \\ 0 & 1 & 2 & 1 \\ 0 & 3 & 1 & 0 \end{pmatrix}$$

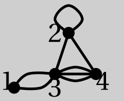

Let G be a planar multigraph. From here we can construct a new planar multigraph, called the *dual* of G, interchanging the roles of vertices and edges.

8 Definition

Let G be a planar multigraph. Let G^* be the graph where

$$V(G^*) = F(G), E(G^*) = E(G), F(G^*) = V(G)$$

We have an edge between two vertices between f_i and f_j, for each common edge in the boundary of the faces f_i and f_j.

9 Example

An example of a graph G in black and its corresponding dual multigraph G^* in red. Notice that the degree of each vertex counts the number of edges in the boundary walk of the graph in the dual, and vice versa. Also notice that the double dual, $(G^*)^*$, is G.

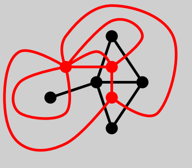

One can similarly define the boundary maps and face maps to the theory of multigraphs, either by taking the counts of edges $\mod 2$ or by using *directed multigraphs.*

Duals and Homology

In both cases, we see that the theory of planar graphs comes with two equally good chain complexes. On one hand, we can look at $\mathscr{F}(G) \to \mathscr{E}(G) \to \mathscr{V}(G)$. On the other hand, we might have instead of chosen to study $\mathscr{F}(G^*) \to \mathscr{E}(G^*) \to \mathscr{V}(G^*)$. By identifying $\mathscr{F}(G^*) = \mathscr{V}(G)$, we get a different set of maps:

$$\mathscr{F} \xleftarrow[d_E]{} \mathscr{E} \xleftarrow[d_V]{} \mathscr{V}$$

which is a kind of "dual complex" to our original theory.

10 Claim

If G is a simple graph, then every vertex of G^* has degree at least 3.

Proof: This is because each vertex of G^* corresponds to a face in G, and the boundary of such faces are cycles of G. Since G is simple, each cycle must have length at least 3. This means that each face borders at least 3 other faces, so the degree of each $f_i \in V(G^*)$ must be at least 3.

Claim 11
Basic Graph Properties

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

– $2|E| \geq |V|$

– $2|E| \geq 3|F|$

Proof: The first claim follows from our argument about average degree and edges in Claim **??**. Since G is connected, the average degree of a vertex is at least 1. The second claim is actually the same as the first claim, just applied to the dual graph. As every vertex in G^* has degree at least 3, we have that

$$2E(G) = 2E(G^*) = (\text{ Average degree in } G^* \,)|V(G^*)| \geq 3|V(G^*)| = 3|F(G)|.$$

The structure of a graph gives us an additional relation between the number of vertices, edges and faces.

Theorem 12
Euler's Formula

Let G be a connected planar graph. Then

$$|V| - |E| + |F| = 2.$$

Proof: This theorem is traditionally prove by induction on the number of edges. Of course, we cannot induct on the number of edges right away, because if we remove an edge from a graph it may not remain connected. However, with every connected graph G, there exists an edge so that $G \setminus e$ is connected, or G is a tree (See Exercise **??**.) Therefore, checking trees suffices as a base case.
In the case of a tree, we know that the number of edges is 1 fewer than the number of vertices, and the number of faces is 1. This gives us

$$|V| - |E| + |F| = |V| - (|V| - 1) + 1 = 2.$$

For the induction step, let G be a planar connected graph, and let e be an edge so that $G \setminus e$ is still connected. Since $G \setminus e$ is still connected, it is the case that e

belongs to 2 different faces. Therefore, $E(G \setminus e) = E(G) - 1$, and $F(G \setminus e) = F(G) - 1$. We conclude that

$$\begin{aligned} |V(G)| - |E(G)| + |F(G)| =& |V(G \setminus e)| - (|E(G \setminus e)| + 1) + (|F(G \setminus e)| + 1) \\ =& |V(G \setminus e)| - |E(G \setminus e)| + |F(G \setminus e)| \\ =& 2. \end{aligned}$$

A nice way to prove this is to use the algebraic graph techniques that we've been working on. See Exercise **??** for details on how one would prove this.

Let G be a simple planar graph. If $|V| \geq 3$, then $|E| \leq 3|V| - 6$.

13 Corollary

Proof: Take a planar simple graph G with at least 3 vertices. Suppose that it is not the case that every face of G is a triangle. Then you may 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*.
Let G' be a maximally planar graph containing G. In this situation we have the additional equality $2|E(G')| = 3|F(G')|$. From Euler's formula, we have

$$|V(G')| - |E(G')| + 2/3|E(G')| = 2$$

Since G' has more edges than G, we can rearrange this equality to $|E| \leq 3|V| - 6$.

These simple inequalities already rule out the existence of a planar embedding for many graphs.

For every $n \geq 5$, the complete graph K_n admits no planar embedding.

14 Corollary

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.

This gives us actually a general criterion we can check to see if a graph is non-planar, as if even G contains K_5 as a topological minor, it cannot be planar.

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

15 Corollary

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 stereographic projection. Therefore, to each platonic solid we should look at a graph which has vertices of degree m, and faces n boundary edges.
Since we know the degree and size of each edge, we can state the exact relations:

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

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

Therefore, $|F| = m/n|V|$. Applying this to Euler's formula tells us:

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

Now, we have some additional geometric bounds we may place on m and n.

- We know that n is at least 3 (all faces have at least 3 sides.)
- We get the bound
$$m\left(\frac{1}{n} - \frac{1}{2}\right) > -1$$
from the Euler characteristic formula. This means that m cannot be greater than 5. Another interpretation of this is that you cannot pack more than 5 regular polygons around a point and have the angles at the vertex sum to less than 2π.
- m must be at least 3. The only valid values of m are now 3, 4, 5.
- By taking a dual polygon, we get similar restraints on n.

Tabulating our results we have:

m	n	$\|V\|$	$\|E\|$	$\|F\|$	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. These dual-polytopes pairs are given by dual-graphs.

Outside of the Platonic solids, there are many different types of polytopes that one can study and represent with graphs. For instance, a *quasi-regular* polyhedra is allowed to have 2 different kinds of faces that alternate around each vertex, and their classification follows a similar argument as the one used above. The general theory of understanding convex polytopes branches substantially into algebraic topology. For instance, understanding the simplicial convex d-polytopes can be understood in just the number of facets it has in every dimension([Sta80]).

2 Graph Colorings

We return to a historic problem: how many colors are needed to color a map. We first explore abstract colorings of maps, including coloring estimates, chromatic polynomials, and properties of 2-colorable graphs. We then later turn to coloring planar graphs, proving the 6 and 5 color theorem.

We're going to take a slight detour from planar graphs to talk about graph colorings, which are a major tool in graph theory. Eventually, we'll bring this back to planar graphs when we discuss colorings of planar graphs.

Definition 16 *Colorings*

Let G be a graph. A *k-coloring* of a graph 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 $\gamma(G)$.

Example 17

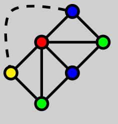

The graph on the right, despite not containing a complete graph on 4 vertices, still requires a minimal of 4 colors to color. The removal of the dashed edge would lower the chromatic number to 3.

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 18

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 19

Let G be a graph. Then

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

Proof: Create a multigraph K_G 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 (see Figure **??** for an example of these graphs.) Since this is a complete graph on $\gamma(k)$ vertices, G must have at least $\frac{1}{2}k(k-1)$ edges.

∂^2 Homological Algebra

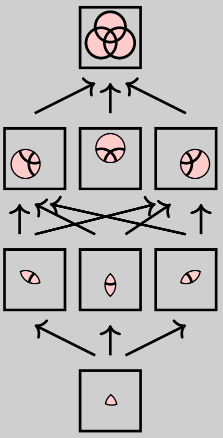

Inclusion Exclusion

These notes are meant as a quick reference guide for the constructions in homological algebra that we will use throughout the course, and are not in any way suppose to be a substitute for a proper set of notes on homological algebra.

Definitions

Theorems and Lemmas

Examples

Vector Spaces, Sets, and Diagrams

Before we start with the development of homological algebra, it is a good idea to set up some common conventions and diagrams for simplifying linear algebra.

1

These are some class notes! Please lets me if you know see any errors. Here we will flesh these methods out in more detail before developing chain complexes.

1 Definition *Direct Sum*

Let V_1 and V_2 be vector spaces. The *direct sum* of V_1 and V_2 is the vector space of pairs of vectors, and is denoted

$$V_1 \oplus V_2 := \{(v_1, v_2) \mid v_1 \in V_1, v_2 \in V_2\}.$$

The vector addition on $V_1 \oplus V_2$ is done component wise,

$$(v_1, v_2) + (w_1, w_2) = (v_1 + w_1, v_2 + w_2).$$

The scalar multiplication acts on all components simultaneously,

$$\lambda \cdot (v_1, v_2) = (\lambda \cdot v_1, \lambda \cdot v_2).$$

2 Example *Real n dimensional space*

The set of n-tuples of real numbers is usually denoted $\mathbb{R}^n$. Another way of presenting this vector space is

$$\mathbb{R}^n = \underbrace{\mathbb{R} \oplus \mathbb{R} \oplus \cdots \oplus \mathbb{R} \oplus \mathbb{R}}_{n}$$

where now each "vector" $r_i \in R^1_i$ is a scalar.

The direct sum operation is commutative, in that the vector spaces $V_1 \oplus V_2$ is isomorphic to $V_2 \oplus V_1$. Additionally, the direct sum of vector spaces is an associative operation so that the vector spaces $(V_1 \oplus V_2) \oplus V_3$ is isomorphic to $V_1 \oplus (V_2 \oplus V_3)$. If this looks suspiciously like addition on the integers to you, you're picking up on an intertwining between these two operations via dimension:

$$\dim(V_1 \oplus V_2) = \dim(V_1) + \dim(V_2.)$$

Example 3

The rank nullity theorem can be restated as: If $f : V \to W$ is a linear map, then

$$V \simeq \ker f \oplus \operatorname{Im} f.$$

Given vector spaces V_1, V_2, W, and maps $f_1 : V_1 \to W$, $f_2 : V_2 \to W$, one can create a new map from $V_1 \oplus V_2 \to W$, which is defined by taking the sum of the two maps:

$$\begin{aligned} f_1 \oplus f_2 : V_1 \oplus V_2 &\to W \\ (v_1, v_2) &\mapsto f_1(v_1) + f_2(v_2). \end{aligned}$$

We will frequently represent this composition either *diagrammatically* or using matrices. This is a useful shorthand, and we will use it throughout this section on chain complexes.

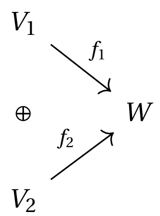

$$\begin{pmatrix} f_1 & f_2 \end{pmatrix} \cdot \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} f_1(v_1) + f_2(v_2) \end{pmatrix}.$$

There is nothing that limits us to taking the direct sum of more than one map along the domain.

Definition 4 *Sum across common target*

Let $f_i : V_i \to W$ be a collection of maps. Then define $\bigoplus_{i=1}^k f_i : \bigoplus_{i=1}^k V_i \to W$ be the map defined on tuples by

$$\left(\bigoplus f_i\right)(v_1, \ldots, v_k) = \sum_{i=1}^k f_i(v_i).$$

Just as we can take the sum along the domains of maps, we are also allowed to take sums along the targets of the maps. Let $g_1 : V \to W_1$ and $g_2 : V \to W_2$ be two linear maps. Then denote the direct sum along the target

$$\begin{aligned} g_1 \oplus g_2 : V &\to W_1 \oplus W_2 \\ v &\mapsto (g_1(v), g_2(v).) \end{aligned}$$

Just as we did for direct sum along the domain, we can represent these maps diagrammatically or with matrices.

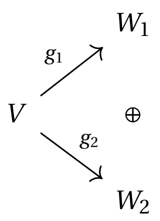

$$\begin{pmatrix} g_1 \\ g_2 \end{pmatrix} \cdot \begin{pmatrix} v \end{pmatrix} = \begin{pmatrix} g_1(v_1) \\ g_2(v_2) \end{pmatrix}.$$

We can quickly do this with many targets at the same time.

5 Definition

Sum across common domain

Let $g_i : V \to W_i$ be a collection of linear maps. Define their direct sum to be

$$\bigoplus_{i=1}^{k} g_i : V \to \bigoplus_{i=1}^{k} W_i$$
$$v \mapsto (g_1(v), g_2(v) \cdots, g_k(v)).$$

By combining both processes, we can create maps from many domains and targets simultaneously.

6 Definition

Sum across domains and targets

Let $f_{ij} : V_i \to W_j$ be a collection of linear maps. Define their direct sum to be

$$\bigoplus_{i,j} f_{ij} : \bigoplus_{i=1}^{m} V_i \to \bigoplus_{j=1}^{n} W_j$$
$$(v_1, \ldots, v_m) \mapsto \left(\sum_{i=1}^{m} f_{i,1}(v_i), \sum_{i=1}^{m} f_{i,2}(v_i), \ldots, \sum_{i=1}^{m} f_{i,n-1}(v_i), \sum_{i=1}^{m} f_{i,n}(v_i) \right).$$

We again have diagrammatic and matrix notations for these maps.

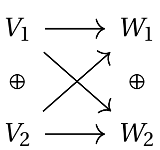

$$\begin{pmatrix} f_{11} & f_{21} \\ f_{12} & f_{22} \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} f_{11}(v_1) + f_{21}(v_2) \\ f_{12}(v_1) + f_{22}(v_2) \end{pmatrix}.$$

7 Inclusion- Exclusion with 2 sets

Suppose that we have a decomposition $A = S_1 \cup S_2$. Then the sizes of these sets are related by the *inclusion-exclusion* formula: $0 = |A| - ((|S_1| + |S_2|)) + |S_1 \cap S_2|$.

