Spectral Graph Theory Lecture 9

Tutte's Theorem

Daniel A. Spielman September 30, 2015

Disclaimer

These notes are not necessarily an accurate representation of what happened in class. The notes written before class say what I think I should say. I sometimes edit the notes after class to make them way what I wish I had said.

There may be small mistakes, so I recommend that you check any mathematically precise statement before using it in your own work.

These notes were last revised on September 30, 2015.

9.1 Overview

We prove Tutte's theorem [Tut63], which says that if one nails down the vertices of a face of a three-connected planar graph in a convex polygon, and make every other vertex be the center of gravity of its neighbors, then the result is a planar embedding of the planar graph. The presentation in this lecture is a based on Tutte's original paper [Tut63], the descriptions of it by Lovász [LV99] and Geelen [Gee12], and an alternative proof of Gortler, Gotsman and Thurston [GGT06].

I begin by recalling some standard results about planar graphs that we will assume.

9.2 3-Connected and Planar Graphs

A graph $G = (V, E)$ is 3-connected if there is no set of two vertices whose removal disconnects the graph. That is, for every $S \subset V$ with $|S| \geq |V| - 2$, $G(S)$ is connected. In a classical graph theory course, one usually spends a lot of time studying things like 3-connectivity.

Planar graphs are the special graphs that can be drawn in the plane without crossing edges. That is, every vertex v can be assigned a point $\boldsymbol{z}(v) \in \mathbb{R}^2$. Every edge (u, v) can be drawn as a curve from $\boldsymbol{z}(u)$ to $\boldsymbol{z}(v)$ so that none of the curves touch except at their endpoints. Such a drawing of a graph in the plane divides the plane into connected regions called *faces*. Each face is identified with the vertices and edges on its boarder.

In a 3-connected planar graph, the vertices and edges identified with each face are fixed. There are planar graphs that are not 3-connected, like the one shown below, in which different planar drawings result in combinatorially different faces. We will only consider 3-connected planar graphs.

In a 3-connected planar graph, the set of faces is uniquely determined and every edge boarders two different faces. Every planar drawing of such a graph will have one *outside face*.

insert figure

We will require two important facts about planar graphs in this lecture. The first is Euler's formula.

Theorem 9.2.1. *Let $G = (V, E)$ be a planar graph and let F be the set of faces in a planar drawing of G. Then,*

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

Proof. Let T be a spanning tree of G. The graph T is planar, has vertex set V, $|V| - 1$ edges, and just one face (the outside face). So, the theorem is true for T. Now, consider adding the edges of E back to T one-by-one. Each time you add an edge of E, it will create a cycle that will break one face into two. As the number of edges increases by one, so does the number of faces. □

Another standard fact about planar graphs is that they remain planar under edge contractions. Contracting an edge (u, v) creates a new graph in which u and v become the same vertex, and all edges that went from other vertices to u or v now go to the new vertex. Contractions also preserve 3-connectivity.

Finally, we recall the two simplest graphs that are not 3-connected: the complete graph on 5 vertices, K_5, and the complete bipartite graph between two sets of 3 vertices, $K_{3,3}$. If you can obtain either of these graphs by contracting and deleting edges from a graph G, then G is not planar.

Claim 9.2.2. *Let $H = (U, D)$ be a graph that is not planar. If the vertices of $G = (V, E)$ may be partitioned into subsets, $S_1, \ldots, S_u$, one for each vertex of U, so that each subgraph $G(S_a)$ is connected and so that for every $(a, b) \in D$, there is an edge between S_a and S_b, then G is not planar.*

We now know everything we need to prove Tutte's theorem.

Theorem 9.2.3. *Let $G = (V, E)$ be a 3-connected planar graph, and let B be the set of vertices on a face of G. If we fix the location of the vertices in B so that the edges between them enclose a strictly convex polygon in $\mathbb{R}^2$ and set every other vertex to the average of its neighbors, then the resulting embedding of G is planar.*

We begin the proof by trying to make it obvious that it must be true. We do this by proving that it is true when certain degeneracies do not occur. We will then do some work to prove that these degeneracies in fact do not occur.

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.

9.3 Setting up Tutte's theorem

Let B be a face of G, let $\boldsymbol{f}(b) \in \mathbb{R}^2$ be the points to which each vertex $b \in B$ is fixed, and assume that these locations are arranged in a convex polygon with the edges of the face on the outside.

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. One can characterize a strictly convex polygon by a property of its boundary: a polygon is convex if and only if every line that crosses its boundary intersects the interior of at most two of its edges. I say "interior" to avoid intersections at vertices.

picture of convex, picture of not convex

We will prove Tutte's theorem by proving that every face of G is embeded as a convex polygon. We will also prove that for every vertex v of V, the edges attached to v are aranged around it in the correct order, a notion that I should now make precise.

Assume that we are given a planar embedding of a 3-connected planar graph. Let v be a vertex, and let $u_1, \dots, u_k$ be the neighbors of v. We can label these neighbors so that they appear in order around v. This means that for each i there is a face that contains the edges (u_i, v) and (v, u_{i+1}). Similarly, there is a face containing (u_k, v) and (v, u_1).

