Lecture 14 Planar and Plane Graphs

Planar graphs have many important applications in computer science, for example in VLSI layout. Many problems that are hard or even *NP*-complete for arbitrary graphs are much easier for planar graphs. In the next lecture we will prove a nice result due to Lipton and Tarjan in 1977 [73] which opens up planar graphs to divide-and-conquer.

In this lecture, we will define planar and plane graphs and develop some of their basic properties. Our treatment will have a more combinatorial flavor than the classical treatment [48, 14]. Edmonds, the same one who showed the greedy algorithm only works for matroids, was the first to give a combinatorial definition of graph embeddings [31].

For the purposes of this lecture and the next, we will allow graphs to have multiple edges and self-loops, but we will prohibit isolated vertices (vertices with no adjacent edges). This assumption is for technical reasons that will become clear.

14.1 Planar and Plane Graphs—Traditional Version

According to the traditional definition, a graph is *planar* if it can be embedded on the plane or sphere in such a way that no two edges cross. A *plane graph* G is a planar graph together with such an embedding.

The complete graph on five vertices K_5 and the complete bipartite graph on two sets of three vertices $K_{3,3}$ are not planar:

$K_{3,3}$ K_5

An amazing result of Kuratowski states that *any* nonplanar graph must contain a subgraph that is topologically equivalent to one of these two graphs.

Theorem 14.1 (Kuratowski) *An undirected graph is nonplanar if and only if it contains a subgraph homeomorphic to K_5 or $K_{3,3}$.*

Here "homeomorphic to" means the edges can be paths. For more on Kuratowski's Theorem, see [48, 14].

14.2 The Plane Dual—Traditional Version

The *plane dual* of a plane graph G is a graph G^* whose vertices are the *faces* of G and whose edges are in one-to-one correspondence with the edges of G. Traditionally, a *face* is defined to be a maximal connected region of $\mathcal{R}^2 - G$, the plane with all vertices and edges of the embedded G removed. The plane dual of G is obtained by placing a vertex in each face and connecting two faces adjacent to a common edge e of G with an edge of G^* that crosses e once and crosses no other edges.

Example 14.2 The following picture shows K_4 and its plane dual K_4^*, which happens to be isomorphic to K_4:

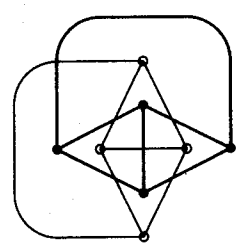

Note that any G^* is connected, and if G is connected, then G^{**} is isomorphic to G. □

14.3 Plane Graphs—Combinatorial Version

An embedding of a planar graph G on the sphere determines an orientation function θ giving a counterclockwise ordering of edges about each vertex. The map θ determines the embedding uniquely (up to rearrangement of the connected components). While we will continue to use the traditional definitions of plane graph and plane dual as intuitive aids, in computational practice it is more convenient to forget the actual embedding and work only with θ.

We will therefore start afresh and give a purely combinatorial definition of plane graphs and duals in terms of θ. This is nice because we can deal with plane graphs purely combinatorially and escape the savage world of real analysis and topology. In addition, this approach works out more nicely when G is not connected. Keep in mind that the two approaches coincide when G is connected, but diverge when G is not.

In our combinatorial formalism, an undirected graph is a tuple

$$G = (E, \theta, \bar{\ })$$

where E is a set of even cardinality, $\bar{\ }$ is an involution on E (permutation of order 2) with no fixpoints, and θ is a permutation on E. The elements of E are thought of as *directed* edges; each undirected edge is represented as a pair $e, \bar{e} \in E$ of directed edges, one in each direction. The map $\bar{\ }$ reverses direction.

