Math 7H	Professor: Padraic Bartlett

Lecture 6: Topology and Graphs

Week 6	*UCSB 2015*

1 Graphs and Induction

A concept that we'll use often in these classes is the idea of a **graph**!

If you haven't seen graphs before, we define them here:

Definition. A **graph** G with n vertices and m edges consists of the following two objects:

1. a set $V = \{v_1, \ldots v_n\}$, the members of which we call G's **vertices**, and
2. a set $E = \{e_1, \ldots e_m\}$, the members of which we call G's **edges**, where each edge e_i is an unordered pair of distinct elements in V, and no unordered pair is repeated. For a given edge $e = \{v, w\}$, we will often refer to the two vertices v, w contained by e as its endpoints.

Example. The following pair (V, E) defines a graph G on five vertices and five edges:

- $V = \{1, 2, 3, 4, 5\}$,
- $E = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 1\}\}$.

Something mathematicians like to do to quickly represent graphs is **draw** them, which we can do by taking each vertex and assigning it a point in the plane, and taking each edge and drawing a curve between the two vertices represented by that edge. For example, one way to draw our graph G is the following:

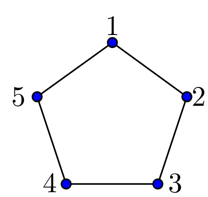

We could also draw our graph like this:

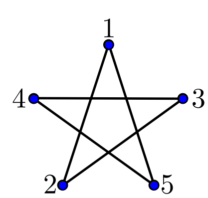

In general, all we care about for our graphs is their vertices and their edges; we don't usually care about how they are drawn, so long as they consist of the same vertices connected via the same edges. Also, we usually will not care about how we "label" the vertices of a graph: i.e. we will usually skip the labelings on our graphs, and just draw them as vertices connected by edges.

Some graphs get special names:

Definition. The **cycle graph** on n vertices, C_n, is the graph on the vertex set $\{v_1, v_2, \ldots v_n\}$ with edge set $E(C_n) = \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots \{v_{n-1}, v_n\}, \{v_n, v_1\}\}$. The cycle graphs C_n can be drawn as n-gons, as depicted below:

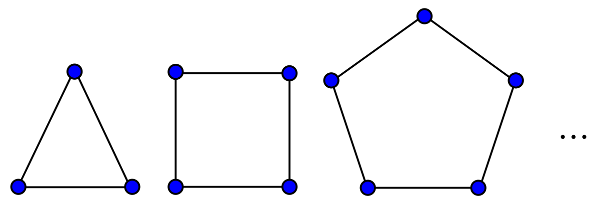

Definition. The **path graph** on n vertices, P_n, is the graph on the vertex set $\{v_1, v_2, \ldots v_n\}$ with edge set $E(C_n) = \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots \{v_{n-1}, v_n\}\}$. The path graphs P_n can be drawn as paths of length n, as depicted below:

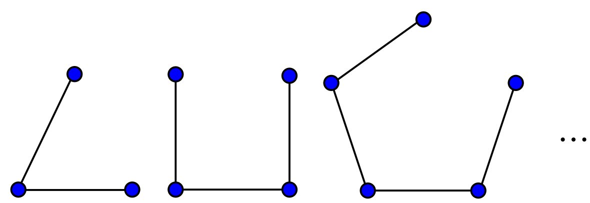

Definition. The **complete graph** K_n**.** The complete graph on n vertices, K_n, is the simple graph on the vertex set $\{v_1, v_2, \ldots v_n\}$ that has every possible edge: in other words, $E(K_n) = \{\{v_i, v_j\} : i \neq j\}$. We draw several of these graphs below:

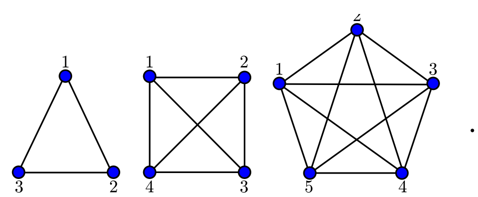

Every vertex in a K_n has degree $n-1$, as it has an edge connecting it to each of the other $n-1$ vertices; as well, a K_n has $n(n-1)/2$ edges in total in it, by the degree-sum formula. (Explicitly: every vertex has degree $n-1$ and there are n vertices, therefore the sum of the degrees of K_n's vertices is $n(n-1)$. We've shown that this quantity is twice the number of edges in the graph; dividing by 2 then tells us that the number of edges in K_n is $n(n-1)/2$, as claimed.)