The embedding also gives an ordering to the vertices and edges arranged around each face. If f is a face and (u, v) and (v, w) are edges on the face f, then v is between u and w in this ordering.

9.4 Non-degenerate Embeddings and Potentials

I will define an embedding of (G, B) to be a function $\boldsymbol{z}$ that maps the vertices of G to $\mathbb{R}^2$ so that

- the vertices of B lie on the corners of a strictly convex polygon and the edges of G between the vertices of B lie on the boundary, and
- every vertex not in B is the weighted average of its neighbors.

I will say that an embedding is *non-degenerate* if

(a) For all $(u, v) \in E$, $\boldsymbol{z}(u) \neq \boldsymbol{z}(v)$, and

(b) For every vertex $u \notin B$ and every $t \in \mathbb{R}^2$, there is a neighbor v of u so that $t^T(\boldsymbol{z}(u) - \boldsymbol{z}(v)) > 0$.

If the first condition holds, then the second condition can only fail to hold if all of the neighbors of u are in one line. Thus, the second condition says that u is in the interior of the convex hull of its neighbors.

We would expect the embedding we obtain to be non-degenerate, and we will eventually prove that it is. Tutte's theorem will be easy to prove if we assume that the embedding is non-degenerate.

We define a potential to be a function $\boldsymbol{x}$ that is harmonic on $V - B$. Given a potential $\boldsymbol{x}$, we say that a vertex v is *extreme* if

$$\boldsymbol{x}(v) = \min_u \boldsymbol{x}(u) \quad \text{or} \quad \boldsymbol{x}(v) = \max_u \boldsymbol{x}(u).$$

We say that a potential $\boldsymbol{x}$ is *non-degenerate* if

(c) For all $u \neq v$, $\boldsymbol{x}(u) \neq \boldsymbol{x}(v)$, and

(d) For every non-extreme vertex v, v has neighbors u and w so that

$$\boldsymbol{x}(u) < \boldsymbol{x}(v) < \boldsymbol{x}(w).$$

We can obtain a non-degenerate potential from a non-degenerate embedding by randomly rotating the embedding and then projecting it onto the x-axis. Algebraically, this corresponds to choosing a unit vector $\boldsymbol{t} \in \mathbb{R}^2$ and then setting

$$\boldsymbol{x}(u) = \boldsymbol{t}^T \boldsymbol{z}(u), \text{ for all } u.$$

If $\boldsymbol{z}$ is a non-degenerate potential, then there are only a finite number of $\boldsymbol{t}$ for which the resulting potential is degenerate, but there are an infinite number to choose from. For vertices $u \notin B$, the property (d) follows from property (b) of non-degenerate embeddings. For the vertices $u \in B$ that are not extreme for $\boldsymbol{x}$ it follows from the fact that the vertices in B are arranged in a convex polygon.

9.5 The proof in the non-degenerate case

We call an ordered triple of vertices (u, v, w) a *corner* if (u, v) and (v, w) are edges on the same face. We call the vertex v the *apex* of the corner. We say that a corner is *normal* for a potential $\boldsymbol{x}$ if $\boldsymbol{x}(v)$ is between $\boldsymbol{x}(u)$ and $\boldsymbol{x}(w)$. We say that a corner is *extreme* for a potential $\boldsymbol{x}$ if either

$$\boldsymbol{x}(v) > \max(\boldsymbol{x}(u), \boldsymbol{x}(w)) \qquad \text{or} \qquad \boldsymbol{x}(v) < \min(\boldsymbol{x}(u), \boldsymbol{x}(w)).$$

If the potential $\boldsymbol{x}$ is non-degenerate, then every corner is either normal or extreme.

If a face f is embedded as a strictly convex polygon, then for every non-degenerate potential f will have exactly two extreme corners. Conversely, if the embedding of f is not convex, then there is a non-degenerate potential for which f has at least 4 extreme corners. We will prove that the embedding of every face is a strictly convex polygon by proving that for every non-degenerate potential every face has exactly 2 extreme corners.

We can similarly count the extreme and normal corners with a common apex. If the neighbors of a vertex $v \notin B$ are arranged around it in order, then exactly two of the corners will be normal with respect to a non-degeneral potential. The rest will be extreme. We will prove that for every non-degenerate potential, every non-boundary vertex is the apex of exactly two normal corners. We prove all these assertions in four steps.

For a face f let $\mathsf{normal}(f)$ and $\mathsf{extreme}(f)$ denote the number of normal and extreme corners around f. Make the same definitions for vertices v.

Lemma 9.5.1.

$$\sum_{v \in V} \mathsf{normal}(v) = \sum_{f \in F} \mathsf{normal}(f).$$

Proof. In a 3-connected planar graph, every corner (u, v, w) lies on one exactly one face f and has exactly one apex. So, every normal corner contributes to both sums once.

