Spectral Graph Theory Lecture 15

Tutte's Theorem: How to draw a graph

Daniel A. Spielman October 22, 2018

15.1 Overview

We prove Tutte's theorem [Tut63], which shows how to use spring embeddings to obtain planar drawings of 3-connected planar graphs. One begins by selecting a face, and then nailing down the positions of its vertices to the corners of a strictly convex polygon. Of course, the edges of the face should line up with the edges of the polygon. Ever other vertex goes where the springs say they should—to the center of gravity of their neighbors. Tutte proved that the result is a planar embedding of the planar graph. Here is an image of such an embedding

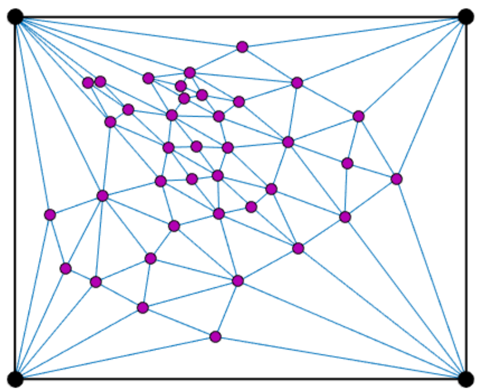

The presentation in this lecture is a based on notes given to me by Jim Geelen. I begin by recalling some standard results about planar graphs that we will assume.

15.2 3-Connected, Planar Graphs

A graph $G = (V, E)$ is k-connected if there is no set of $k-1$ vertices whose removal disconnects the graph. That is, for every $S \subset V$ with $|S| \geq |V| - (k-1)$, $G(S)$ is connected. In a classical graph theory course, one usually spends a lot of time studying things like 3-connectivity.

A *planar drawing* of a graph $G = (V, E)$ consists of mapping from the vertices to the plane, $\boldsymbol{z} : V \to \mathbb{R}^2$, along with interior-disjoint curves for each edge. The curve for edge (a, b) starts at $\boldsymbol{z}(a)$, ends at $\boldsymbol{z}(b)$, never crosses itself, and its interior does not intersect the curve for any other edge. A graph is *planar* if it has a planar drawing. There can, of course, be many planar drawings of a graph.

If one removes the curves corresponding to the edges in a planar drawing, one divides the plane into connected regions called *faces*. In a 3-connected planar graph, the sets of vertices and edges that border each face are the same in every planar drawing. There are planar graphs that are not 3-connected, like those in Figures 15.2 and 15.2, in which different planar drawings result in combinatorially different faces. We will only consider 3-connected planar graphs.

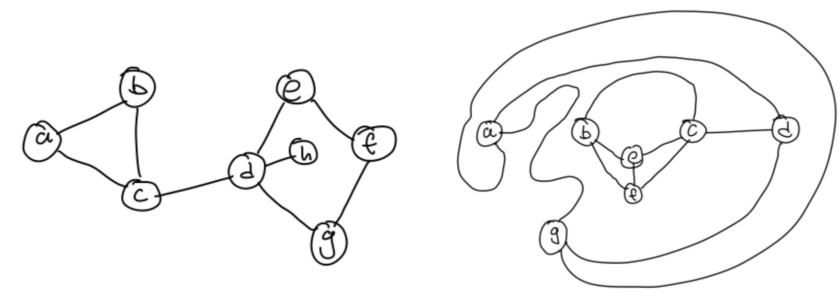

Figure 15.1: Planar graphs that are merely one-connected. Edge (c, d) appears twice on a face in each of them.

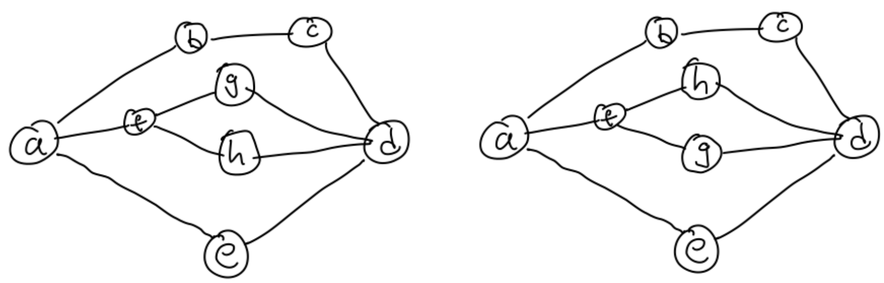

Figure 15.2: Two different planar drawings of a planar graph that is merely two-connected. Vertices g and h have switched positions, and thus appear in different faces in each drawing.