We will first translate the sets A, S_1, S_2 and $S_1 \cap S_2$ into vector spaces. We take a slightly different approach than before. To each set U, let $\mathscr{F}(U) := \hom(U, \mathbb{Z}_2)$. Note that $\mathscr{F}(U) \cong \mathscr{U}$, the $\mathbb{Z}_2$ vector space whose basis is given by U, but not canonically isomorphic. The advantage with working with the vector space $\mathscr{F}(U)$ is that it is canonically defined (i.e. doesn't come with a preferred basis.) Each element $\phi \in \mathscr{F}(U)$ can be thought of as an assignments of 0's and 1's to the elements of U.

A slightly confusing feature of working with this vector space is that functions between sets translate into functions going the other direction on the vector spaces,

$$f : U \to V$$
$$\mathscr{F}(U) \leftarrow \mathscr{F}(V) : f^*.$$

The map f^* is called the *pullback* map, and it is defined via precomposition. Given an element $\phi \in \mathscr{F}(V)$, the pullback along f is the map $(\phi \circ f) \in \mathscr{F}(U)$. I find the clearest way to think about this is interpret $\mathscr{F}(V)$ as the space of measurements on V. Then a function $f : U \to V$ yields for each measurement $\phi : V \to \mathbb{Z}_2$ a new measurement $f^*(\phi)$ on the space U. The way this measurement $f^*(\phi)$ works is by taking elements $u \in U$, sending them to V, and then performing the measurement ϕ there:

$$f^*(\phi)(u) := \phi(f(u)).$$

Remark.The function $\mathscr{F} : \mathbf{Sets} \to \mathbf{Vect}$ turns problems about sets into problems of vector spaces. This function is an example of a *functor.* Because $\mathscr{F}$ reverses the directions of functions, we call this a *contravariant functor.* The general theory of functors belongs to a branch of mathematics called category theory, which studies mathematics from the perspective of general properties of functions.

An important feature of the functor $\mathscr{F}$ is that it exchanges cardinality with dimension:

$$|U| = \dim(\mathscr{F}(U)).$$

Let's return to the setting of inclusion-exclusion. Suppose that we have a decomposition $A = S_1 \cup S_2$. We can encode this decomposition in the following maps between sets:

$$\begin{array}{ccc} S_1 \cap S_2 & \xrightarrow{i_1} & S_1 \\ \downarrow{\scriptstyle i_2} & & \downarrow{\scriptstyle j_1} \\ S_2 & \xrightarrow{j_2} & A \end{array} \qquad \begin{array}{ccc} \mathscr{F}(S_1 \cap S_2) & \xleftarrow[i_1^*]{} & \mathscr{F}(S_1) \\ {\scriptstyle i_2^*}\uparrow & & {\scriptstyle j_1^*}\uparrow \\ \mathscr{F}(S_2) & \xleftarrow[j_2^*]{} & F(A) \end{array}.$$

Theorem.Let $A^0 = \mathscr{F}(A)$, $A^1 = (\mathscr{F}(S_1) \oplus \mathscr{F}(S_2))$ and $A^2 = \mathscr{F}(S_1 \cap S_2)$. Let $i^* := i_1^* \oplus i_2^* : A^1 \to A^2$, and let $j^* := j_1^* \oplus j_2^* : A^0 \to A^1$ as drawn below:

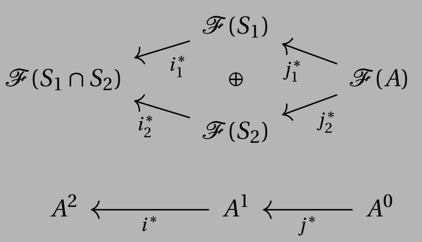

The map j^* is an inclusion, the map i^* is surjective, and $\ker(i^*) = \operatorname{Im}(j^*)$.

Proof: We show that the map j^* is an inclusion. Let $\phi \in \mathscr{F}(A)$ be a non-zero element, and let $a \in A$ be the element so that $\phi(a) = 1$. Since $A = S_1 \cup S_2$, there is an element $b \in S_1$ or $b \in S_2$ so that $j_1(b) = a$ or $j_2(b) = a$. Without loss of generality, suppose $b \in S_1$. We can then compute that $j^*(\phi) = (\phi \circ j_1, \phi \circ j_2)$ and $\phi \circ j_1(b) \neq 0$. This proves that $j^*(\phi)$ is nonzero, so the map j^* has trivial kernel and is therefore injective. A similar proof shows that i^* is surjective.

We now show that $\ker(i^*) = \operatorname{Im}(j^*)$. For any element $a \in S_1 \cap S_2$, we note that

$$(i^* \circ j^*(\phi))(a) = \phi((j_1 \circ i_1)(a)) + \phi((j_2 \circ i_2)(a))$$

Since $(j_1 \circ i_1)(a) = (j_2 \circ i_2)(a)$,

$$= 2\phi(j_1 \circ i_1(a)) = 0$$

This shows that $\operatorname{Im}(j^*) \subset \ker(i^*)$. The reverse inclusion is by a similar argument.

We can now prove Inclusion-Exclusion for two sets. We will instead show that $\dim A^0 - \dim A^1 + \dim A^2 = 0$ using two applications of the rank-nullity theorem.

$$\dim A^0 - \dim A^1 + \dim A^2 = (\dim\ker(j^*) + \dim \operatorname{Im}(j^*)) - (\dim\ker(i^*) + \dim \operatorname{Im}(i^*)) + \dim A^2$$

As the map j^* is injective and i^* is surjective

$$\begin{aligned} &= (0 + \dim \operatorname{Im}(j^*)) - (\dim\ker(i^*) + \dim \operatorname{Im}(i^*)) + \dim \operatorname{Im}(i^*)) \\ &= \dim \operatorname{Im}(j^*) - \dim\ker(i^*) \\ &= 0. \end{aligned}$$

∎

2 Connected Components, and Toplogy

In this section we introduce some basic notions from topology which will motivate some of our future discussions.

It's beyond the scope of this course to define what a topological space is, and the functions between those topological spaces. The main framework that we'll need is to know the following facts about topological spaces.

- Topological spaces are sets with some additional structure (called a topology.)
- There are certain functions between these sets, called *continuous functions*, which preserve the useful properties of the topology.
- The composition of continuous functions is again continuous.
- If X is a topological space, and $\mathbb{Z}_2$ is the topological space with two points, then the set of continuous functions $C^0(X,\mathbb{Z}_2)$ is the a vector space. Furthermore, $\dim(C^0(X,\mathbb{Z}_2))$ is the number of connected components of X.

These are the only properties of topological spaces which we will need to continue this discussion.

Example 8

The basic example of a topological space is $X := \mathbb{R}$. The functions from $f : \mathbb{R} \to \mathbb{R}$ which are continuous are exactly the continuous functions you know and love, satisfying the property

$$\lim_{x_i \to x} f(x_i) = f(x).$$

This property is fondly phrased as "when you draw the graph of $f(x)$, there are no jumps in the graph. "

Some more interesting examples of topological spaces are things like circles, tori, disks, spheres, graphs.

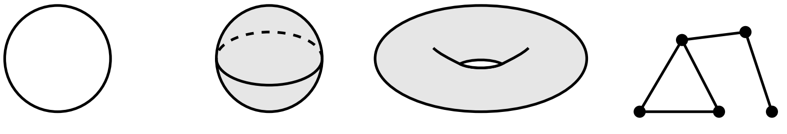

Our intuition for continuous maps is that they are the functions between topological spaces which send nearby points to nearby points. We give a very brief overview of some concepts from topology in Example 10 .

We define the *connected component space of X* to be the vector space

$$C^0(X) := \hom(X, \mathbb{Z}_2)$$

of continuous functions from X to the two point set. One can think of this as assigning a color to each connected component of the space X, and the number of colorings (determined by the dimension $\dim C^0(X)$) tells you how many connected components there are.

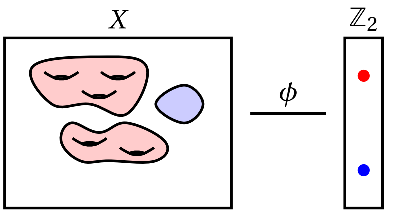

Claim 9 — *Pullback Map*

Given a continuous $f : X \to Y$ between topological spaces, there is a map

$$f^* : C^0(Y) \to C^0(X).$$

Proof: The pullback function is defined as before:

$$\begin{aligned} f^* : C^0(Y) &\to C^0(X) \\ \phi &\mapsto (\phi \circ f) \end{aligned}$$

The only thing to check is that $\phi \circ f$ is a continuous map from $X \to \mathbb{Z}_2$; this follows from the composition of continuous maps being continuous. ∂^2

What this claim means is that we can track how the connected components of X are mapped to connected components of Y by using the pullback map. One interpretation of this is that given a map $f : X \to Y$, we can "color" the connected components of X by the connected components of Y.

This framework should look very familiar– it is the same set-up that we used to describe the number of elements in sets. The connected component space $C^0(X)$ turns questions about connected components into problems in linear algebra instead. Let us take the annulus, and decompose it into two sets as drawn below. This configuration does not respect an inclusion-exclusion like property in the usual sense, in that U_1, U_2, X each have one connected component, but $U_1 \cap U_2$ has two connected components.

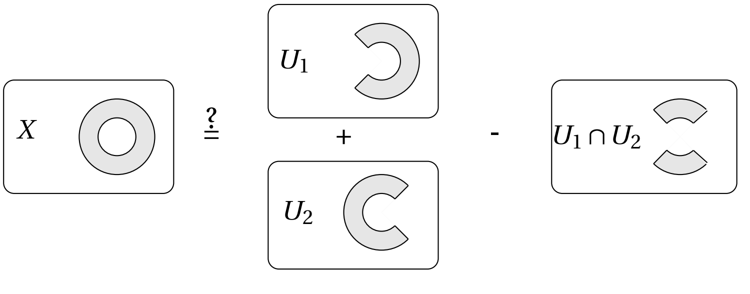

Topology 101

A topological space is a set, equipped with the additional data of *open sets* which determine which points on the topological space are close to each other. In this section, we give a quick overview of point-set topology.

Definition.A *topological space* is a pair $(X, \mathscr{U})$, where X is a set, and $\mathscr{U}$ is a specified collection of subsets of X, called *open sets* satisfying the following axioms:

– The empty set and whole space X are open sets.

$$\emptyset, X \in \mathscr{U}$$

– Any union of open sets is an open set.

$$U_\alpha \subset \mathscr{U} \Rightarrow \left(\bigcup_{\alpha \in A} U_\alpha\right) \in \mathscr{U}.$$

– Any finite intersection of open sets is an open subset.

$$\mathscr{B} \subset \mathscr{U}, |B| < \infty \Rightarrow \left(\bigcap_{\beta \in B} U_\beta\right) \in \mathscr{U}.$$

Open sets are kind of strange things. Roughly speaking, if x and y mutually belong to an open set, then we know that they are close to each other in *some* sense, but unlike in the metric space a topology doesn't tell you *how* near two points are two each other. It just tells you that there is something containing both of them. We still get some relative idea of closeness– if two points mutually belong to many open sets, then we think of them being closer to each other.
Let's introduce a few examples of topologies.

Example (*The Discrete Topology*). Let X be a set. The *discrete topology* has every subset of X as an open set:

$$\mathscr{U} = \{U \mid U \subset X\}$$

This topology has too many open subsets, and all of the points are very far away from each other!

A common example of a topological space comes from metric spaces. We'll say that a U is open if every point in x is contained within an open ball inside of U.

Example.Let (X, ρ) be a metric space. Say that a set U is *ρ-open* if for every point $x \in U$, there exists an open ball $B_\epsilon(y)$ with

$$x \in B_\epsilon(y) \subseteq U.$$

Then the collection of sets

$$\mathscr{U} = \{U \subset X \mid U \text{ is } \rho\text{-open}\}$$

makes $(X, \mathscr{U})$ a topology. For example, on the real numbers every open interval is an example of an open set with this topology.

The interesting maps between topological spaces are those which preserve the topological structure.

Definition (*Continuous Maps*). Let $f : X \to Y$ be a function, and $U \subset Y$. The *pre-image* of Y is all the elements of X which get mapped to U,

$$f^{-1}(U) := \{x \in X \mid f(x) \in U\}.$$

A function $f : X \to Y$ is continuous if and only if for every open set $U \subset Y$, the preimage

$$f^{-1}(U) \subset X$$

is an open set of X.

Suppose that $f : X \to Y$ and $g : Y \to Z$ are continuous maps. Then for any $U \in Z$, $(g \circ f)^{-1}(U)$ is again an open set, which shows that the composition of continuous maps is continuous.

A topological space is called *disconnected* if $X = U_1 \sqcup U_2$, with U_1, U_2 nonempty open sets. The *connected components* of a topological space are the smallest nonempty open sets $\{U_i\}$ so that $X = \bigsqcup_{i=1}^k U_i$. We say that in this case that X has k-connected components.