Definition. The **complete bipartite graph** $K_{n,m}$**.** The complete bipartite graph on $n+m$ vertices with part sizes n and m, $K_{n,m}$, is the following graph:

- $V(K_{n,m}) = \{v_1, v_2, \dots v_n, w_1, w_2, \dots w_m\}$.
- $E(K_{n,m})$ consists of all of the edges between the n-part and the m-part; in other words, $E(K_{n,m}) = \{(v_i, w_j) : 1 \leq i \leq n, 1 \leq j \leq m\}$.

The vertices v_i all have degree m, as they have precisely m edges leaving them (one to every vertex w_j); similarly, the vertices w_j all have degree n. By either the degree-sum formula or just counting, we can see that there are nm edges in $K_{n,m}$.

Definition. Given a graph G and another graph H, we say that H is a **subgraph** of G if and only if $V(H) \subset V(G)$ and $E(H) \subset E(G)$.

Definition. Given a graph G, we call G **connected** if for any two vertices $x, y \in V(G)$, there is a path that starts at x and ends at y in our graph G.

Definition. If a graph G has no subgraphs that are cycle graphs, we call G **acyclic**. A **tree** T is a graph that's both connected and acyclic. In a tree, a **leaf** is a vertex whose degree is 1.

Example. The following graph is a tree:

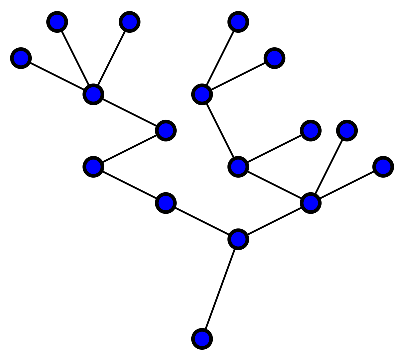

2 The Four-Color Theorem and Euler Characteristic

Graph theory got its start in 1736, when Euler studied the Seven Bridges of Königsberg problem. However, I claim that it first blossomed in earnest in 1852 when Guthrie came up with the Four-Color Problem.

Theorem. Take any map, which for our purposes is a way to partition the plane $\mathbb{R}^2$ into a collection of connected regions $R_1, \dots R_n$ with continuous boundaries. There is some way to assign each region R_i to a color in the set $\{R, G, B, Y\}$, such that if two regions R_i, R_j are "touching" (i.e. they share some nonzero length of boundary between them,) then those two regions must receive different colors.

We're going to prove this theorem… in our next class! To do this, however, we need some topological fundamentals: namely, the idea of the Euler characteristic! We introduce this concept by first giving a few definitions:

Definition. We say that a graph G is **planar** if we can draw it in the plane so that none of its edges intersect.

Sometimes, it will help to think of planarity in the following way:

Definition. We call a connected graph G **planar** if we can draw it on the sphere S^2 in the following fashion:

- Each vertex of G is represented by a point on the sphere.
- Each edge in G is represented by a continuous path drawn on the sphere connecting the points corresponding to its vertices.
- These paths do not intersect each other, except for the trivial situation where two paths share a common endpoint.

We call such a drawing a **planar embedding** of G on the sphere.

It is not hard to see that this definition is equivalent to our earlier definition of planarity. Simply use the stereographic projection map (drawn below) to translate any graph on the plane to a graph on the sphere:

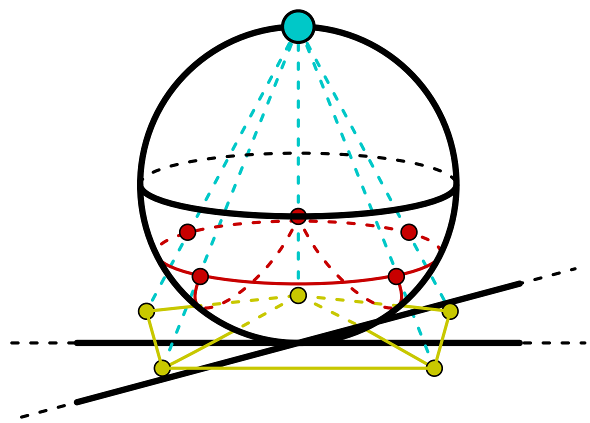

By drawing lines from the "north pole" (0,0,1) through points either in the xy-plane or on the surface of the sphere, we can translate graphs drawn on the sphere (in red) to graphs drawn in the plane (in yellow.)