We state a few properties of 3-connected planar graphs that we will use. We will not prove these properties, as we are more concerned with algebra and these properly belong in a class on combinatorial graph theory.

Claim 15.2.1. *Let $G = (V, E)$ be a planar graph. Then, there exists a set of faces F, each of which corresponds to a cycle in G, so that no vertex appears twice in a face, no edge appears twice in a face, and every edge appears in exactly two faces.*

We call the face on the outside of the drawing the *outside face*. The edges that lie along the outside face are the *boundary edges*.

Another standard fact about planar graphs is that they remain planar under edge contractions. Contracting an edge (a, b) creates a new graph in which a and b become the same vertex, and all edges that went from other vertices to a or b now go to the new vertex. Contractions also preserve

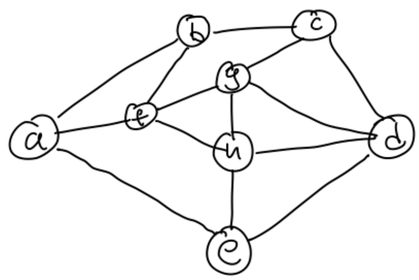

Figure 15.3: 3-connected planar graphs. Some faces of the graph on the left are abf, fgh, and $afhe$. The outer face is $abcde$. The graph on the right is obtained by contracting edge (g, h).

3-connectivity. Figure 15.2 depicts a 3-connected planar graph and the result of contracting an edge.

A graph $H = (W, F)$ is a minor of a graph $G = (V, E)$ if H can be obtained from G by contracting some edges and possibly deleting other edges and vertices. This means that each vertex in W corresponds to a connected subset of vertices in G, and that there is an edge between two vertices in W precisely when there is some edge between the two corresponding subsets. This leads to Kuratowski's Theorem [Kur30], one of the most useful characterizations of planar graphs.

Theorem 15.2.2. *A graph G is planar if and only if it does not have a minor isomorphic to the complete graph on 5 vertices, K_5, or the bipartite complete graph between two sets of 3 vertices, $K_{3,3}$.*

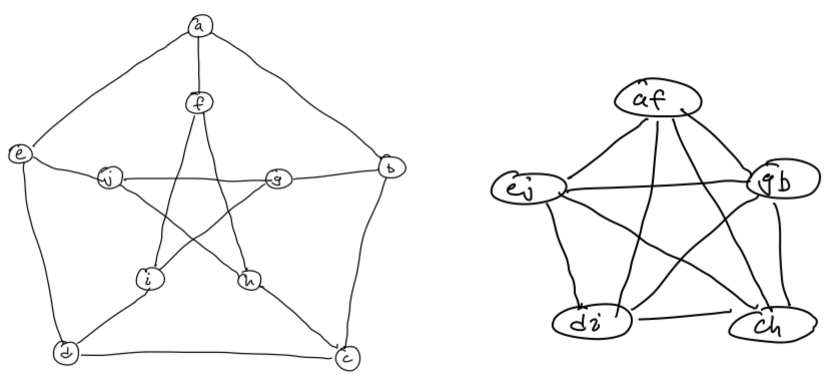

Figure 15.4: The Peterson graph appears on the left. On the right is a minor of the Peterson graph that is isomorphic to K_5, proving that the Peterson graph is not planar.

We will use one other important fact about planar graphs, whose utility in this context was observed by Jim Geelen.

Lemma 15.2.3. *Let (a, b) be an edge of a 3-connected planar graph and let S_1 and S_2 be the sets of vertices on the two faces containing (a, b). Let P be a path in G that starts at a vertex of $S_1 - \{a, b\}$,* *ends at a vertex of $S_2 - \{a, b\}$, and that does not intersect a or b. Then, every path in G from a to b either intersects a vertex of P or the edge (a, b).*

Proof. Let s_1 and s_2 be the vertices at the ends of the path P. Consider a planar drawing of G and ¡the closed curve in the plane that follows the path P from s_1 to s_2, and then connects s_1 to s_2 by moving inside the faces S_1 and S_2, where the path only intersects the curve for edge (a, b). This curve separates vertex a from vertex b. Thus, every path in G that connects a to b must intersect this curve. This means that it must either consist of just edge (a, b), or it must intersect a vertex of P. See Figure 15.2. □

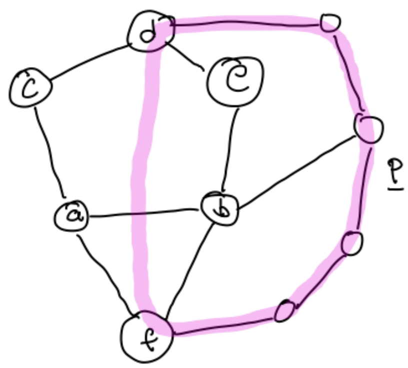

Figure 15.5: A depiction of Lemma 15.2.3. $S_1 = abcde$, $S_2 = abf$, and the path P starts at d, ends at f, and contains the other unlabeled vertices.