The map θ is supposed to give an orientation of the edges around each vertex. But, you may well ask, where are the vertices? They are ***defined*** to be the cycles of θ. A *cycle* of the permutation θ is a minimal nonempty subset of E closed under θ. It is not to be confused with a cycle of the graph G. An edge e is considered directed out of vertex u if u is the unique cycle of θ containing e. Correspondingly, e is considered directed into vertex v if v is the unique cycle of θ containing $\bar{e}$. Thus θ cyclically permutes the edges out of any vertex. From this definition it becomes clear why isolated vertices were disallowed: you cannot have empty cycles. The tail and head functions $\mathbf{t} : E \to V$ and $\mathbf{h} : E \to V$ giving the source and sink, respectively, of each edge are defined by

$$\begin{aligned} \mathbf{t}(e) &= \{\text{the unique cycle of } \theta \text{ containing } e\} \\ \mathbf{h}(e) &= \{\text{the unique cycle of } \theta \text{ containing } \bar{e}\} \\ &= \mathbf{t}(\bar{e}) \,. \end{aligned}$$

With these definitions the tuple

$$(V, E, \mathbf{h}, \mathbf{t}, \theta, \bar{\ })$$

gives a more conventional representation of the graph G.

Definition 14.3 Define the function $\theta^* : E \to E$ by:

$$\theta^*(e) = \theta(\bar{e}) \,.$$

A *face* of G is a cycle of the permutation θ^*. The set of faces of G is denoted V^*. □

Note that this definition makes sense even for nonplanar graphs.

According to Definition 14.3, to compute $\theta^*(e)$, we first reverse the direction of e to get $\bar{e}$, then rotate about the tail of $\bar{e}$. Intuitively, for plane graphs, the operation θ^* moves an edge *clockwise* around the face to its right:

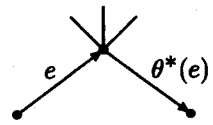

Definition 14.4 A *connected component* of G is an orbit of E under the permutation group generated by θ and $^-$. That is, it is a minimal nonempty subset of E closed under θ and $^-$. □

Definition 14.5 Let $m = \frac{1}{2}|E|$, the number of *undirected* edges of G; $n = |V|$; $n^* = |V^*|$; and c the number of connected components of G. The *characteristic* of G is the quantity

$$\chi(G) \;=\; 2c + m - n - n^* \,.$$

The graph G is said to be *plane* if $\chi(G) = 0$. □

Theorem 14.6 *A graph G is plane according to Definition 14.5 iff it is plane according to the traditional definition (i.e., if θ corresponds to the counterclockwise ordering induced by an embedding of G in the plane with no edges crossing).*

Proof. Miscellaneous Exercise 11. □

Definition 14.5 is similar to the traditional definition of the *Euler characteristic*

$$c + 1 + m - n - n^* \,.$$

Euler's Theorem states that plane graphs have Euler characteristic 0. Our Definition 14.5 and Theorem 14.6 agree with the traditional version when $c = 1$, *i.e.* when G is connected. The difference comes from the definition of the dual—in our version, disconnected graphs have more faces than in the traditional version.

14.4 The Plane Dual—Combinatorial Version

Definition 14.7 Let G be the graph

$$G \;=\; (E,\ \theta,\ ^-)\,.$$

The *dual* of G is the graph

$$G^* \;=\; (E,\ \theta^*,\ ^-)\,.$$

□

Note that this definition makes sense for graphs that are not plane.

The following theorem is immediate.

Theorem 14.8

(i) If G is plane, then so is G^*.

(ii) G^{**} and G are isomorphic (in fact they are equal).

This theorem is where our combinatorial definition wins out: (ii) is false for disconnected plane graphs under the traditional definition.

For computational purposes, a convenient representation of the undirected graph

$$G = (V, E, \mathbf{h}, \mathbf{t}, \theta, \bar{\ })$$

consists of a set of list elements, one for each vertex and one for each directed edge (element of E). The vertices are arranged in a linked list. The vertex v points to a circular list of edges e such that $\mathbf{t}(e) = v$ arranged in the order θ. The edge e points to $\mathbf{t}(e)$ and $\bar{e}$. This representation can be produced in linear time from a conventional adjacency list representation (Miscellaneous Homework 8).