Theorem.Suppose that X has k-connected components. Let $\hom(X, \mathbb{Z}_2)$ denote the set of linear maps from X to the space with two points. Then

$$\dim(\hom(X, \mathbb{Z}_2)) = k.$$

Let's see exactly how the argument from that worked in the proof that $|X| - (|U_1| + |U_2|) + (|U_1 \cap U_2|) = 0$ fails when we now try to understand the number of connected components. The spaces U_1, U_2, X all have one connected component, so

$$C^0(X) = C^0(U_1) = C^0(U_2) = \mathbb{Z}_2.$$

On the other hand, $U_1 \cap U_2$ has two connected components, so $C^0(U_1 \cap U_2) = \mathbb{Z}_2 \oplus \mathbb{Z}_2$. We now look at the inclusions of topological spaces

$$\begin{array}{ccc} U_1 \cap U_2 & \xrightarrow{i_1} & U_1 \\ \downarrow{\scriptstyle i_2} & & \downarrow{\scriptstyle j_1} \\ U_2 & \xrightarrow{j_2} & X \end{array} \qquad \begin{array}{ccc} C^0(U_1 \cap U_2) & \xleftarrow[i_1^*]{} & C^0(U_1) \\ {\scriptstyle i_2^*}\uparrow & & {\scriptstyle j_1^*}\uparrow \\ C^0(U_2) & \xleftarrow[j_2^*]{} & C^0(X) \end{array} \qquad \begin{array}{ccc} \mathbb{Z}_2 \oplus \mathbb{Z}_2 & \xleftarrow[i_1^*]{} & \mathbb{Z}_2 \\ {\scriptstyle i_2^*}\uparrow & & {\scriptstyle j_1^*}\uparrow \\ \mathbb{Z}_2 & \xleftarrow[j_2^*]{} & \mathbb{Z}_2 \end{array}.$$

We then condense this down into a sequence of vector spaces by defining $C^1(X) := C^0(U_1) \oplus C^0(U_2)$, and $C^2(X) := C^0(U_1 \cap U_2)$. Similarly, we define the maps

$$j^* := j_1^* \oplus j_2^* : C^0(X) \to C^1(X)$$
$$i^* := i_1^* \oplus i_2^* : C^1(X) \to C^2(X).$$

as before to give us a sequence of vector spaces and maps between them.

$$C^0(X) \xrightarrow{j^*} C^1(X) \xrightarrow{i^*} C^2(X)$$

This entire set-up so far follows the same steps as the inclusion-exclusion set up for sets. At this point, we deviate from that example.

Claim 11

For the maps and sets above, the map j^* is injective and $\operatorname{Im}(j^*) \subset \ker(i^*)$.

Proof: Let $\phi : X \to \mathbb{Z}_2$ be any continuous function. Then $j^*(\phi)$ is $(j_1)^*\phi \oplus (j_2)^*\phi$, where $(j_1)^*\phi : U_1 \to \mathbb{Z}_2$ and $(j_2)^*\phi : U_2 \to \mathbb{Z}_2$ are the restriction of ϕ to the subsets U_1, U_2. Then

$$(i^* \circ j^*)\phi = (i_1^* \circ j_1^*)\phi + (i_2^* \circ j_2^*)\phi$$

Since $i_1^* j_1^* = i_2^* j_2^*$,

$$= 2(i_1^* \circ j_1^*)\phi = 0.$$

This proves that $i^* \circ j^* = 0$, which is equivalent to $\operatorname{Im}(j^*) \subset \ker(i^*)$.

This claim is weaker than the statement that we had for the complex involving sizes of sets. That claim stated that $\operatorname{Im}(j^*) = \ker(j^*)$, instead of only having an inclusion, and that i^* was a surjection. The discrepancy between these two statements – equality of image and kernel versus inclusion of image into kernel – gives us an exact measurement of how the inclusion exclusion principle fails.

Chain Complexes, Homology, and Chain Maps

Homological Algebra is a algebraic tool that we'll return to at several points throughout the course, and it makes sense to combine the general facts of the theory in one place.

12 Definition
Cochain complexes

A *cochain complex* is a sequence of vector spaces, $\dots C^{-1}, C^0, C^1 \dots$ and boundary maps $d^n : C^n \to C^{n+1}$ with the condition that

$$d^{n+1} \circ d^n = 0.$$

Frequently, we represent a chain complex with the following diagram of vector spaces and maps:

$$\cdots \xleftarrow[d^1]{} C^1 \xleftarrow[d^0]{} C^0 \xleftarrow[d^{-1}]{} C^{-1} \xleftarrow[d^{-2}]{} \cdots$$

We will usually denote the chain complex as $(C^\bullet, d^\bullet)$, where $C^\bullet$ is the sequence of modules and $d^\bullet$ the sequence of boundary maps. [1]

Abelian Categories

In principle, all of the tools that we are developing with cochain complexes can be defined with rings and modules instead of just vector spaces. In fact, the field of homological algebra generally works over any *Abelian category*, which is a category equipped with the necessary structures to make linear algebra-like constructions.

13 Example

Let's look at a first example of a chain complex. Let $C^1 = C^2 = C^3 = \mathbb{R}^2$, so that we may represent our boundary maps by matrices. Consider the sequence of maps

$$0 \xrightarrow{0} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{0} 0$$

This is an example of a chain complex, as the composition of the differential is zero:

$$d^3 \circ d^2 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

The boundary squaring to zero is equivalent to the statement that the image of the boundary map d^k is in the kernel of the map d^{k+1}.

The theory of cohomology was developed and inspired from techniques in topology, but it is a very useful algebraic framework to have in mind. Abstractly, the chain complexes and cohomology are a tool that explains the relations, and relations of relations, and higher meta-relations. For example, let V be a set with a relation $E \subset V \times V$ on it. Let $\mathscr{F}(V)$ and $\mathscr{F}(E)$ be the vector spaces given by maps to the field of two elements. One might state the relationship now in terms of a map $d : \mathscr{F}(V) \to \mathscr{F}(E)$, where the image of a function $\phi : V \to \mathbb{Z}_2$ consists of all relations E which have a member evaluating under ϕ.

However, the framework of homology allows us to put relations on the set of relations, by introducing maps' $\mathscr{F}(E) \to \mathscr{F}(V)$, and so on.

Example 14

'' The example we considered in Theorem 46 is more than just a cochain complex; it satisfies the stronger condition of being *exact* in that $\operatorname{Im} d^k = \ker d^{k+1}$. We'll explore exact complexes in more detail in the future.

Example 15

The examples considered in Station 2 of topological spaces covered with sets, and the $\mathscr{F}(-)$ functor give another example of cochain complexes.

Before we study the general theory of cochain complexes, we would like to build a combinatorial framework for describing topological spaces, which will give us something concrete to stand on when we start describing cochain complexes in this class. The natural extension of vertices, edges and faces are building blocks called **simplices.**

Definition 16
Geometric Simplex

For $k \geq 0$, a **geometric k-simplex** α^k is the set of points in $\mathbb{R}^{k+1}$ whose coordinates are non-negative and sum to 1.

$$\{(x_1, x_2, \ldots, x_{k+1}) \mid x_1 + x_2 + \cdots x_{k+1} = 1, x_i \geq 0\}.$$

Given a simplex, we say that k is the **dimension** of α^k.

17 Example

We've already seen a couple of geometric simplices before, and given them some common names.

Dim	Name	Notes	Graphical Representation
0	Vertex	By the above definition, it specifically the point $1 \in \mathbb{R}^1$.	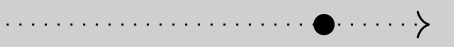
1	Edge	Drawn with the above notation, it is the line segment in the first quadrant. Notice that the restriction of the line to either axis gives us a point.	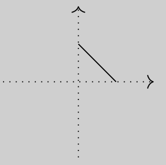
0	Face	A 2-simplex is a (filled in) triangle, filling the first quadrant. Again, the restriction to either the coordinate planes or axis gives us edges and vertices respectively.	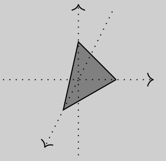

Simplices have the property that their boundaries are created of smaller simplices. For instance, a 2-simplex (triangle) has 3 boundary 1-simplices (edges.) A 3-simplex (tetrahedron) has 4 boundary 1-simplices. In general a k-simplex has $k+1$ boundary $k-1$-simplices, called *facets*.

A simplex has more than just $k-1$ dimensional facets; it also has boundary components of dimension $k-l$. Each boundary component is uniquely specified by the $k-l+1$ corner vertices it uses. If we wanted to build more complicated spaces by gluing together simplices, one would imagine that we would take these simplices and join them together along boundary strata picked out by identifying their vertices.

18 Example

A simplicial complex

Here is an example of a topological space constructed from simplices. It uses 8 vertices, has 13 edges, 8 faces, and 1 3-simplex (the right simplex is not filled in.) Notice that this topological space doesn't have a consistent notion of "dimension"– the dimension varies from 1-3 dimensional depending on which part of the complex you look at.

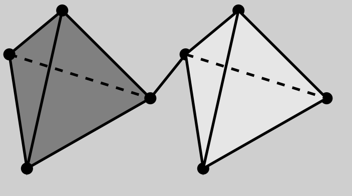

In practice, it is simpler to build in this identification of simplices from the very beginning.

Definition 19
Abstract Simplicial Complex

A *finite abstract simplicial complex* is a pair $X = (\Delta, \mathscr{S})$ where

- $\mathscr{S}$ is a base set of vertices
- $\Delta \subset \mathscr{P}(S)$ is a finite set of *simplices*

where the simplices are downward closed. This means that whenever $\sigma \in \Delta$ and $\tau \subset \sigma$, then $\tau \in \Delta$. We say that $\sigma \in \Delta$ is a k*-simplex* if $|\sigma| = k+1$. We will in this case write that $\dim(\sigma) = k$. If $\sigma \subset \tau$, and $\dim \sigma = \dim \tau - 1$, then we say that σ is a *facet* of τ and write $\sigma \lessdot \tau$.

Claim 20
Covers from Simplices

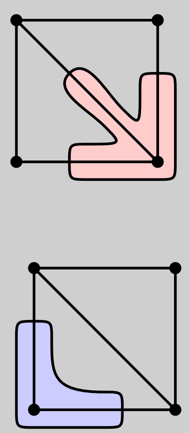

Let $X = (\Delta, \mathscr{S})$. There is a collection of sets $\{U_s\}_{s \in \mathscr{S}}$ so that $\bigcup_{s \in \mathscr{S}} U_s = X$. Define for each simplex $\sigma \in \Delta$ the associated covering set

$$U_I = X \cap \bigcap_{s \in I} U_s.$$

Furthermore, for every indexing set I, U_I is contractible, and is non-empty if and only if $I = \sigma$ for some simplex in our complex.

Note that for each $\sigma \lessdot \tau$, there exists an inclusion map $i_{\sigma\tau} : U_\tau \to U_\sigma$, and subsequently a map

$$i^*_{\sigma\tau} : \hom(U_\sigma, \mathbb{Z}_2) \to \hom(U_\tau, \mathbb{Z}_2).$$

We now define the *reduced Cech cochain complex*. For each i, let

$$\underline{C}^{-1}(X, \mathbb{Z}_2) := \hom(X, \mathbb{Z}_2)$$

$$\underline{C}^{i}(X, \mathbb{Z}_2) := \bigoplus_{\sigma \mid \dim(\sigma) = i} \hom(U_\sigma, \mathbb{Z}_2).$$

Define the differential maps

$$d^i : \underline{C}^i(X, \mathbb{Z}_2) \to C^{i+1}(X, \mathbb{Z}_2)$$
$$d^i := \bigoplus_{\sigma \lessdot \tau, \dim \sigma = i} i^*_{\sigma\tau}.$$

$\underline{C}^{\bullet}(X, \mathbb{Z}_2)$ with differential d^i is a cochain complex. Furthermore, a basis of the C^i can be indexed by the i-dimensional simplices of X, and the differential defined on a basis element e_σ can be written as

$$d(e_\sigma) = \sum_{\tau \mid \sigma \lessdot \tau} e_\tau.$$

21 Claim
Simplicial Cochains are a complex

It is rarely the case that this will be an example of an exact chain complex. The difference between $\operatorname{Im} d^{i+1}$ and $\ker d^i$ will be an interesting thing to measure. Because we are loathsome to leave the land of vector spaces, we will measure this difference with a new vector space.

Let $(C, \partial_\bullet)$ be a chain complex. The *cohomology* of $C^\bullet$ at n is defined to be the module

$$H^n(C) = \frac{\ker d^n}{\operatorname{Im} d^{n-1}}$$

As the composition $d^{n+1} \circ d^n = 0$, this is well defined.

22 Definition
Cohomology Groups

For convenience, we will often call the kernel of d^n the set of cocycles, and write it Z^n. The image of d^{n-1} is the set of coboundaries and will be written B^n. Then $H^n(C) = Z^n / B^n$. The names cycles and boundaries correspond to the geometric interpretation of the homology as given above.

We say that a chain complex is *bounded* if there exists n such that $C^i = 0$ if $|i| \geq n$.

23 Definition

While it doesn't make sense to ask about the dimension of a chain complex, there is a generalization of dimension which applies to chain complexes.

Let (C, d) be a bounded cochain complex with each C^i of finite dimension. Then the *Euler Characteristic* of (C, d) is the integer

$$\chi(C, d) := \sum_{k=-\infty}^{\infty} (-1)^k \dim(C^k).$$

24 Definition
χ of a complex

Notice that the Euler Characteristic has no dependence on the differential of a chain complex. However, it is intimately related to the chain structure through an application of the rank-nullity theorem.

Lemma 25

Euler via Homology

Suppose that the chain complex is bounded. Then

$$\chi(C,d) = \sum_{k=-\infty}^{\infty} (-1)^k \dim H^k.$$

Proof: Because our complex is bounded, there exists n such that $|k| \geq n$ implies that $C^k = H^k = 0$. Then we proceed by computing the sum:

$$\chi(C,d) = \sum_{k=-i}^{i} (-1)^k C^k$$

Applying the Rank-Nullity theorem

$$= \sum_{k=-i}^{i} (-1)^k (\dim(\ker d^k) + \dim(\operatorname{Im} d^k))$$

Shifting the sum

$$= \sum_{k=-i}^{i} (-1)^k (\dim(\ker d^k) - \sum_{k=-i}^{i} (-1)^{k-1} \dim(\operatorname{Im} d^k))$$

$$= \sum_{k=-i}^{i} (-1)^k \dim(\ker d^k) - \dim(\operatorname{Im} d^{k-1})$$

$$= \sum_{k=-i}^{i} (-1)^k \dim H^k$$

∂^2

One interpretation of homology is that it is an algebraic measure of how far a sequence strays from being *exact*.