15.3 Strictly Convex Polygons

This is a good time to remind you what exactly a convex polygon is. A subset $C \subseteq \mathbb{R}^2$ is convex if for every two points $\boldsymbol{x}$ and $\boldsymbol{y}$ in C, the line segment between $\boldsymbol{x}$ and $\boldsymbol{y}$ is also in C. A convex polygon is a convex region of $\mathbb{R}^2$ whose boundary is comprised of a finite number of straight lines. It is *strictly* convex if in addition the angle at every corner is less than π. We will always assume that the corners of a strictly convex polygon are distinct. Two corners form an edge of the polygon if the interior of the polygon is entirely on one side of the line through those corners. This leads to another definition of a strictly convex polygon: a convex polygon is strictly convex if for every edge, all of the corners of the polygon other than those two defining the edge lie entirely on one side of the polygon. In particular, none of the other corners lie on the line.

Definition 15.3.1. *Let $G = (V, E)$ be a 3-connected planar graph. We say that $\boldsymbol{z} : V \to \mathbb{R}^2$ is a Tutte embedding if*

a. *There is a face F of G such that $\boldsymbol{z}$ maps the vertices of F to the corners of a strictly convex polygon so that every edge of the face joins consecutive corners of the polygon;*

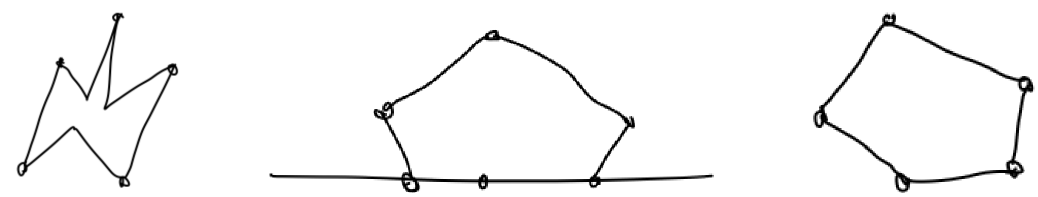

(a) A polygon (b) A convex polygon (c) A strictly convex polygon

Figure 15.6: Polygons

b. Every vertex not in F lies at the center of gravity of its neighbors.

We will prove Tutte's theorem by proving that every face of G is embedded as a strictly convex polygon. In fact, we will not use the fact that every non-boundary vertex is exactly the average of its neighbors. We will only use the fact that every non-boundary vertex is inside the convex hull of its neighbors. This corresponds to allowing arbitrary spring constants in the embedding.

Theorem 15.3.2. *Let $G = (V, E)$ be a 3-connected planar graph, and let $\boldsymbol{z}$ be a Tutte embedding of G. If we represent every edge of G as the straight line between the embedding of its endpoints, then we obtain a planar drawing of G.*

Note that if the graph were not 3-connected, then the embedding could be rather degenerate. If there are two vertices a and b whose removal disconnects the graph into two components, then all of the vertices in one of those components will embed on the line segment from a to b.

Henceforth, G will always be a 3-connected planar graph and $\boldsymbol{z}$ will always be a Tutte embedding.

15.4 Possible Degeneracies

The proof of Theorem 15.3.2 will be easy once we rule out certain degeneracies. There are two types of degeneracies that we must show can not happen. The most obvious is that we can not have $\boldsymbol{z}(a) = \boldsymbol{z}(b)$ for any edge (a, b). The fact that this degeneracy can not happen will be a consequence of Lemma 15.5.1.

The other type of degeneracy is when there is a vertex a such that all of its neighbors lie on one line in $\mathbb{R}^2$. We will rule out such degeneracies in this section.

We first observe two simple consequences of the fact that every vertex must lie at the average of its neighbors.