14.5 Triangulation

Definition 14.9 A graph G is *triangulated* if every face of G is a triangle, *i.e.* has degree exactly three. A *triangulation* of G is a triangulated graph of which G is a subgraph. □

Theorem 14.10 *Let G be a graph such that all faces have degree at least three. A triangulation $\widehat{G}$ of G can be produced in linear time such that $\chi(\widehat{G}) = \chi(G)$; in particular, if G is plane then so is its triangulation. If G is plane, then*

$$m \leq 3n - 6 ,$$

with equality holding when G is triangulated.

Proof. We triangulate G as follows. First find all connected components and connect them in a treelike fashion by adding edges between components. Two components can be connected by adding an edge between any vertex u in one component and any vertex v in the other, and the edge can go anywhere in the ordered edge lists of u and v without changing the characteristic or the property that each face is of degree at least three. This takes linear time using DFS, and at most $c - 1 = O(m)$ new edges are added. Then traverse each face, adding chords as necessary to break up faces of degree greater than three into triangles (don't worry about multiple edges). At most $O(m)$ time is needed since each edge is traversed at most once in each direction.

Now it will suffice to prove that

$$m = 3n - 6$$

for triangulated plane graphs. Since the graph is connected, $c = 1$. Since every edge is adjacent to exactly two faces and every face is adjacent to exactly three edges, the number of adjacent face-edge pairs is $3n^* = 2m$. The result now follows from Euler's Theorem. □

Lecture 15 The Planar Separator Theorem

The Planar Separator Theorem of Lipton and Tarjan [73] says that in any undirected planar graph G there exists a small *separator* $S \subseteq V$ whose removal leaves two disjoint sets of vertices $A, B \subseteq V$ with no edge between them; moreover, each of A and B is at most a constant fraction of the size of V.

This theorem opens up planar graphs to divide-and-conquer. One can often solve a problem on a planar graph G recursively by splitting the graph into two subgraphs of size at most $\frac{2}{3}$ the size of G, recursively solving the problem on these two subgraphs, and then combining the two solutions into a solution for G. Because the sizes of the subproblems diminish geometrically, the depth of the recursion will be $O(\log n)$.

Theorem 15.1 (Planar Separator Theorem) *Let G be an undirected planar graph. There exists a partition of V into disjoint sets A, B and S such that:*

1. $|A|, |B| \leq \frac{2n}{3}$
2. $|S| \leq 4\sqrt{|V|}$
3. $(A \times B) \cap E = \emptyset$ *(S is a separator).*

Moreover, such a partition can be found in linear time.

Proof. Assume the graph is connected. (If not, perform the algorithm on the connected components and recombine the partitions into a solution for the whole graph; details omitted.)

First find a plane embedding in linear time using the algorithm of Hopcroft and Tarjan [52].

Choose an arbitrary vertex s and perform a breadth-first search (BFS) starting from s. Assign a level to each vertex, so that s is at level 0, any vertices adjacent to s are at level 1, any vertices adjacent to them that have not already been assigned a level are at level 2, and so forth. For technical reasons, we include an empty level $\ell+1$, where ℓ is the level of the last vertex encountered. Let $L(t)$ denote the set of vertices at level t.

A property of BFS traversal is that no edge ever crosses two or more levels—all edges must connect vertices in the same or consecutive levels. This means that any $L(t)$, $0 < t < \ell$, is a separator.

Let t_1 be the middle level, *i.e.* the one such that $L(t_1)$ contains vertex $n/2$ in the breadth-first numbering. The set $L(t_1)$ has some of the properties of the separator we are looking for:

$$\begin{aligned} |\bigcup_{t<t_1} L(t)| &< n/2 \\ |\bigcup_{t>t_1} L(t)| &< n/2 \,. \end{aligned}$$

So if $|L(t_1)| \leq 4\sqrt{n}$, we are done. The trouble is that $L(t_1)$ may be too large. However, there exist levels with $\sqrt{n}$ or fewer vertices on either side of t_1 and not too far away:

Lemma 15.2 *There exist levels $t_0 \leq t_1$ and $t_2 > t_1$ such that $|L(t_0)| \leq \sqrt{n}$, $|L(t_2)| \leq \sqrt{n}$, and $t_2 - t_0 \leq \sqrt{n}$.*

Proof. Let t_0 be the largest number such that $t_0 \leq t_1$ and $|L(t_0)| \leq \sqrt{n}$. Such a t_0 exists since $|L(0)| = 1$. Let t_2 be the smallest number such that $t_2 > t_1$ and $|L(t_2)| \leq \sqrt{n}$. Such a t_2 exists since $|L(\ell+1)| = 0$. Every level strictly between t_0 and t_2 contains more than $\sqrt{n}$ elements, so there must be fewer than $\sqrt{n}$ of them, otherwise there would be more than n vertices. □

Now let

$$\begin{aligned} C &= \bigcup_{t<t_0} L(t) \\ D &= \bigcup_{t_0<t<t_2} L(t) \\ E &= \bigcup_{t_2<t} L(t) \,. \end{aligned}$$

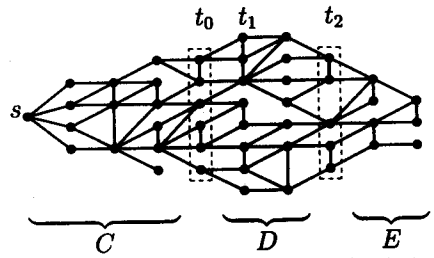

If $|D| \leq \frac{2}{3}n$ then we are done: take $S = L(t_0) \cup L(t_2)$, A the largest of C, D, E, and B the union of the other two.

We should be so lucky. If $|D| > \frac{2}{3}n$, then we at least have $|D|$ in a better shape, and this will make it easier to cut D up. The sets C and E are small: $|C|, |E| \leq \frac{n}{2}$. If we can find a $\frac{1}{3}$-$\frac{2}{3}$ separator for D with $2\sqrt{n}$ vertices or fewer, we will combine this with $L(t_0)$ and $L(t_2)$ to get a separator S of size at most $4\sqrt{n}$, combine the larger of C and E with the smaller of the two pieces of D to get A, and combine the smaller of C and E with the larger piece from D to get B. Both A and B will have no more than $\frac{2}{3}n$ vertices.

To construct a separator for D of size at most $2\sqrt{n}$, we will remove the rest of the graph, but add back the starting vertex s and connect it to everything on level $t_0 + 1$. We can do this maintaining the planarity of the graph because there were non-crossing paths back from each of those vertices to s in the original graph. Some paths may have joined on the way back to s, but they can be separated without violating planarity.

The main property of the new graph D that we will exploit is that *it has a spanning tree T of diameter at most $2\sqrt{n}$*. This is because every vertex is reachable from s by a path of length at most $\sqrt{n}$. We can construct T as follows: start with the vertices at the last level; for each such vertex, choose one edge back to the next-to-last level; repeat for the vertices on the next-to-last level, and so on all the way back to s. The $\frac{1}{3}$-$\frac{2}{3}$ separator for D will turn out to be a path in T.

We will need a useful property of plane duals. (Here we revert to the traditional definition since we need isolated vertices.)

Lemma 15.3 *Let $G = (V, E)$ be a connected plane graph with dual G^*. For any $E' \subseteq E$, the subgraph (V, E') of G has a cycle iff the subgraph $(V^*, E - E')$ of G^* is disconnected.*

Proof. ($\rightarrow$) Suppose there is a cycle in (V, E'). Choose any edge e of the cycle, and let $f, g \in V^*$ be the endpoints of e in G^*. One of f, g is inside the cycle and the other is outside. Then there is no path from f to g in $E - E'$, since no such path can cross the boundary of the cycle.