Definition 26

Exact Sequences

A chain complex (C,d) is called *exact* if $H^k(C) = 0$ for all k.

Notice by Lemma 25, whenever (C,d) is exact, the Euler characteristic $\chi(C,d) = 0$.

Corollary 27

Inclusion-Exclusion holds for sets.

Proof: In Theorem 46 we showed that the chain complex dictating inclusion-exclusion for sets was exact. Furthermore, we showed that the inclusion-exclusion principle for sets was equivalent to $\chi(A,d) = 0$. ∂^2

Maps between chain complexes, addition and subtraction

4

Now that we have chain complexes, we want to look at functions that can go between them. Just like when we study vector spaces and groups, it is only useful to study the maps between these objects which preserve their structure. We want the function between chain complexes to be compatible with the differential.

28 Definition

Chain map

Let (A, d_A) and (B, d_B) be chain complexes, and let $f^i : A^i \to B^i$ be a collection of maps. Then we say that $f^\bullet = \{f^i\}$ is a *cochain map* if the following diagram commutes for all i:

$$\begin{array}{ccc} A^i & \xrightarrow{d_A^i} & A^{i+1} \\ \downarrow{f^i} & & \downarrow{f^{i+1}} \\ B^i & \xrightarrow{d_B^i} & B^{i+1} \end{array} .$$

A chain map not only preserves the boundary structure of the chain complex, it also gives us maps between their homology groups.

29 Claim

Induced map on cohomology

Let $f^\bullet : (A, d_A) \to (B, d_B)$ be a chain map. Then there is a well defined map between the cohomology of (A, d_A) and (B, d_B) given by

$$\begin{aligned} f^k : H^k(A) &\to H^k(B) \\ [a] &\mapsto [f^k(a)]. \end{aligned}$$

Proof: In order to show that this map is well defined, we need to check two things. First we must show that elements representing homology classes in A get sent to elements representing homology classes in B. Second, we must show that resulting map does not depend on the choice of representative for a.

- For the first part, let $[a] \in H^k(A)$ be an element of homology. In order for $[f^k(a)]$ to be an element of $H^k(B)$, we need that $f^k(a) \in \ker d_B$. We make a computation:

  $$d_B(f^k(a)) = f^{k+1}(d_A(a))$$

  Since $[a] \in H^k(A)$, we know that $a \in \ker d_A$.

  $$= f^{k+1}(0) = 0.$$

– For the second part, suppose we have 2 different representatives of the same cohomology class $[a] = [a'] \in H^k(A)$. We would like to show that $[f^k(a)] = [f^k(a')] \in H^k(B)$.
  Two classes in homology are equivalent if they differ by an element in the image of d^{k-1}. Therefore, we can prove the statement by finding an element $\beta \in B^{k-1}$ which satisfies:

$$[f^k(a)] - [f^k(a')] = d^{k-1}(\beta).$$

  We can construct this β by looking at the difference $a - a'$. Since $[a] = [a']$, there is an element $\alpha \in C^{k-1}(A)$ so that $d_A(\alpha) = a - a'$.
  We now are in the place to make a computation.

$$\begin{aligned} f^k(a) - f^k(a') =& f^k(a - a') \\ =& f^k(d_A(\alpha)) \\ =& d_B(f^{k-1}(\alpha)). \end{aligned}$$

  We set $\beta = f^{k-1}(\alpha)$ to realize the equivalence relation between the two homology classes $[f^k(a)], [f^k(a')]$.

∂^2

The most useful example of exact complexes are *short exact sequences*, which are exact complexes of the form:

$$0 \longrightarrow A \xrightarrow{i} B \xrightarrow{\pi} C \longrightarrow 0 \ .$$

From the definition of exactness $i : A \to B$ must be injective, and $\pi : B \to C$ must be surjective. If we were only interested in vector spaces, then $B = A \oplus C$ would be the only interesting data about this exact complex. If we think of A, B, and C as being the generalizations of the numbers $\dim(A), \dim(B)$ and $\dim(C)$, then a short exact sequence is a way to encode that $\dim(A) + \dim(C) = \dim(B)$.

In the world of chain complexes, B could contain more data than just that of the vector spaces $A \oplus C$ – we need to additionally consider the information that comes from a differential.

Definition 30

Let $(A, d_A), (B, d_B), (C, d_C)$ be chain complexes. Let $i^\bullet : A^\bullet \to B^\bullet$ and $\pi^\bullet : B^\bullet \to C^\bullet$ be maps of cochain complexes. We say that

$$0 \longrightarrow A^\bullet \xrightarrow{i^\bullet} B^\bullet \xrightarrow{\pi^\bullet} C_\bullet \longrightarrow 0$$

is a short exact sequence of chain complexes if for all k,

$$0 \longrightarrow A^k \xrightarrow{i^k} B^k \xrightarrow{\pi^k} C^k \longrightarrow 0$$

is a short exact sequence of vector spaces.

The theory of short exact sequences of chain complexes is a lot richer than the theory for vector spaces, because chain complexes contain much more internal structure. We will now associate to each map $f^\bullet : A^\bullet \to B^\bullet$ a canonical short exact sequence.

Let $f^\bullet : A^\bullet \to B^\bullet$ be a map of cochain complexes. Define the *cone of f*, to be the cochain complex with

- Chain groups $\text{cone}^k(f) = A^{k+1} \oplus B^k$
- Differential defined by $d^k_{\text{cone}}(a,b) = (-d_A^{k+1}(a), d_B^k(b) + f^{k+1}(a))$.

31 Definition

Cone of Chain morphism

Note that for each k, $A^{k+1} \to \text{cone}^k(f) \to B^k$ is a short exact sequence. We should think of $\text{cone}^\bullet(f)$ as being the chain complex created by "attaching" $A^{\bullet+1}$ to $B^\bullet$.

$\text{cone}^\bullet(f)$ is a cochain complex.

32 Claim

Mapping cone is complex

Proof: A convenient notation for this proof will be to think of d^k_{cone} as having the form of a matrix:

$$d^k_{\text{cone}} = \begin{pmatrix} -d_A^{k+1} & 0 \\ f^{k+1} & d_B^k \end{pmatrix}.$$

We can then compute $d^{k+1}_{\text{cone}} \circ d^k_{\text{cone}}$ by using matrix multiplication.

$$\begin{aligned} d^{k+1}_{\text{cone}} d^k_{\text{cone}} &= \begin{pmatrix} -d_A^{k+2} & 0 \\ f^{k+2} & d_B^{k+1} \end{pmatrix} \begin{pmatrix} -d_A^{k+1} & 0 \\ f^{k+1} & d_B^k \end{pmatrix} \\ &= \begin{pmatrix} d_A^{k+2} \circ d_A^{k+1} & 0 \\ d_B^{k+1} \circ f^{k+1} - f^{k+2} \circ d_A^{k+1} & d_B^{k+1} \circ d_B^k \end{pmatrix} \end{aligned}$$

Using the definitions of chain map and chain differential,

$$= \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

∂^2

The cone of a morphism $f^\bullet : A^\bullet \to B^\bullet$ fits into a short exact sequence of chain complexes,

$$0 \longrightarrow B^\bullet \xrightarrow{i} \text{cone}^\bullet(f) \xrightarrow{\pi} A^{\bullet+1} \longrightarrow 0$$

where i, π are the natural inclusion and projection maps. Notice the shift in the index on the left hand side. A piece of notation that we will use for this shift in index is

$$C^{\bullet-1} = C^\bullet[-1].$$

The way that $A^{\bullet+1}$ is glued to $B^\bullet$ is dictated by the map $f^\bullet$. In this way, the exact sequence of chain complexes not only remembers that we can put $A^{\bullet+1}, B^\bullet$ together to build cone$^\bullet$, but also *how* these things were glued together.

From this short exact sequence, we surprisingly get a *long exact sequence* of homology groups.

Theorem 33 *SES-LES for mapping cones*

Let $f^\bullet : A^\bullet \to B^\bullet$ be a chain map. We have a short exact sequence of chain complexes

$$0 \longrightarrow B^\bullet \xrightarrow{i} \text{cone}^\bullet(f) \xrightarrow{\pi} A^\bullet[1] \longrightarrow 0$$

And we have the following long exact sequence of homology groups:

$$\cdots \xrightarrow{f} H^k(B) \xrightarrow{i} H^k(\text{cone}(f)) \xrightarrow{\pi} H^k(A[1]) \xrightarrow{f} H^{k+1}(B) \longrightarrow \cdots.$$

Proof: Showing that this is a long exact sequence amounts to checking that the sequence is exact at $H^k(B), H^k(\text{cone}(f)), H^k(A[1])$. We will show that the function is exact at $H^k(\text{cone}(f)) \to H^k(A[1]) \to H^{k+1}(B)$, which is perhaps the most surprising statement in the proof. To show the isomorphism

$$\ker(f : H^k(A[1]) \to H_{k+1}(B)) \simeq \text{Im}\,(\pi : H^k(\text{cone}(h)) \to H^k(A[1]),$$

we will show two inclusions.

We prove that $\ker(f : H^k(A[1]) \to H^{k+1}(B)) \subset \text{Im}\,(\pi : H^k(\text{cone}(f)) \to H^k(A[1])$. Take a cohomology class $[a] \in H^k(A[1])$ which is in the kernel of f so that

$$f([a]) = [0].$$

Since $\text{cone}^k(f) = A^k[1] \oplus B^k$, a natural candidate for an element of $\text{cone}^k(f)$ whose image under π is a would be $(a, 0)$. However, it may not be the case that this a homology class, as

$$d_{\text{cone}}(a, 0) = (d_A a, f(a))$$

which is not necessarily zero. As $[a] \in H^k(A[1])$, we are guaranteed that $d_A a = 0$. However, the only data that we have about $f(a)$ is that it is *cohomologous* to 0. Since $f([a]) = [0]$, there is an element $b \in B^k$ realizing the equivalence relation via $f(a) = d_B b$. Replacing our candidate element[2] with

$$\pi^{-1}(a) := (a, -b)$$

[2] The notation $\pi^{-1}(a)$ means that we have picked *an* inverse image of a under π. However, the map π is usually not invertible, and choices were made to produce this inverse image. In short, π^{-1} is not a map.

we can compute

$$\begin{aligned} \pi(\pi^{-1}(a)) =& \pi(a,-b) = a \\ d^{\text{cone}}(\pi^{-1}(a)) =& d_{\text{cone}}(a,-b) = 0 \end{aligned}$$

Therefore, $\ker(f : H^k(A[1]) \to H^{k+1}(B)) \subset \text{Im}\,(\pi : H^k(\text{cone}(f)) \to H^k(A[1])$.

The other direction is that $\ker(f : H^k(A[1]) \to H^{k+1}(B)) \supset \text{Im}\,(\pi : H^k(\text{cone}(f)) \to H^k(A[1])$. To show this, we need to show that the composition of $f \circ \pi = 0$ on cohomology. Let $[(a,b)] \in H^k(\text{cone}(f))$ be any element of homology. Since this is an element of homology, $d_{\text{cone}}(a,b) = 0$, and in particular,

$$f(a) = -d_B b.$$

We can use this when computing:

$$f \circ \pi[(a,b)] = f[(a)] = [-d_B b] = [0].$$

We omit the arguments for showing exactness at the other portions of the sequence.

This is sometimes notated in the following way:

$$\begin{array}{l} \cdots \xrightarrow{i} H^n(\text{cone}(f)) \xrightarrow{\pi} H^n(A[1]) \xrightarrow{f} \\ H^{n+1}(B) \xrightarrow{i} H^{n+1}(\text{cone}(f)) \xrightarrow{\pi} H^{n+1}(A[1]) \xrightarrow{f} \\ H^{n+2}(B) \xrightarrow{i} H^{n+2}(\text{cone}(f)) \xrightarrow{\pi} \cdots \end{array}$$

There is a useful corollary that follows from this construction:

34 Corollary *2-out of 3*

Suppose that $A^\bullet, B^\bullet$ are exact, and let $f^\bullet : A^\bullet \to B^\bullet$ be any map. Then $\text{cone}^\bullet(f)$ is exact.

Proof: By assumption $H^k(A) = H^k(B) = 0$ for all k. Therefore, we have the long exact sequence

$$\begin{array}{l} \cdots \xrightarrow{i} H^n(\text{cone}(f)) \xrightarrow{\pi} 0 \xrightarrow{f} \\ 0 \xrightarrow{i} H^{n+1}(\text{cone}(f)) \xrightarrow{\pi} 0 \xrightarrow{f} \\ 0 \xrightarrow{i} H^{n+2}(\text{cone}(f)) \xrightarrow{\pi} \cdots \end{array}$$

from which it follows that $H^k(\text{cone}(f)) = 0$ for all k. Therefore $\text{cone}^\bullet(f)$ is exact.

Inclusion-Exclusion

Let X be a set with a decomposition into smaller subsets, $X = \bigcup_{i\in I} U_i$. Let $U_J = \cap_{j\in J} U_i$. There exists an exact chain complex $CR^{\bullet}(\mathscr{U})$ with $CR^{\bullet}(\mathscr{U}) = \bigoplus_{J\subset I, |J|=k} \mathscr{F}(U_J)$.