Claim 15.4.1. *Let a be a vertex and let ℓ be any line in $\mathbb{R}^2$ through $\boldsymbol{z}(a)$. If a has a neighbor that lines on one side of ℓ, then it has a neighbor that lies on the other.*

Claim 15.4.2. *All vertices not in F must lie strictly inside the convex hull of the polygon of which the vertices in F are the corners.*

Proof. For every vertex a not in F, we can show that the position of a is a weighted average of the positions of vertices in F by eliminating every vertex not in $F \cup \{a\}$. As we learned in Lecture 13, this results in a graph in which all the neighbors of a are in F, and thus the position of a is some weighted average of the position of the vertices in F. As the graph is 3-connected, we can show that this average must assign nonzero weights to at least 3 of the vertices in F. $\square$

Note that it is also possible to prove Claim 15.4.2 by showing that one could reduce the potential energy by moving vertices inside the polygon. See Claim 8.8.1 from my lecture notes from 2015.

Lemma 15.4.3. *Let H be a halfspace in $\mathbb{R}^2$ (that is, everything on one side of some line). Then the subgraph of G induced on the vertices a such that $\boldsymbol{z}(a) \in H$ is connected.*

Proof. Let $\boldsymbol{t}$ be a vector so that we can write the line ℓ in the form $\boldsymbol{t}^T\boldsymbol{x} = \mu$, with the halfspace consisting of those points $\boldsymbol{x}$ for which $\boldsymbol{t}^T\boldsymbol{x} \geq \mu$. Let a be a vertex such that $\boldsymbol{z}(a) \in H$ and let b be a vertex that maximizes $\boldsymbol{t}^T\boldsymbol{z}(b)$. So, $\boldsymbol{z}(b)$ is as far from the line defining the halfspace as possible. By Claim 15.4.2, b must be on the outside face, F.

For every vertex c, define $t(c) = \boldsymbol{t}^T\boldsymbol{z}(c)$. We will see that there is a path in G from a to b along which the function t never decreases, and thus all the vertices along the path lie in the halfspace. We first consider the case in which $t(a) = t(b)$. In this case, we also know that $a \in F$. As the vertices in F embed to a strictly convex polygon, this implies that (a, b) is an edge of that polygon, and thus the path from a to b.

If $t(a) < t(b)$, it suffices to show that there is a path from a to some other vertex c for which $t(c) > t(a)$ and along which t never decreases: we can then proceed from c to obtain a path to b. Let U be the set of all vertices u reachable from a for which $t(u) = t(a)$. As the graph is connected, there must be a vertex $u \in U$ that has a neighbor $c \notin U$. By Claim 15.4.1 u must have a neighbor c for which $t(c) > t(u)$. Thus, the a path from a through U to c suffices. $\square$

Lemma 15.4.4. *No vertex is colinear with all of its neighbors.*

Proof. This is trivially true for vertices in F, as no three of them are colinear.

Assume by way of contradiction that there is a vertex a that is colinear with all of its neighbors. Let ℓ be that line, and let S^+ and S^- be all the vertices that lie above and below the line, respectively. Lemma 15.4.3 tells us that both sets S^+ and S^- are connected. Let U be the set of vertices u reachable from a and such that all of us neighbors lie on ℓ. The vertex a is in U. Let W be the set of nodes that lie on ℓ that are neighbors of vertices in U, but which themselves are not in U. As vertices in W are not in U, Claim 15.4.1 implies that each vertex in W has neighbors in both S^+ and S^-. As the graph is 3-connected, and removing the vertices in W would disconnect U from the rest of the graph, there are at least 3 vertices in W. Let w_1, w_2 and w_3 be three of the vertices in W.

We will now obtain a contradiction by showing that G has a minor isomorphic to $K_{3,3}$. The three vertices on one side are w_1, w_2, and w_3. The other three are obtained by contracting the vertex sets S^+, S^-, and U. $\square$

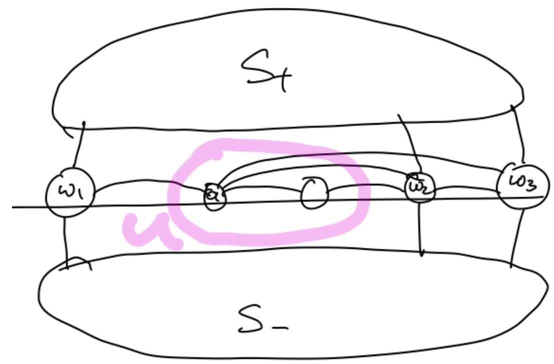

Figure 15.7: An illustration of the proof of Lemma 15.4.4.