Note: this statement can be made true even if the graph is not 3-connected. You just need to count each corner once for each face in which it appears. □

Lemma 9.5.2. *Every face has at least two extreme corners with respect to a non-degenerate potential.*

Proof. The vertices of the face that are the maxima and minima for the potential are the apexes of extreme corners. □

Lemma 9.5.3. *Every vertex $v \notin B$ is the apex of at least two normal corners in a non-degenerate potential.*

Proof. Let $u_1, \ldots, u_k$ be the neighbors of v, in order. As $\boldsymbol{x}(v)$ is strictly between the maximum and minimum of $\boldsymbol{x}(u_i)$, there must be at least two normal corners at which v is the apex. □

Lemma 9.5.4.

$$\sum_{v \in V} \mathsf{normal}(v) + \sum_{f} \mathsf{extreme}(f) = 2\left(|V| + |F| - 2\right).$$

Proof. By Lemma 9.5.1,

$$\begin{aligned}
\sum_{v \in V} \mathsf{normal}(v) + \sum_{f} \mathsf{extreme}(f) &= \sum_{v \in V} \mathsf{normal}(v) + \sum_{f} (\mathsf{sides}(f) - \mathsf{normal}(f)) \\
&= \sum_{f} \mathsf{sides}(f) \\
&= 2\,|E|\,.
\end{aligned}$$

The lemma now follows from Theorem 9.2.1. □

Lemma 9.5.5. *For every non-degenerate potential $\boldsymbol{x}$, every face has exactly two extreme vertices and every non-extreme vertex is the apex of exactly two normal corners.*

Proof. From Lemmas 9.5.3 and 9.5.2,

$$\sum_{\text{non-extreme } v} \mathsf{normal}(v) + \sum_{f} \mathsf{extreme}(f) \geq 2(|V| - 2) + 2\,|F| = 2\,(|V| + |F| - 2)\,.$$

As the extreme vertices do not have any normal corners, Lemma 9.5.4 implies that this inequality must be tight. So, it must be the case that for every face, $\mathsf{extreme}(f) = 2$ and for every non-extreme vertex v, $\mathsf{normal}(v) = 2$. $\square$

9.6 Degeneracy

We will now show that an embedding $\boldsymbol{z}$ of a 3-connected planar graph cannot have a vertex v whose neighbors all lie in one line. This will guarantee condition (d). We will assume that condition (c) holds in our proof—that is, that all pairs of vertices connected by an edge are mapped to distinct points. We will justify that later.

We know from last lecture that every internal vertex must lie in the interior of the convex hull of B.

Assume by way of contradiction that there is an internal vertex v so that all of its neighbors lie in a line with v. Then, there is a potential $\boldsymbol{x}$ so that $\boldsymbol{x}(v) = 0$ and $\boldsymbol{x}(u) = 0$ for all neighbors u of v. You can obtain it by projecting orthogonal to the line. We will obtain a contradiction by proving that G contains a $K_{3,3}$ minor. That is, we can contract and remove vertices of G to obtain $K_{3,3}$. By Claim 9.2.2, this would contradict the planarity of G.

Let L be the set of vertices w for which $\boldsymbol{x}(w) = 0$. If $w \in L$ has a neighbor on which $\boldsymbol{x}$ is negative, then it must have a neighbor on which $\boldsymbol{x}$ is positive. This is true for the internal vertices, because $\boldsymbol{x}$ is harmonic at those vertices. You can prove it for the vertices in B by using the fact that there is some vertex in B for which $\boldsymbol{x}$ is negative and another for which it is positive, as v is internal.

Now, let S_- be the set of all vertices for which $\boldsymbol{x}$ is negative, and let S_+ be the set of all vertices for which $\boldsymbol{x}$ is positive. I claim that both S_- and S_+ are connected. For S_+, you can show this by proving that there is a path from every vertex in S_+ to each vertex at which $\boldsymbol{x}$ is maximized on which $\boldsymbol{x}$ never decreases. Note that a vertex at which $\boldsymbol{x}$ is maximized must lie on B. To show this, it suffices to prove that there is a path from every vertex in S_+ to a vertex with a higher value of $\boldsymbol{x}$, and so that $\boldsymbol{x}$ never decreases on that path. Let $u \in S_+$ and let U be the set of all vertices w reachable from u for which $\boldsymbol{x}(u) = \boldsymbol{x}(w)$. As G is connected, if $\boldsymbol{x}(u)$ is not maximal, there must be some vertex in U with a neighbor not in U. This vertex will have a neighbor for which $\boldsymbol{x}$ is larger.

Let C be the subset of vertices in L that have neighbors that are not in L. Finally, let S_0 be the component of v in $G(L - C)$. That is, S_0 is the set of vertices that you can reach from v without touching a vertex that has a neighbor that is not on L.

As S_0 is a subgraph of G, and as G is 3-connected, there must be at least 3 vertices on the boundary of S_0. These will be vertices that are in L, but which have neighbors that are not, and which are neighbors of vertices in S_0. Let w_1, w_2 and w_3 be three such vertices.

We will obtain our $K_{3,3}$ minor by letting S_0, S_+ and S_- be the subsets of vertices on one side of the graph. We then let w_1, w_2 and w_3 be the other 3. This is a $K_{3,3}$ because each w_i has a neighbor in S_0, S_+ and S_-.

9.7 The other degeneracy

Is easy, and will be finished later.

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.

[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.