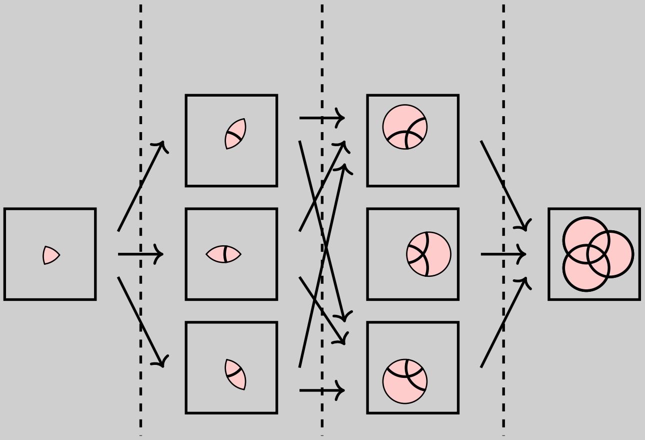

We will prove this theorem by using the tools of homological algebra, and induct on the size of I.

Definition.Let $\mathscr{U} = \{U_i\}_{i\in I}$ be a collection of subsets which cover X. Denote by $\mathscr{U}_\cap := \{U_j\}$

A covering $\mathscr{U} = \{U_i\}$ of X is a collection of subsets $U_i \subset X$ so that

$$X = \bigcup_{i\in I} U_i.$$

To each covering of X we will create an *resolution complex* $CR_{\bullet}(\mathscr{U})$.

Definition.Let $\mathscr{U} = \{U_i\}_{i\in I}$ be a covering of X. For each $J \subset I$, define the subset $U_J := X \cap (\bigcap_{i\in J} U_i)$. Suppose that J and K differ by a single index. We will then write $J \lessdot K$. Notice that whenever $K \gtrdot J$ we have an inclusion map $i_{K\gtrdot J} : U_K \to U_J$, and therefore we get an associated map

$$i^*_{K\gtrdot J} : \mathscr{F}(U_J) \to \mathscr{F}(U_K).$$

We define the chain groups

$$CR^k(\mathscr{U}) := \bigoplus_{K\subset I, |K|=k} \mathscr{F}(U_K)$$

and define the differential map to be

$$d^k_{CR} := \bigoplus_{K\gtrdot J} i^*_{K\gtrdot J}.$$

We will show that this gives us a chain complex by constructing it in a different fashion.

Lemma.Let $\hat{U}_1$ be the elements of X which only belong to U_1, Let $\mathscr{U}_X = \{U_i\}_{i\in I}$ be a cover of X. Let $\mathscr{U}_\cap = \{U_i \cap U_1\}_{1\neq i\in I}$ be a cover for $U_1 \setminus \hat{A}_1$. Let $\mathscr{U}_\setminus = \{U_i\}_{1\neq i\in I}$ be a cover for $X \setminus \hat{U}_1$. Then there is a natural maps $i_J : \bigcap_{i\in J}(U_i \cap U_1) \to \bigcap_{i\in J}(U)i)$ for each J, inducing a map

$$i^* : CR_\bullet(\mathscr{U}_\setminus) \to CR_\bullet(\mathscr{U}_\cap)$$

and $CR_\bullet(\mathscr{U}_X) = \text{cone}(i^*) \oplus (\mathscr{F}(\hat{U}_1) \to \mathscr{F}(\hat{U}_1))$

As always, a diagram explains the core concept of this proof:

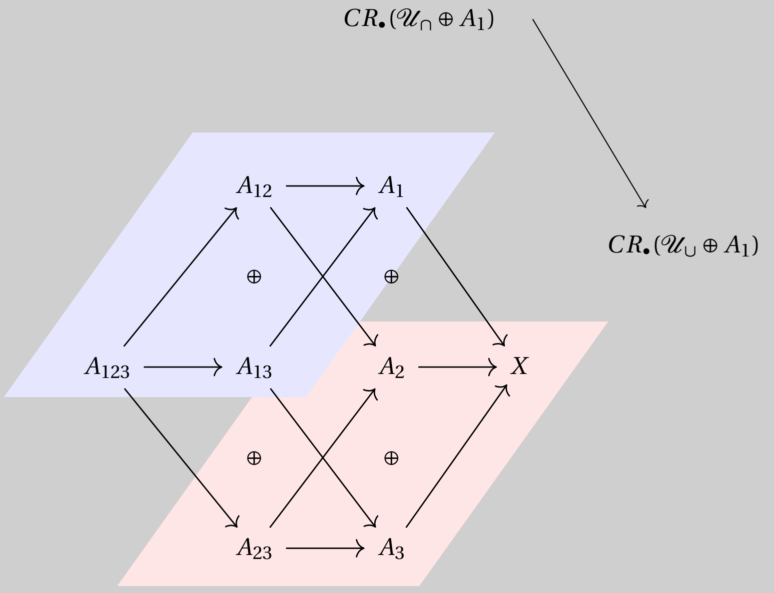

Corollary.The homology of the resolution complexes are trivial: $H_\bullet(CR_\bullet(\mathscr{U})) = 0$, i.e. $CR_\bullet(\mathscr{U})$ is exact.

Proof: We again prove by induction on the size of the cover. As a base case, we can let $\mathscr{U} = \{X\}$, then $H_\bullet(\mathscr{U}) = 0$ trivially.
Now assume that we know by induction that $CR_\bullet(\mathscr{U}_\cap)$ and $CR_\bullet(\mathscr{U}_\cup)$ have trivial homology. Since the cone of exact chain complexes is exact, we get $CR_\bullet(\mathscr{U})$ is exact. ∂^2

5 Mayer Vietoris

We finally return to one of the core concepts of this course: given a decomposition of a space $X = A \cup B$, what can we tell about the topology of X in terms of the topology of A and B?

At the start of the course, we alluded that we would like an algorithm to compute the number of connected components via an inclusion-exclusion principle on a decomposition of X into smaller topological spaces. Let's look at an example where this works, and an example that shows that our theory requires some more depth.

Example 36

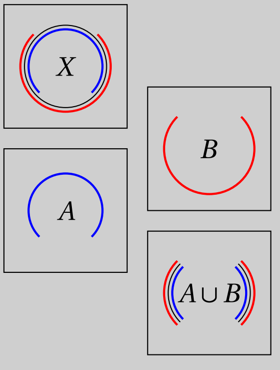

Let $S^1 = A \cup B$ as drawn in the figure. Let's try to compute the number of connected components of S^1 using this decomposition. $A \cap B$ has two connected components, so we would have that

$$b_0(A) + b_0(B) - b_0(A \cap B) = 0$$

which means that we cannot use the principle of inclusion-exclusion to compute the number of connected components of the circle. The obstruction in this case to the principle of inclusion-exclusion working is the presence of nontrivial homology in $H^1(S^1)$.

While we cannot use the principle of inclusion-exclusion to compute the number connected components, we can get an inclusion-exclusion like principle to work homologically. For full details on how to generalize inclusion-exclusion like principles to general settings, see Appendix **??**.

Theorem 37 *Mayer-Vietoris*

Let A, B, X be topological spaces. Let

$$j_A : A \to X$$
$$j_B : B \to X$$

be two inclusions of topological spaces so that $A \cup B = X$. Let $A \cap B$ be the common intersection of A and B in X, with the natural inclusions

$$i_A : A \cap B \to A$$
$$i_B : A \cap B \to B$$

Then there is a short exact sequence of chain complexes

$$0 \longrightarrow \underline{C}^\bullet(X) \begin{array}{c} \xrightarrow{j_A^*} \underline{C}^\bullet(A) \xrightarrow{i_A^*} \\ \oplus \\ \xrightarrow{j_B^*} \underline{C}^\bullet(B) \xrightarrow{-i_B^*} \end{array} \underline{C}^\bullet(A\cap B) \longrightarrow 0$$

This in turn gives us a long exact sequence on homology from Lemma **??**.

$$\cdots \to H^{k-1}(A\cap B) \to H^k(X) \to H^k(A)\oplus H^k(B) \to H^k(A\cap B) \to H^{k+1}(X) \to \cdots$$

Proof: To show that this is an exact sequence, we need to check that the chain maps form exact sequences of vector spaces at each grading k:

$$0 \longrightarrow \underline{C}^k(X) \xrightarrow{j_A^*\oplus j_B^*} \underline{C}^k(A)\oplus \underline{C}^k(B) \xrightarrow{i_A^*\oplus(-i_B^*)} \underline{C}^k(A\cap B) \longrightarrow 0.$$

Let's start by checking exactness at the first position of the sequence.

$$0 \longrightarrow \underline{C}^k(X) \xrightarrow{j_A^*\oplus j_B^*} \underline{C}^k(A)\oplus \underline{C}^k(B)$$

The statement of exactness at this point is that $\ker(j_A^*\oplus j_B^*) = 0$, or that the map is injective. Recall that $\underline{C}^k(X), \underline{C}^k(A)$ and $\underline{C}^k(B)$ are continuous $\mathbb{Z}_2$ labellings of the k-intersections of the covering sets U_i. Given $U_\sigma \subset X$ a k-fold intersection of open sets, it is either the case that $U_\sigma \subset A$ or $U_\sigma \subset B$. As a result, given $\phi \in \underline{C}^\bullet(X)$, the labelling of U_σ can be determined by its image under the map j_A^* or j_B^*. This means that the labelling ϕ can be recovered from $(j_A^*\oplus j_B^*)(\phi)$, so $(j_A^*\oplus j_B^*)$ is injective.

At the last position of the sequence,

$$\underline{C}^k(A)\oplus \underline{C}^k(B) \xrightarrow{i_A^*\oplus(-i_B^*)} \underline{C}^k(A\cap B) \longrightarrow 0.$$

exactness means that $\operatorname{Im} i_A^* \oplus i^* B_B = \underline{C}^k(X)$ i.e. $i_A^*\oplus i_B^*$ is surjective. In fact, i_A^* is already surjective, as $U_\sigma \subset A\cap B$ is contained in $U_\sigma \subset A$, and therefore every labelling of an open set in $\underline{C}^k(A\cap B)$ can be lifted to a labelling of open sets in $\underline{C}^k(A)$ and extended by zero over $\underline{C}^k(B)$.

The remaining tricky part of the argument is on the middle section,

$$\underline{C}^k(X) \xrightarrow{j_A^*\oplus j_B^*} \underline{C}^k(A)\oplus \underline{C}^k(B) \xrightarrow{i_A^*\oplus(-i_B^*)} \underline{C}^k(A\cap B)$$

Here, the statement is that $\ker(i_A^*\oplus(-i_B^*)) = \operatorname{Im}(j_A^*\oplus j_B^*)$. The kernel of the map $(j_A\oplus(-j_B))$ consists exactly of labellings of the k-fold intersections on A and B

which agree on the intersection. These are exactly the labellings which are in the image of $j_A^* \oplus j_B^*$.

Once we know that the short sequence of chain complexes is exact, the long exact sequence of homology groups

$$\cdots \to H^{k-1}(A \cap B) \to H^k(X) \to H^k(A) \oplus H^k(B) \to H^k(A \cap B) \to H^{k+1}(X) \to \cdots$$

follows from the application of the Zig-Zag Lemma (Station 6.)

We usually represent the Mayer-Vietoris long exact sequence with the following diagram of homology groups :

$$\cdots \xrightarrow{j_A^* \oplus j_B^*} H^k(A) \oplus H^k(B) \xrightarrow{i_A^* \oplus i_B^*} H^k(A \cap B) \xrightarrow{\delta}$$

$$H^k(X) \xrightarrow{j_A^* \oplus j_B^*} H^{k+1}(A) \oplus H^{k+1}(B) \xrightarrow{i_A^* \oplus i_B^*} H^{k+1}(A \cap B) \xrightarrow{\delta}$$

$$H^{k+1}(X) \xrightarrow{j_A^* \oplus j_B^*} H^{k+2}(A) \oplus H^{k+2}(B) \xrightarrow{i_A^* \oplus i_B^*} \cdots$$

The maps i^* and j^* somewhat act in a normal way: cycles in the spaces X, A, B and $A \cap B$ are related to each other. We now will try to figure out what the map δ does.

This requires a better geometric understanding of what each homology class means. Each element of $\underline{C}^k(X)$ represents a labelling of the k-simplices of X, and the differential map "pushes" those labellings to the higher simplices.

A label represents a non-trivial class in $H^k(X)$ if, when pushed to the higher dimensional simplices it cancels out, and the labelling itself does not arise from a lower-dimensional labelling.

Suppose that we have a labelling ϕ of the simplices of $A \cap B$ giving us a cohomology class. This means that the "push" of the labelling on $A \cap B$ to the higher simplices *inside of* $A \cap B$ will cancel out. Let us take ϕ some labelling of the k-simplices on $A \cap B$ representing some cohomology class. Use this to create a labelling ϕ_A on A and a labelling ϕ_B on B. Even though $d_{A\cap B}\phi$ equals zero, the extended labellings may not have this property, and so $d_A\phi_A$ and $d_B\phi_B$ are some interesting labellings to talk about. They, in some sense, represent the "boundary" $A \cap B$ inside of A and B.

Let's now use both $d_A\phi_A$ and $d_B\phi_B$ to create a labelling for all of X. We take $d_A\phi_A + d_B\phi_B$ as a labelling on all of X. This element is, surprisingly, closed.