15.5 All faces are convex

We now prove that every face of G embeds as a strictly convex polygon.

Lemma 15.5.1. *Let (a, b) be any non-boundary edge of the graph, and let ℓ be a line through $\boldsymbol{z}(a)$ and $\boldsymbol{z}(b)$ (there is probably just one). Let F_0 and F_1 be the faces that border edge (a, b) and let S_0 and S_1 be the vertices on those faces, other than a and b. Then all the vertices of S_0 and S_1 lie on opposite sides of ℓ, and none lie on ℓ.*

Note: if $\boldsymbol{z}(a) = \boldsymbol{z}(b)$, then we can find a line passing through them and one of the vertices of S_0. This leads to a contradiction, and thus rules out this type of degeneracy.

Proof. Assume by way of contradiction that the lemma is false. Without loss of generality, we may then assume that there are vertices of both S_0 and S_1 on or below the line ℓ. Let s_0 and s_1 be such vertices. By Lemma 15.4.4 and Claim 15.4.1, we know that both s_0 and s_1 have neighbors that lie strictly below the line ℓ. By Lemma 15.4.3, we know that there is a path P that connects s_0 and s_1 on which all vertices other than s_0 and s_1 lie strictly below ℓ.

On the other hand, we can similarly show that that both a and b have neighbors above the line ℓ, and that they are joined by a path that lies strictly above ℓ. Thus, this path cannot consist of the edge (a, b) and must be disjoint from P. This contradicts Lemma 15.2.3. $\square$

So, we now know that the embedding $\boldsymbol{z}$ contains no degeneracies, that every face is embedded as a strictly convex polygon, and that the two faces bordering each edge embed on opposites sides of that edge. This is all we need to know to prove Tutte's Theorem. We finish the argument in the proof below.

Proof of Theorem 15.3.2. We say that a point of the plane is *generic* if it does not lie on any $\boldsymbol{z}(a)$ for on any segment of the plane corresponding to an edge (a, b). We first prove that every generic point lies in exactly one face of G.

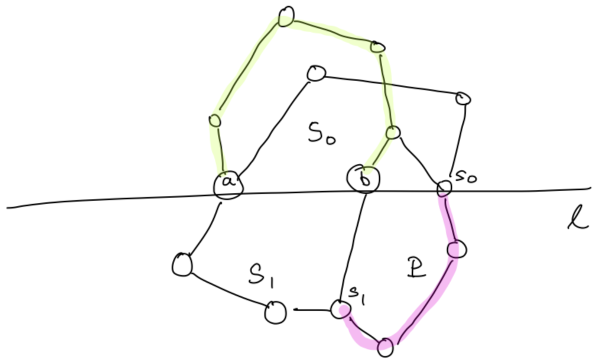

Figure 15.8: An illustration of the proof of Lemma 15.5.1.

Begin with a point that is outside the polygon on which F is drawn. Such a point lies only in the outside face. For any other generic point we can draw a curve between these points that never intersects a $\boldsymbol{z}(a)$ and never crosses the intersection of the drawings of edges. That is, it only crosses drawings of edges in their interiors. By Lemma 15.5.1, when the curve does cross such an edge it moves from one face to another. So, at no point does it ever appear in two faces.

Now, assume by way of contradiction that the drawings of two edges cross. There must be some generic point near their intersection that lies in at least two faces. This would be a contradiction. □

15.6 Notes

This is the simplest proof of Tutte's theorem that I have seen. Over the years, I have taught many versions of Tutte's proof by building on expositions by Lovász [LV99] and Geelen [Gee12], and an alternative proof of Gortler, Gotsman and Thurston [GGT06].

References

[Gee12] Jim Geelen. On how to draw a graph. https://www.math.uwaterloo.ca/∼jfgeelen/Publications/tutte. 2012.

[GGT06] Steven J Gortler, Craig Gotsman, and Dylan Thurston. Discrete one-forms on meshes and applications to 3d mesh parameterization. *Computer Aided Geometric Design*, 23(2):83–112, 2006.

[Kur30] Casimir Kuratowski. Sur le probleme des courbes gauches en topologie. *Fundamenta mathematicae*, 15(1):271–283, 1930.

[LV99] László Lovász and Katalin Vesztergombi. Geometric representations of graphs. *Paul Erdos and his Mathematics*, 1999.

[Tut63] W. T. Tutte. How to draw a graph. *Proc. London Mathematical Society*, 13:743–768, 1963.