($\leftarrow$) Suppose $(V^*, E - E')$ is disconnected. Let A, B be a partition of V^* such that no edge in $E - E'$ connects A and B in G^*. Since G^* is connected, there exists at least one edge in E connecting A and B, and all such edges are in E'. These edges form a cycle in G. □

Lemma 15.4 *Let $G = (V, E)$ be a connected plane graph with dual $G^* = (V^*, E)$, and let $E' \subseteq E$. Then (V, E') is a spanning tree in G iff $(V^*, E - E')$ is a spanning tree in G^*.*

Proof. The subgraph (V, E') forms a spanning tree in G iff it is connected and has no cycles. By Lemma 15.3, this occurs iff the subgraph $(V^*, E - E')$ of G^* is connected and has no cycles, *i.e.* is a spanning tree. □

Now back to the Planar Separator Theorem. We have a plane graph D with spanning tree $T = (V, E_T)$ of diameter at most $2\sqrt{n}$. We can assume without loss of generality that D is triangulated; if not, we can triangulate it in linear time as described in the last lecture. We then construct the plane dual D^* (Miscellaneous Homework 9). This can also be done in linear time. Call the edges in $E - E_T$ *fronds*; according to Lemma 15.4, the fronds form a spanning tree $T^\dagger$ in D^*. We arbitrarily pick one face of D for the root of $T^\dagger$, say the outside face, and orient all the edges of $T^\dagger$ away from the root.

Let $e = (u, v)$ be a frond. There exists a unique path from u to v in T, which along with e forms a cycle $c(e)$.

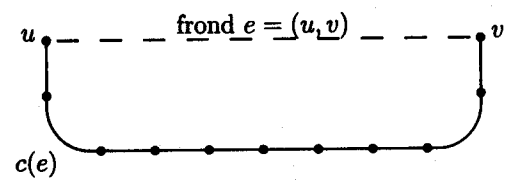

We now perform a DFS on $T^\dagger$, calculating the following information for each frond e inductively from leaves up:

- $I(e)$ = number of vertices strictly inside $c(e)$
- $|c(e)|$ = number of vertices on $c(e)$
- a list representation of $c(e)$.

There are four cases to consider, the first the base case in which e is a leaf of $T^\dagger$, and the remaining three cases induction steps:

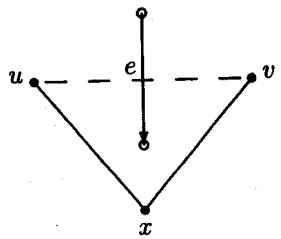

Case 1

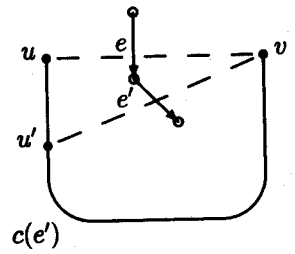

Case 2

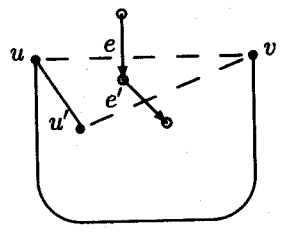

Case 3

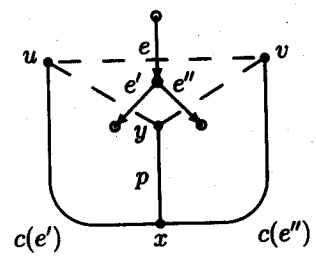

Case 4

Case 1 In this case, we are at a leaf in $T^\dagger$ (this can be detected by counting adjacencies). Then

- $I(e) = 0$
- $|c(e)| = 3$ (T is triangulated)
- $c(e) = [u, x, v]$.

Case 2 We have calculated the information for the frond $e' = (u', v)$, e is a frond in the same triangle as e', and u' is on the cycle $c(e)$; this can be detected by checking that u is not on the list $c(e')$. Then

- $I(e) = I(e')$
- $|c(e)| = |c(e')| + 1$
- $c(e) = [u] \cdot c(e')$.

Case 3 We have calculated the information for the frond $e' = (u', v)$, e is a frond in the same triangle as e', and u' is not on the cycle $c(e)$; this can be detected by checking that u is on the list $c(e')$. Then

- $I(e) = I(e') + 1$
- $|c(e)| = |c(e')| - 1$
- $c(e)$ is $c(e')$ with u' cut off the front, *i.e.* $c(e') = [u'] \cdot c(e)$.