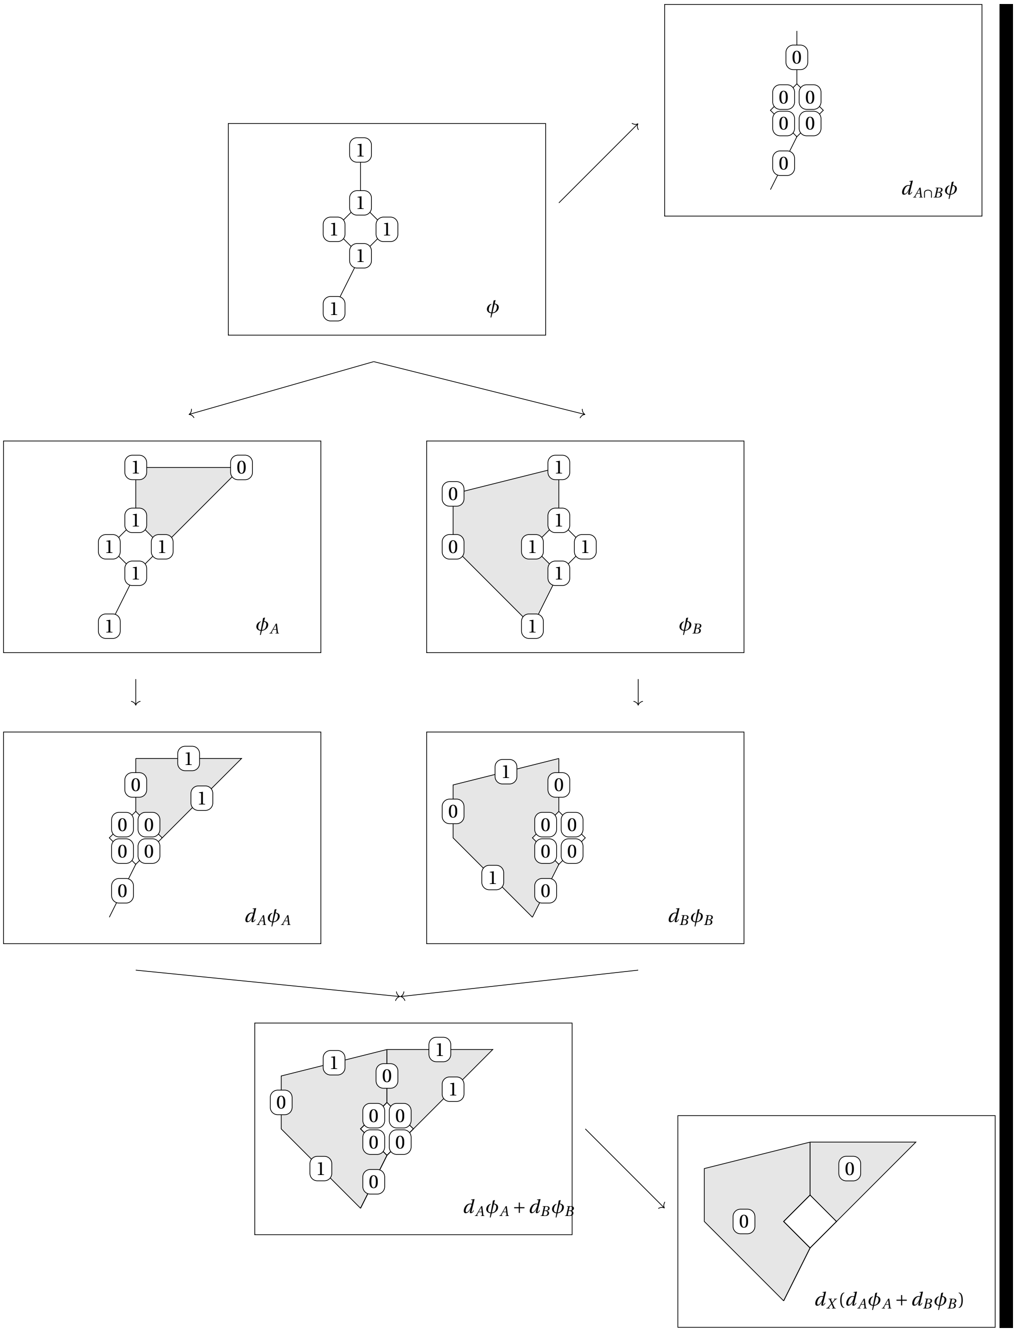
$d_{A\cap B}\phi$
ϕ
ϕ_A
ϕ_B
$d_A\phi_A$
$d_B\phi_B$
$d_A\phi_A + d_B\phi_B$
$d_X(d_A\phi_A + d_B\phi_B)$

38 Homology of Sphere

Let's compute the homology of sphere S^n by using Mayer-Vietoris and induction. For this example, we will start with the assumptions that we know the homology of a disk **??**.

We will prove that $H^k(S^n) = \mathbb{Z}_2$ if and only if $k = n, 0$ by induction on n. Here, we will run the Mayer-Vietoris argument on a the decomposition of S^n into two disks, $A, B = D^n$, which are suppose to represent the upper and lower hemispheres. Notice that the intersection of the two hemispheres is the equatorial sphere, which is a sphere of 1-dimension lower.

191figures/simp_spheredecomp.pdf

So, we have a short exact sequence of chain complexes:

$$0 \to C^\bullet(S^n) \to C^\bullet(D^n) \oplus C^\bullet(D^n) \to C_\bullet(S^{n-1}) \to 0$$

This short exact sequence gives us a long exact sequence of homology groups :

$$\begin{array}{ccccccc}
H^0(S^n) & \longrightarrow & H^0(D^n) \oplus H^0(D^n) & \longrightarrow & H^0(S^{n-1}) & \xrightarrow{\delta} & \\
H^1(S^n) & \longrightarrow & H^1(D^n) \oplus H^1(D^n) & \longrightarrow & H^1(S^{n-1}) & \xrightarrow{\delta} & \\
H^2(S^n) & \longrightarrow & \cdots & \longrightarrow & H^{n-2}(S^{n-1}) & \xrightarrow{\delta} & \\
H^{n-1}(S^n) & \longrightarrow & H^{n-1}(D^n) \oplus H^{n-1}(D^n) & \longrightarrow & H^{n-1}(S^{n-1}) & \xrightarrow{\delta} & \\
H^n(S^n) & \xrightarrow{i_0 \oplus j_0} & H^n(D^n) \oplus H^n(D^n) & \longrightarrow & H^n(S^{n-1}) & \longrightarrow & 0
\end{array}$$

Substituting in the groups we know from induction and our assumptions

$$H^0(S^n) \longrightarrow \mathbb{Z}_2 \oplus \mathbb{Z}_2 \longrightarrow \mathbb{Z}_2 \xrightarrow{\delta}$$

$$H^1(S^n) \longrightarrow 0 \longrightarrow 0 \xrightarrow{\delta}$$

$$H^2(S^n) \longrightarrow \cdots \longrightarrow 0 \xrightarrow{\delta}$$

$$H^{n-1}(S^n) \longrightarrow 0 \longrightarrow \mathbb{Z}_2 \xrightarrow{\delta}$$

$$H^n(S^n) \xrightarrow{i_0 \oplus j_0} 0 \longrightarrow 0 \longrightarrow 0$$

We therefore may now look at these shorter exact sequences instead:

$$0 \to \mathbb{Z}_2 \to H^n(S^n) \to 0$$

$$0 \to H^k(S^n) \to 0 \qquad k \neq n, 0$$

$$0 \to H^0(S^n) \to Z_2 \oplus Z_2 \to Z_2 \to H^1(S^n) \to 0$$

Running through the properties of exactness at each part shows confirms our computation of the homology of S^n.

6 Inclusion-Exclusion principles: The Zig-Zag Lemma

Let's now use Inclusion-Exclusion to build up some more intuition on what homological algebra can get us. We will now work a little abstractly.

Let $\mathscr{C}$ be a collection of objects. Let's suppose that objects in this collection admit decompositions, so that we may write [3]

$$X = A \cup B$$

and for every such decomposition, we may also associate an objects called $A \cap B$.

A *property* is a function $P : \mathscr{C} \to \mathbb{N}$ which assigns to each object a number.

Definition 39 Let $\mathscr{C}$ be a category, and $P : \mathscr{C} \to \mathbb{N}$ be a property. We say that P obeys the *homological inclusion-exclusion principle* if for all X, there exists a chain complex $P_\bullet(X)$ satisfying the following conditions:

- *Recovery of P:* We have that $\dim H_0(P_\bullet(X)) = P(X)$.
- *Inclusion-Exclusion:* Whenever $X = A \cup B$, we have a short exact sequence:

$$0 \to P_\bullet(A \cap B) \to P_\bullet(A) \oplus P_\bullet(B) \to P_\bullet(X) \to 0.$$

Notice that satisfying a homological inclusion-exclusion principle is in a lot of ways like satisfying a inclusion-exclusion principle, in that

$$\dim(P_0(X)) = \dim(P_0(A)) + \dim(P_0(B)) - \dim(P_0(A \cap B)).$$

While we don't get an actual inclusion exclusion principle from a homological inclusion-exclusion principle, we get something very close to the principle holding. In order to see the relation between inclusion-exclusion and homological inclusion-exclusion, we need a powerful lemma from homological algebra.

Theorem 40 *Zig-Zag Lemma* Let $A_\bullet, \partial_\bullet^A$, $B_\bullet, d_\bullet^B$ and $C_\bullet, d_\bullet^C$ be chain complexes. Given

$$0 \longrightarrow A_\bullet \xrightarrow{f} B_\bullet \xrightarrow{g} C_\bullet \longrightarrow 0$$

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

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

[3] We also adopt homological grading conventions in this section, as opposed to cohomological grading conditions.

Before we get into a proof of this theorem, let's quickly make a remark on the map δ. On the one hand, the map is remarkable, as there is no reason to expect a map connecting $C \to A$. However, we've seen the existence of a long exact sequence that arises from a short exact sequence before when we looked at cones.

In a certain sense, this theorem says that *all short exact sequences of chain complexes essentially arise from the cone sequence.* While we will not be able to prove this result in this class, one can make a version of this statement true by exploring the derived category and triangulated structures.

Foo

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

$$
\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow{\scriptstyle \partial} & & \downarrow{\scriptstyle \partial^B} & & \downarrow{\scriptstyle \partial^C} & & \\
0 & \longrightarrow & A_{n+1} & \xrightarrow{f_{n+1}} & B_{n+1} & \xrightarrow{g_{n+1}} & C_{n+1} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial_{n+1}} & & \downarrow{\scriptstyle \partial^B_{n+1}} & & \downarrow{\scriptstyle \partial^C_{n+1}} & & \\
0 & \longrightarrow & A_n & \xrightarrow{f_n} & B_n & \xrightarrow{g_n} & C_n & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial^A_n} & & \downarrow{\scriptstyle \partial^B_n} & & \downarrow{\scriptstyle \partial^C_n} & & \\
0 & \longrightarrow & A_{n-1} & \xrightarrow{f_{n-1}} & B_{n-1} & \xrightarrow{g_{n-1}} & C_{n-1} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial^A_{n-1}} & & \downarrow{\scriptstyle \partial^B_{n-1}} & & \downarrow{\scriptstyle \partial^C_{n-1}} & & \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}
$$

We want to construct a function δ from $H_n(C)$ to $H_{n-1}(A)$. The following argument is an *element chasing argument*, which can be a bit difficult to follow through; it's suggested that the reader write out the argument step-by-step at some point on their own to see where the maps come from.
Since this lemma contains several statements, we will check some of them and leave the remainder as exercises.

There exists a canonical map $\delta : H_k(C) \to H_{k-1}(A)$.

41 Claim

As mentioned before, we should somewhat expect the existence of this map from our studies of cones. First, let's try and show that to a homology class $[\gamma] \in H_k(C)$, we can find an element in A_{k-1}

– As the map g_n is surjective, we know that we can pick an element in the preimage β so that $g_n(\beta) = \gamma$. *Notice that this is not a canonical choice!*

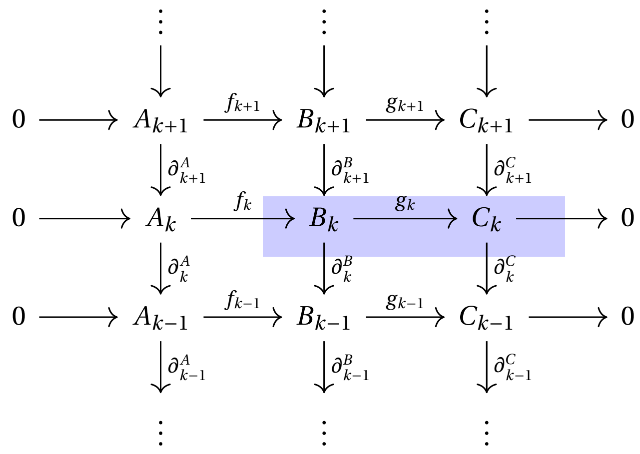

– We can apply $\partial_n^B(\beta)$ and we wind up with an element in B_{n-1}. Using that g_{n-1} is a chain map, we get that

$$g_{n-1}\partial_n^B(\beta) = \partial^C g_n(\beta) = \partial^C \gamma = 0$$

where the second equality comes from the fact that γ represents a homology class.

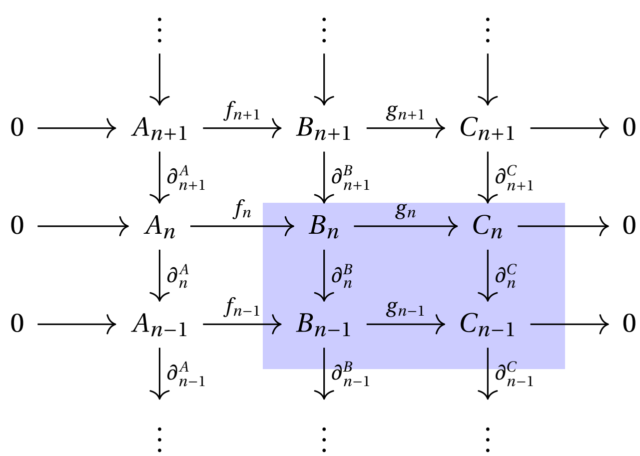

– Since $\partial_n^B(\beta) \in \ker g_{n-}1$, and the sequence is chain complexes is exact, we know that $\partial^B_(\beta) \in$ Im f_{n-1}. Since f_{n-1} is injective, we know that there is unique α corresponding to this β so that $f_{n-1}(\alpha) = \partial^B(\beta)$.