Definition. For any connected planar graph G, we can define a **face** of G to be a connected region of $\mathbb{R}$ whose boundary is given by the edges of G.

For example, the following graph has four faces, as labeled:

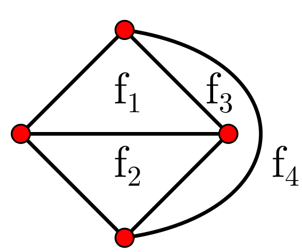

Notice that we always have the "outside" face in these drawings, which can be easy to forget about when drawing our graphs on the plane. This is one reason why I like to think about these graphs as drawn on the sphere; in this setting, there is no "outside" face, as all of the faces are equally natural to work with.

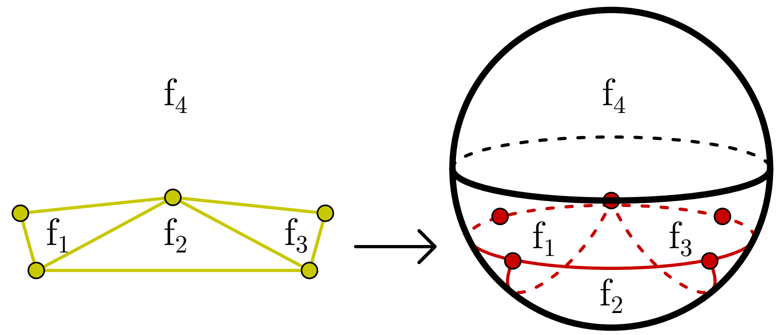

This observation has a nice accompanying lemma:

Lemma. Take any connected planar graph G, and any face F of G. Then G can be drawn on the plane in such a way that F is the outside face of G.

Proof. Take a planar embedding of G on the unit sphere. Rotate this "drawn-upon" sphere so that the face F contains the north pole $(0, 0, 1)$ of the sphere. Now, perform stereographic projection to create a planar embedding of G in $\mathbb{R}^2$. By construction, the face F is now the outside face, which proves our claim. $\square$

It bears noting that not all graphs are planar:

Proposition. The graph K_5 is not planar.

Proof. Draw a 5-cycle on the sphere. If the edges of this 5-cycle do not intersect each other, then the resulting pentagon partitions the sphere into two parts, each part of which is bounded by this pentagon. Take either one of these parts; notice that within that part, we can draw at most two nonintersecting edges connecting nonadjacent vertices in that part. Consequently, it is impossible to draw the additional 5 edges required to create K_5 without using overlapping edges. Therefore it is impossible to find a planar embedding of K_5 on the sphere, as claimed! $\square$

Planar graphs have many particularly beautiful properties! One of them is the Euler characteristic:

Theorem. (Euler characteristic.) Take any connected graph that has been drawn in $\mathbb{R}^2$ as a planar graph. Then, if V is the number of vertices, E is the number of edges, and F is the number of faces in this graph, we have the following relation:

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

Proof. We will actually prove a stronger claim: we will show that any planar **multigraph** (a graph, but where we allow multiple edges between vertices, and also edges that start and end at the same vertex) satisfies the $V - E + F = 2$ formula. For the rest of this proof, we will assume that graph and multigraph are synonymous; once we are done with this proof, though, we will stop assuming this.

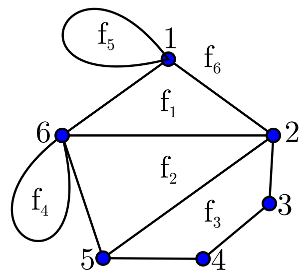

We proceed by induction on the number of vertices. Suppose that $V = 1$. Then our graph looks like something of the following form:

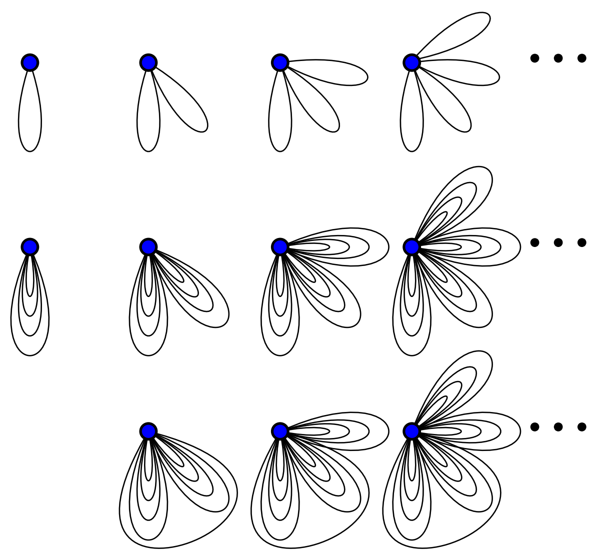

I claim that $V - E + F = 2$ for any of these graphs, and prove it by a second induction on the number of edges. For a zero-edge graph, this is easy; we have one vertex, no edges and one face, we have $V - E + F = 1 - 0 + 1 = 2$. Now, assume via induction that every one-vertex multigraph on n edges has $V - E + F = 2$. Take any graph on one vertex with $n+1$ edges. Pick one of these edges, and look at it.

I claim that this edge borders exactly two faces. To see why, take any edge, and assign an orientation to it (i.e. if our edge is $\{x, y\}$, then orient the edge so that we travel from x to y.) If you do this, then our edge has two "sides," the left- and right-hand sides, if we travel along it via this orientation.

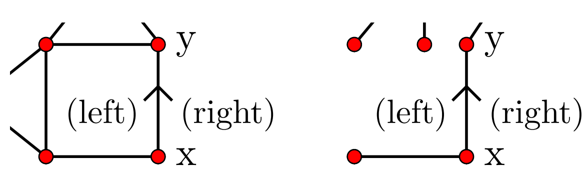

There are two possibilities, as drawn above: either the left- and right-hand sides are different, or they are the same. This tells us that our edge either borders one or two faces! To see that we have exactly two, we now recall that our edge (because our graph has exactly one vertex) must start and end at the same vertex. In other words, it is a **closed** loop: i.e.

its outside is different from its inside! In other words, our left- and right-hand sides are different, and our edge separates two distinct faces.

Therefore, deleting this edge does the following things to the graph: it decreases our edge count by 1, and also decreases our face count by 1 (as we merge two faces when we delete this edge.) In other words, deleting this edge does not change $V - E + F$! But by induction we know that $V - E + F = 2$ for all 1-vertex graphs on n edges, which is what we get if we delete this edge from a $n + 1$-edge graph. So we're done!

This settles our base case for our larger induction on V, the number of vertices. We now go to the second phase of an inductive proof: we show how to reduce larger cases to smaller cases!

To do this, consider the following operation, called **edge contraction**. Take any edge with two distinct endpoints. Delete this edge, and combine its two endpoints together: this gives us a new graph! We draw examples of this process below: we start with a graph on six vertices, and contract one by one the edges labeled in red at each step.

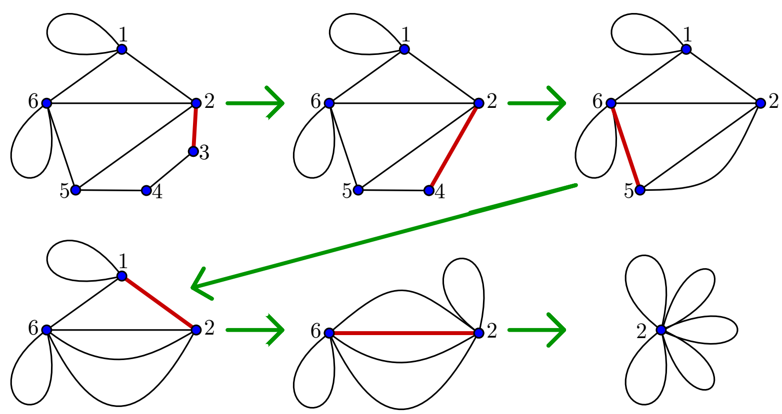

Contracting an edge decreases the number of vertices by 1 at each step, as it "squishes together" two adjacent vertices into one vertex. It also decreases the number of edges by 1 at each step, as we are contracting an edge to a point! Finally, it never changes the number of faces; if two faces were distinct before this process happens, they stay distinct, as we're not making any cuts in any of our boundaries (and instead are just shrinking them partially a bit!)

But this means that $V - E + F$ is still constant! Therefore, by induction, if $V - E + F$ holds for every n-vertex multigraph, it holds for any $n + 1$-vertex multigraph by just contracting an edge! This finishes our induction, and thus our proof. $\square$