Case 4 We have calculated the information for the fronds $e' = (u, y)$ and $e'' = (y, v)$, and e is a frond in the same triangle as both e', e''. Let p be the path common to $c(e')$ and $c(e'')$ and let x be the other endpoint of p besides y.

- $I(e) = I(e') + I(e'') + |p| - 1$ (all vertices of p except x are inside $c(e)$)
- $|c(e)| = |c(e')| + |c(e'')| - 2|p| + 1$
- $c(e) = c' \cdot [x] \cdot c''$, where c' is $c(e')$ with p removed and c'' is $c(e'')$ with p removed.

We can compute $|p|$ and construct a list representation of $c(e)$ by scanning $c(e')$ and $c(e'')$ starting at y until we encounter the last common vertex, which is x. This does not destroy the linear time complexity, since we do this for the edges on p only once.

It remains to prove that there exists a frond e such that

$$\begin{aligned} I(e) &\leq \frac{2n}{3} \\ n - (I(e) + |c(e)|) &\leq \frac{2n}{3} \,. \end{aligned}$$

Then we can just take $c(e)$ as the separator, the vertices inside $c(e)$ as A, and the vertices outside $c(e)$ as B.

Take the first frond e encountered on the way out from the leaves of $T^\dagger$ to the root such that $I(e) + |c(e)| \geq \frac{n}{3}$. Then the set of vertices outside of $c(e)$ is of cardinality $n - (I(e) + |c(e)|) \leq \frac{2n}{3}$, so it remains to show that $I(e) \leq \frac{2n}{3}$. The argument depends on the case 1 through 4 above in which e fell:

1. $I(e) = 0 \leq \frac{2n}{3}$.
2. $I(e)+|c(e)| = I(e')+|c(e')|+1$ and $I(e')+|c(e')| < \frac{n}{3}$, so $I(e)+|c(e)| \leq \frac{2n}{3}$ (for $n \geq 3$).
3. $I(e) + |c(e)| = I(e') + |c(e')|$, so e could not have been the first frond encountered such that $I(e) + |c(e)| \geq \frac{n}{3}$.
4. Both $I(e') + |c(e')| \leq \frac{n}{3}$ and $I(e'') + |c(e'')| \leq \frac{n}{3}$, so

$$\begin{aligned} & I(e) + |c(e)| \\ &= I(e') + I(e'') + |p| - 1 + |c(e')| + |c(e'')| - 2|p| + 1 \\ &= I(e') + I(e'') + |c(e')| + |c(e'')| - |p| \\ &\leq \frac{2n}{3} - |p| \\ &\leq \frac{2n}{3} \,. \end{aligned}$$

This completes the proof of the Planar Separator Theorem. □

Here is the entire algorithm:

Algorithm 15.5

1. Embed G in the plane using Hopcroft/Tarjan.
2. Do BFS on G, assigning level numbers.
3. Find t_0 and t_2 such that $|L(t_0)| \leq \sqrt{n}$, $|L(t_2)| \leq \sqrt{n}$, and $t_2 - t_0 \leq \sqrt{n}$. Divide the graph into C, D, E. If $|D| \leq \frac{2n}{3}$, we are done.
4. Otherwise, construct the spanning tree T of D of diameter at most $2\sqrt{n}$.
5. Triangulate if necessary.
6. Construct the plane dual D^* and spanning tree $T^\dagger$.
7. Do DFS on $T^\dagger$ to compute $I, |c|, c$.
8. Find the frond e such that $c(e)$ gives a $\frac{2}{3}$-$\frac{1}{3}$ separator. Let X and Y be the two sets into which D is separated.
9. Let A be the union of the larger of X, Y and the smaller of C, E, let B be the union of the smaller of X, Y and the larger of C, E, and let the separator be the union of $c(e)$, $L(t_0)$, and $L(t_2)$.