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & A_{n+1} & \xrightarrow{f_{n+1}} & B_{n+1} & \xrightarrow{g_{n+1}} & C_{n+1} & \longrightarrow & 0 \\
 & & \downarrow \partial^A_{n+1} & & \downarrow \partial^B_{n+1} & & \downarrow \partial^C_{n+1} & & \\
0 & \longrightarrow & A_n & \xrightarrow{f_n} & B_n & \xrightarrow{g_n} & C_n & \longrightarrow & 0 \\
 & & \downarrow \partial^A_n & & \downarrow \partial^B_n & & \downarrow \partial^C_n & & \\
0 & \longrightarrow & A_{n-1} & \xrightarrow{f_{n-1}} & B_{n-1} & \xrightarrow{g_{n-1}} & C_{n-1} & \longrightarrow & 0 \\
 & & \downarrow \partial^A_{n-1} & & \downarrow \partial^B_{n-1} & & \downarrow \partial^C_{n-1} & & \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}$$

– We initially define $\delta[\gamma] = \alpha$.

We now need to show that α is a homology class, that is, that $\partial^A_{k-1}(\alpha) = 0$.

– Look at $\partial^A_{k-1}(\alpha)$. Since this diagram is commutative, we have that $f_{k-2}\partial^A_{k-1}(\alpha) = \partial^B_{k-1} f_{k-1}(\alpha)$.

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & A_k & \xrightarrow{f_k} & B_k & \xrightarrow{g_k} & C_k & \longrightarrow & 0 \\
 & & \downarrow \partial^A_k & & \downarrow \partial^B_k & & \downarrow \partial^C_k & & \\
0 & \longrightarrow & A_{k-1} & \xrightarrow{f_{k-1}} & B_{k-1} & \xrightarrow{g_{k-1}} & C_{k-1} & \longrightarrow & 0 \\
 & & \downarrow \partial^A_{k-1} & & \downarrow \partial^B_{k-1} & & \downarrow \partial^C_{k-1} & & \\
0 & \longrightarrow & A_{k-2} & \xrightarrow{f_{k-2}} & B_{k-2} & \xrightarrow{g_{k-2}} & C_{k-2} & \longrightarrow & 0 \\
 & & \downarrow \partial^A_{k-2} & & \downarrow \partial^B_{k-2} & & \downarrow \partial^C_{k-2} & & \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}$$

– Recalling or definition of α, we know that $f_{k-1}(\alpha) = \partial^B_k(\beta)$, so $\partial^B_{k-1}(\partial^B_k(\beta) =$

$f_{k-2}(\partial_{k-1}\alpha) = 0$. Since f_{k-2} is injective, we get that $\partial_{k-1}\alpha) = 0$.

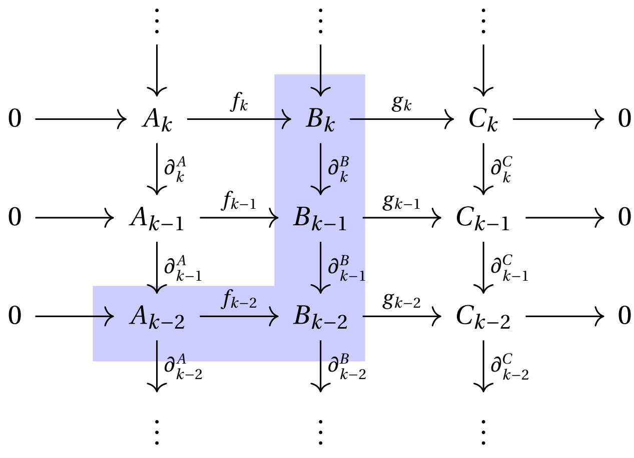

Finally, when we constructed the class α, we had to make a choice of $\beta = g_k^{-1}(\gamma)$. Let's show that the homology class of α does not depend on the choice of β lifting α.

– Suppose that β, β' are two different liftings of γ so that $g_k(\beta) - g_k(\beta') = 0$. We want to show that the associated classes $[\alpha], [\alpha']$ are homologous. Since $g_k(\beta - \beta') = 0$, there exists a class $f_k^{-1}(\beta - \beta')$ due to exactness of the row.

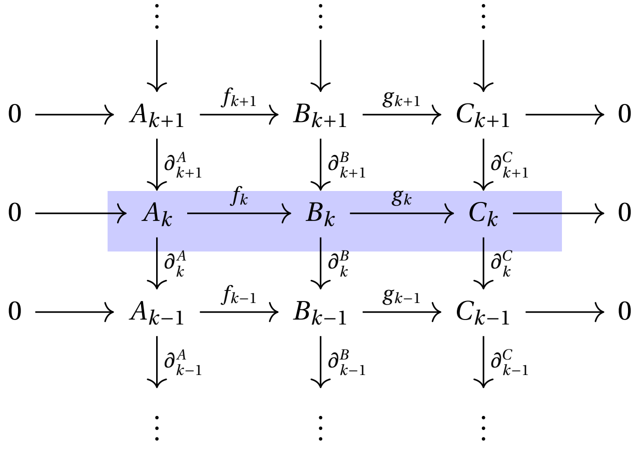

– Due to commutativity of the highlighted square, we have that $f_{k-1}\partial_k^A(f_k^{-1}(\beta - \beta') = \partial_k^B(\beta - \beta') = f_{k-1}(\alpha - \alpha')$. Due to the injectivity of f_{k-1}, we conclude that $\alpha - \alpha' = \partial_k^A(f_k^{-1}(\beta - \beta')$, so these two classes are cohomologous.

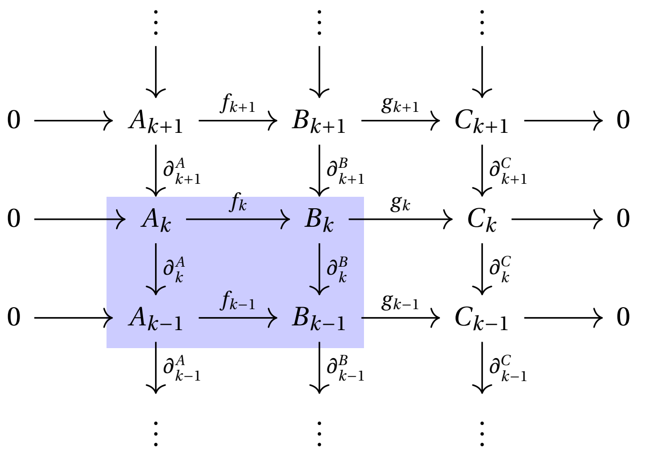

This completes the proof that the map δ is well defined on homology. Now we will show some of the exactness statements.

42 Claim

The sequence of homology groups

$$H_k(B) \xrightarrow{g_k} H_k(C) \xrightarrow{\delta_k} H_{k-1}(A)$$

is exact.

In order to prove this claim, we need to show that $\ker(\delta) \subset \text{Im}\,(g_k)$, and $\text{Im}\,(g_k) \subset \ker\delta$.

- To show that $\text{Im}\,(g_k) \subset \ker\delta$, it suffices to show that the composition $\delta_k \circ g_k = 0$. Let $[\beta] \in H_k(B)$ be a homology class. Then $[\delta_k g_k(\beta)] = [f_{k-1}^{-1}(\partial_k^B \beta)]$. Since $[\beta]$ is a class in homology, the boundary map starts by computing $\partial_k^B \beta = 0$, and we conclude that $\delta_k(g_k(\beta)) = 0$.

- To show that the $\ker(\delta_k) \subset \text{Im}\,(g_k)$, let γ be an element so that $\delta_k[\gamma] = 0$. Since the map $g_k : B_k \to C_k$ is surjective, we might hope that $\beta = g_k^{-1}\gamma$, a choice of lift of γ, is a class in homology. So we need to show that $\partial_k^B(\beta) = 0$. By commutativity of the lower right square, we have that $\partial_k^B(\beta) =$

$f_{k-1}(\delta(\gamma)) = 0.$

$$
\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & A_{k+1} & \xrightarrow{f_{k+1}} & B_{k+1} & \xrightarrow{g_{k+1}} & C_{k+1} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial^A_{k+1}} & & \downarrow{\scriptstyle \partial^B_{k+1}} & & \downarrow{\scriptstyle \partial^C_{k+1}} & & \\
0 & \longrightarrow & A_k & \xrightarrow{f_k} & B_k & \xrightarrow{g_k} & C_k & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial^A_{k}} & & \downarrow{\scriptstyle \partial^B_{k}} & & \downarrow{\scriptstyle \partial^C_{k}} & & \\
0 & \longrightarrow & A_{k-1} & \xrightarrow{f_{k-1}} & B_{k-1} & \xrightarrow{g_{k-1}} & C_{k-1} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial^A_{k-1}} & & \downarrow{\scriptstyle \partial^B_{k-1}} & & \downarrow{\scriptstyle \partial^C_{k-1}} & & \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}
$$

Claim 43

The sequence of homology groups

$$H_{k+1}(C) \xrightarrow{\delta_{k+1}} H_k(A) \xrightarrow{f_k} H_k(B)$$

is exact.

Chain Homotopy

Homological algebra is ultimately the study of which chain complexes are isomorphic to each other in a homological way

7

44 Definition *Quasi-isomorphism*

Let $f : A_\bullet \to B_\bullet$ be a chain map. Then we say that f is a *quasi-isomorphism* if the induced map on homology, $f_* : H_i(\mathscr{A}) \to H_i(\mathscr{B})$ are isomorphisms of homology groups.

Notice that while every isomorphism which is a chain map gives us a quasi-isomorphism, a chain map need not be an isomorphism to be a quasi-isomorphism $\blacksquare$

45 Example *Non-isomorphic, but quasi-isomorphic*

Not isomorphic, but quasi-isomorphic.

Similarly, even if (A, ∂) and (B, ∂) have isomorphic homology groups, they need to not be quasi-isomorphic.

46 Example

It is not necessarily the case that if two chain complexes have isomorphic homology that those two complexes are quasi isomorphic.

Even though $f_\bullet : A_\bullet \to B_\bullet$ is a quasi-isomorphism, there is no guarantee that there exists $g_\bullet : B_\bullet \to A_\bullet$ so that the maps $(g \circ f)_k : H_k(A) \to H_k(A)$ is the identity. In other words, there is no need for inverses to exists to quasi-isomorphism on either the chain or homological level. If such a map exists, we call it a *quasi-inverse*.

47 Example *Non-inverticble quasi-isomorphism*

Chain Complexes with no quasi-inverse.

It is usually hard come up with an interpretation of where a quasi-isomorphism comes from; in general the question if two maps $f, g : A_\bullet \to B_\bullet$ do the same thing on homology is hard to get some intuition on. As a proxy to showing that two maps have the same definition on homology, we introduce an idea from topology: that of a *homotopy*.

Definition 48
Chain Homotopy

Let (A, ∂^A) and (B, ∂^B) be chain complexes. Let $f_\bullet : A_\bullet \to B_\bullet$ and $g_\bullet : A_\bullet \to B_\bullet$ 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

$$f - g = \partial^B h_i + h_{i-1}\partial^A$$

We write that $f \sim g$.

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

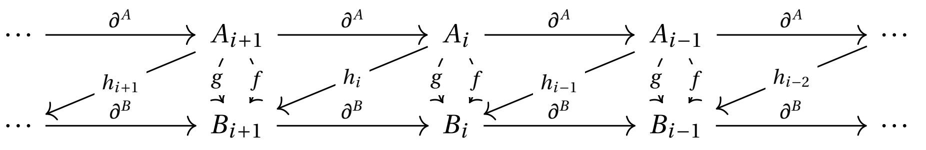

It can be difficult to get an intuition on what a chain homotopy between two map constitutes. One interpretation comes from topology; for every element x, the difference between $f(x)$ and $g(x)$ can be expressed as a cylinder connecting $f(x), g(x)$. This bears resemblance to the definition of a homotopy between two maps in point-set topology.

Example 49

An example of a chain homotopy will go here!

One thing worth pointing out is that we don't have any condition of compatibility with the differential for the homotopy maps $h_k : A_k \to B_{k+1}$; they are allowed to be as crazy as need be. Chain homotopy is especially useful for the following lemma:

Lemma 50
Homotopic maps agree on homology

Suppose that $f_\bullet, g_\bullet : A_\bullet \to B_\bullet$ are chain homotopic chain maps. Then they are the same map on homology, in the sense that

$$f_k[x] = g_k[x]$$

for every $[x] \in H_k(A)$.

Proof: What we want to show is that $f_k - g_k = \partial^B_{k+1} h_k + h_{k-1}\partial^A_k$, then for every $[a] \in H_k(A)$, there exists $b \in C_{k+1}(B)$ with

$$f_k(a) - g_k(a) = \partial_{k+1}(b).$$

The homotopy gives us a natural for b is; we can let $b = h_k(a)$. Taking our definition of homotopy shows

$$f_k(a) - g_k(a) = \partial^B_{k+1} h_k(a) + h_{k-1}\partial^A_k(a)$$

As a represents a class in homology

$$= \partial^B_{k+1} h_k(a) = \partial^B_{k+1}(b).$$

∂^2

This next claim shows the usefulness of chain homotopies:

51 Lemma *Homotopic to Identity*

Let $f : \mathscr{A} \to \mathscr{B}$ and $g : \mathscr{B} \to \mathscr{A}$ be chain maps. Suppose that $g \circ f \sim 1_{\mathscr{A}}$ and $g \circ f \sim 1_{\mathscr{B}}$. Then f and g are quasi-isomorphisms.

Proof: Let's start with a diagram.

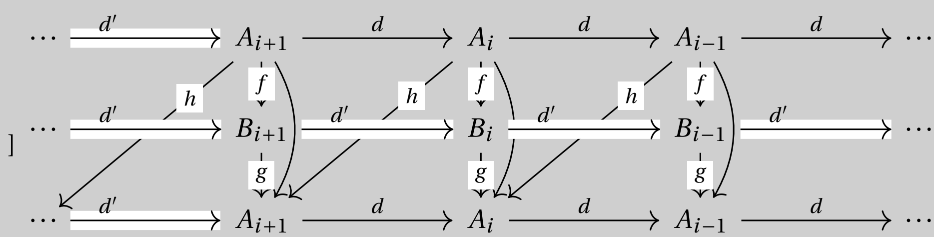

The homotopy to the identity map gives us that there exists h such that $g \circ f - 1_{\mathscr{A}} = dh^i + h^{i-1}d$. Suppose that $v \in H_i A$. Then v is in the kernel of d, so $h^{i-1}d(v) = h^{i-1}(0) = 0$. We have that therefore $g \circ f - 1_{\mathscr{A}} \in \operatorname{Im}(d)$, which is to say that on homology $g \circ f = 1_{\mathscr{A}}$, as we mod out by $\operatorname{Im}(d)$ when we take homology. Of course, a similar proof shows that $f \circ g = 1_{\mathscr{A}}$ ∂^2

P Exercises

The zero vector space, 0, is the vector space which only has one element in it.

Exercise P15

Let V_1 and V_2 be vector spaces. Suppose that $f : V_1 \to V_2$ is a linear map. Show that $\ker(f) = \{0\}$ if and only if the map $f : V_1 \to V_2$ is injective.

Exercise P16

Suppose we have 5 vectors spaces and maps between them.

$$V^0 \xrightarrow{d^0} V^1 \xrightarrow{d^1} V^2 \xrightarrow{d^2} V^3 \xrightarrow{d^3} V^4$$

and suppose that $\operatorname{Im} d^i = \ker d^{i+1}$ for each i.

- Show that $V^0 = 0$, then d^1 is injective.

$$0 \xrightarrow{d^0} V^1 \xrightarrow{d^1} V^2$$

- Show that if $V^4 = 0$, then d^2 is surjective.

$$V^2 \xrightarrow{d^2} V^3 \xrightarrow{d^3} 0$$

- Show that if $V^0 = V^3 = 0$, then $d^1 : V^1 \to V^2$ is an isomorphism.

$$0 \xrightarrow{d^0} V^1 \xrightarrow{d^1} V^2 \xrightarrow{d^2} 0$$

- Show that if $V^0 = V^4 = 0$, then $\dim(V^1) + \dim(V^3) = \dim(V^2)$.

$$0 \xrightarrow{d^0} V^1 \xrightarrow{d^1} V^2 \xrightarrow{d^2} V^3 \xrightarrow{d^3} 0$$

- Furthermore, show that there is a non-canonical isomorphism of vector spaces, $V^2 = V^1 \oplus V^3$.

P17 Exercise
Translating Sets in to Vector Spaces

Let A be any finite set. Let $\mathscr{F}(A)$ be the set of functions $\phi : A \to \mathbb{Z}_2$.

- Prove that there are $2^{|A|}$ such functions.
- Prove that $\mathscr{F}(A)$ is a $\mathbb{Z}_2$ vector space.
- Prove that $\dim(\mathscr{F}(A)) = |A|$.

P18 Exercise
Categories and Functors

Show that if $f : A \to B$ and $g : B \to C$ are two maps of sets, then

$$(g \circ f)^* = f^* \circ g^*,$$

i.e. the pullback relation preserves compositions.

P19 Exercise

Let $S_1, S_2 \subset A$ be two subsets as before.

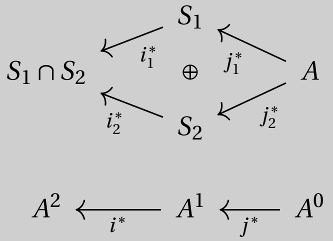

Prove that the map i^* is surjective.

P20 Exercise
Open Ended Exercise

Suppose that S_1, S_2 and S_3 are three sets, and $A = S_1 \cup S_2 \cup S_3$. Describe how one would extend the Inclusion-Exclusion formula to this setting using the linear algebra machinery that we set up before.

Exercise P21

Let $U \subset V$ be a subspace of a vector space. Consider the equivalence relation

$$v_1 \sim_U v_2 \text{ if and only if } v_1 - v_2 \in U.$$

Show that the quotient space $V/U := \{[v]_{\sim_U}\}$ given by the set of equivalence classes is a vector space.

Exercise P22

Let $U \subset V$ be a subspace of a vector space. Construct an exact chain complex

$$0 \to U \to V \to V/U \to 0$$

Exercise P23

Let G be a graph – a simplicial complex with only 0 and 1 dimensional simplices. The spaces $C^0(G, \mathbb{Z}_2)$ and $\underline{C}^1(G, \mathbb{Z}_2)$ have basis given by the vertices and edges of the graph. Describe d^0 as a matrix in terms of this basis.

Exercise P24

Show that whenever $e_1, \ldots, e_k$ sequence of edges with k odd which form a cycle in G, then one of $e_1 + \ldots + e_k \in C^1(G, \mathbb{Z}_2)$ is not in the image of d^0. Make a similar conclusion for when k is even. Conclude that if G has a cycle, $\underline{H}^1(G) := H^1(\underline{C}^\bullet(G, \mathbb{Z}_2))$ is at least 1-dimensional.

Exercise P25

Show that $\underline{H}^0(G)$ is one fewer than the number of connected components in G .

Exercise P26

Show that $\underline{H}^1(G) = 0$ if and only if G is a tree.

Suppose that G has one connected component. Compute the dimension of $H^1(G)$ in terms of the number of edges and vertices of G.

P27 Exercise

Let S^2 be the simplicial complex defined by the tetrahedron (do not include the interior 3-simplex, but only the 4 faces.) Show that $\underline{H}^0(S^2) = 0, \underline{H}^2(S^2) = \mathbb{Z}_2$ and $H^1(S^2) = 0$.

P28 Exercise

Let $C^i(X, \mathbb{Z}_2)$ be the cochain complex associated to a simplicial space. Show that if X has only one connected component then $\underline{H}^0(\mathbb{Z}_2) = 0$.

P29 Exercise

In class, we looked at one configuration of open sets which covered the circle. We will look at some examples where we use multiple sets to cover a topological space.

Let X be the line segment drawn below, covered by two sets U_1 and U_2. Repeat the connected component construction for the line covered with two sets.

P30 Exercise

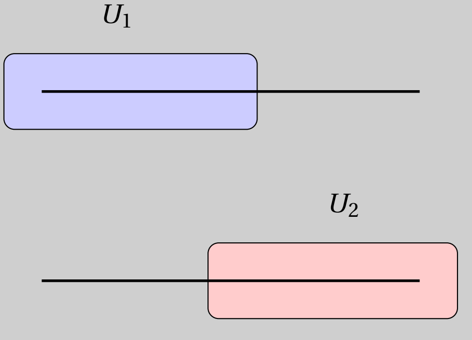

Show that the map $i^* : C^1(X) \to C^2(X)$ is surjective, and so

$$\dim(C^2(X)) - \dim(\text{ Im }(i^*)) = \dim(\ker(0_{C^2(X)\to 0})) - \dim(\text{ Im }(i^*)) = 0.$$

Exercise P31

Let X be the line segment, covered with n open intervals which overlap as in the diagram below:

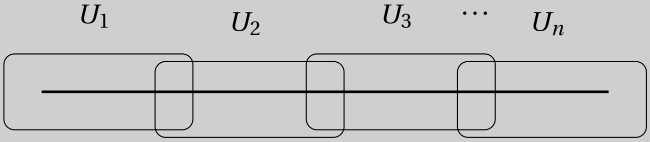

Define a sequence

$$C^0(X) \xrightarrow{j^*} C^1(X) \xrightarrow{i^*} C^2(X)$$

where $C^1(X)$ is based on the connected components of the U_i, and the $C^2(X)$ is based on the intersections $U_i \cap U_{i+1}$. Again, show that

$$\dim(C^2(X)) - \dim(\text{ Im }(i^*)) = \dim(\ker(0_{C^2(X)\to 0})) - \dim(\text{ Im }(i^*)) = 0.$$

Exercise P32

Let X be the circle, covered with n intervals which overlap end to end as drawn below.

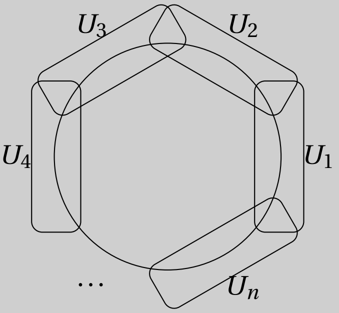

Define $C^1(X)$ and $C^2(X)$ as in the previous problem.

- Pick a basis for $C^1(X)$ and $C^2(X)$ given by functions which map a single connected component to 1, and all other components to zero. Write down the map i^* in this basis.
- Show that for this cycle,

$$\dim(C^2(X)) - \dim(\text{ Im }(i^*)) = \dim(\ker(0_{C^2(X)\to 0})) - \dim(\text{ Im }(i^*)) = -1.$$

P33 Exercise

Cover this figure eight with sets so that

- Each set is connected
- Each pair of sets intersect in one connected component
- No three sets have common overlap.

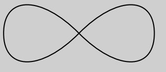

Define a sequence

$$C^0(X) \xrightarrow{j^*} C^1(X) \xrightarrow{i^*} C^2(X)$$

where $C^1(X)$ is based on the connected components of the U_i, and the $C^2(X)$ is based on intersection on the intersections $U_i \cap U_k$. Then compute

$$\dim(\text{ Im } (i^*)) - \dim(C^2(X)).$$

P34 Exercise

Let $A^\bullet$ be a chain complex, and let $B^k := H^k(A)$ be the chain complex whose cochain groups are given by the cohomology groups $H^k(A)$ and whose differential is always zero. Verify that $\pi : A^\bullet \to B^\bullet$ which sends each element of A to its cohomology class is a cochain map, and $\pi : H^k(A^\bullet) \to H^k(B^\bullet)$ is an isomorphism.

P35 Exercise

Let $X = (\Delta_X, \mathscr{S}_X)$ be a simplicial complex. A *simplicial subcomplex* is a simplicial complex $Y = (\Delta_Y, \mathscr{S}_Y)$ with $\mathscr{S}_Y \subset \mathscr{S}_X$ and

$$\sigma \in \Delta_Y \Rightarrow \sigma \in \Delta_X.$$

Show that if Y is a subcomplex of X, there is a cochain map

$$i^* : \underline{C}^\bullet(X, \mathbb{Z}_2) \to \underline{C}^\bullet(Y, \mathbb{Z}_2).$$

Exercise P36

Let $Y \subset X$ be a simplicial subcomplex. Denote the corresponding map of topological spaces $i : Y \to X$. Construct a new simplicial complex, cone(i) whose vertex set is

$$\mathscr{S}_{\text{cone}} := \mathscr{S} \cup \{x\},$$

and whose simplifies are:

$$\Delta_{\text{cone}} := \Delta_X \cup \{\sigma \cup \{x\} \mid \sigma \in \Delta_Y\}.$$

Draw a picture for cone(i) when X is an interval, and Y is the two boundary vertices of the interval. Furthermore, explain why this operation is called the cone.

Exercise P37

Let $i^* : \underline{C}^\bullet(X, \mathbb{Z}_2) \to \underline{C}^\bullet(Y, \mathbb{Z}_2)$ be the map considered above. Prove that

$$\underline{C}^\bullet(\text{cone}(i), \mathbb{Z}_2) = \text{cone}^\bullet(i^*)[-1]$$

Exercise P38

The n-disk (denoted D^n) is the simplicial complex where $\mathscr{S}_{D^k} := \{0, \dots, n\}$ and

$$\Delta_{D^n} = \{\sigma \mid \sigma \subset \mathscr{S}_{D^n}\}.$$

Let $\text{id}_{D^n} : D^n \to D^n$ be the inclusion of D^k into itself as a subcomplex. Show that

$$\text{cone}(\text{id}_{D^n}) = D^{n+1}.$$

When X is a simplicial complex, we denote by $H^i(X, \mathbb{Z}_2)$ to be the i-th cohomology group of $\underline{C}^\bullet(X, \mathbb{Z}_2)$.

Exercise P39

Use the previous characterization of D^{n+1} to compute the homology groups $\underline{H}^i(D^k)$ inductively.

The n-sphere (denoted S^n) is the simplicial complex where $\mathscr{S}_{S^n} = \{0, \ldots, n+1\}$ and

$$\Delta_{S^n} = \{\sigma \mid \sigma \subset \mathscr{S}_{S^n}, \sigma \neq \{0, \ldots, n+1\}.$$

Show that there is a map $i_{S^n} : S^n \to D^{n+1}$, and that

$$\text{cone}(i_{S^n}) = S^{n+1}.$$

P40 Exercise

Use the previous characterization of S^{n+1} to compute the cohomology groups $\underline{H}^i(S^n)$ inductively.

P41 Exercise