Chapter 3

Planar Graphs

This section introduces planar graphs and some of their properties.

3.1 Definitions

A graph is *planar* if it can be drawn in the plane so that no two edges cross. Not all graphs are planar, for example, we will show soon that K_5 is not planar.

Planar drawings and faces

Given a planar drawing of a planar graph G, we have the following information:

- The clockwise ordering of edges around each vertx v. These n orderings together are called a *combinatorial embedding* of G.
- The *faces*, which are the connected pieces of the plane after the drawing is removed. More precisely, draw the graph on a piece of paper. Now cut the paper along every edge of the graph. The resulting pieces of paper each constitute one face. These pieces include the infinite piece (the piece that contains the boundary of the paper), which is called *outer-face*.
- The *degree of a face F* is the number of incident edges, and denoted $\deg(F)$. Every edge belongs to exactly two faces, although these faces are not necessarily distinct. Thus $\sum \deg(F) = 2m$, where the sum is over all faces F.
- The clockwise order of edges around vertices also gives an ordering of the edges around each face. This order is counter-clockwise except for the outer-face, where it is clockwise. See Figure 3.1.
- If the graph is connected, then the boundary of each face is one circuit in the graph. Conversely any circuit in the graph for which any two consecutive edges are clockwise consecutive at their common endpoint is the boundary of a face. Thus, for a connected graph, the faces are defined by the combinatorial embedding alone, and we don't need a planar drawing to find them.

Note that a planar graphs does not necessarily have a unique planar drawing. Figure 3.2 provides an example of different planar drawings of the same graph.

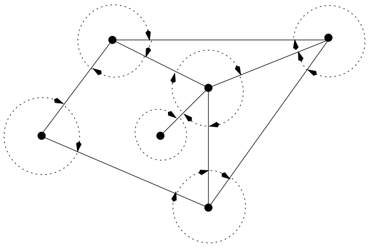

Figure 3.1: The clockwise order around each vertex defines faces, too.

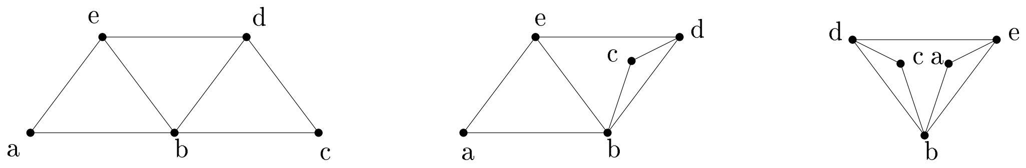

Figure 3.2: Different planar drawings of the same planar graph.

Here, the leftmost drawing is truly different from the middle drawing: the leftmost drawing has a face of degree 6 (the outer-face), while the middle drawing has no such face. On the other hand, close inspection reveals that the leftmost drawing is not so different from the rightmost drawing, because the combinatorial embedding is the same in both, and only the outer-face has changed.

It turns out that for any planar drawing, we can always change the outer-face without changing the combinatorial embedding.

Lemma 3.1 *Let G be a connected planar graph with a planar drawing Γ, and let F be one of its faces. Then there exists a planar drawing of G with the same combinatorial embedding as Γ such that F is the outer-face.*

Proof: Let p be any point inside face F in Γ, and let r be a ray emanating from p that does not cross any vertex. (Since there are only finitely many vertices, there exists such a ray.)

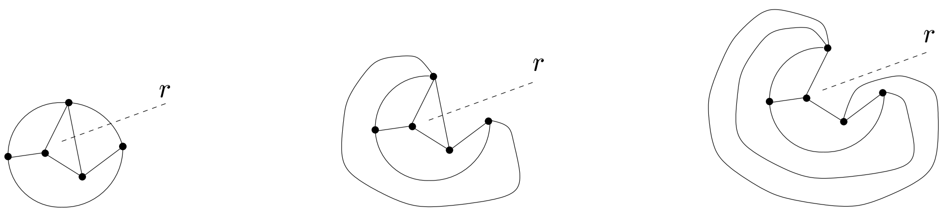

Figure 3.3: Making a face the outer-face.

The proof is now by induction on the number of edges that r crosses. If it crosses none, then F is already the infinite face, i.e., the outer-face, and we are done. If ray r crosses some edges, then it must also cross an edge $e = (u, w)$ on the outer-face of the current drawing. However it is possible to connect vertices u and w so that edge (u, w) will not cross ray r (see Figure 3.3). In this way we get planar drawing Γ' in which ray r crosses fewer edges, and are done by induction. $\square$

Plane graphs

A *plane graph* G is a planar graph with a *fixed* planar drawing, up to deformations of the plane. [If the plane is deformed the appearance of the graph might vary to the human eye, but the combinatorial embedding of the vertices and faces will not.] A plane drawing of a connected graph G is specified by a combinatorial embedding of the vertices of G and the outer-face, which we describe by two consecutive edges of the outer-face. The embedding coupled with the two consecutive edges on the outer-face will allows us to reconstruct the entire planar drawing.

A plane graph is stored using almost the same data structure as is used for graphs in general. Associated with each vertex is a list of incident edges, which in this case are sorted according to the combinatorial embedding. In addition, pointers to the two consecutive edges of the outer-face are also stored.

Dual graphs

Consider a map. The map will consist of a number of countries, who share some common borders. We can draw a graph of this map by placing a vertex at every point where three or more borders meet in a point. Edges are formed by following the borders from vertex to vertex. Each country has a capital, which we can place squarely in the middle of its country. Now across each border, draw a line joining the two capital cities of the countries adjoining the border. If we add one vertex representing the ocean, and edges from the ocean to the capital of every country adjacent to the sea, then we have in effect created the dual map.

More formally, if G is a plane connected graph, the *dual graph* G^* of a graph G is built by first placing a vertex v_F in every face F of G, and then adding an edge $e^* = (v_F, v_{F'})$ to G^* for every edge e in G that is incident to the two faces F and F'. Edge e^* is called the *dual edge* of edge e. It follows trivially that G^* is planar, and that the clockwise order of the dual edges around vertex v_F is the same as the clockwise order of the corresponding edges around face F.

Even if G is simple, it does not necessarily follow that G^* will be simple. For example, consider K_3, a triangle. The dual of K_3 will contain two vertices, but three edges going from one vertex to the other. In short, G^* can have both multiple edges and loops. If G^* has a loop e^*, then the edge e is incident to the same face twice, which means that only this edge connects two subgraphs; therefore e is a bridge in G.

The same planar graph may have different dual graphs, depending on the embedding. Consider for example Figure 3.2: In the left drawing, the dual graph has a vertex of degree 6, which is not the case in the middle drawing. However, if two drawings of a planar graph have the same combinatorial embedding, then the dual graph is the same.

The dual graph has a vertex for every face of the original graph. But what becomes of the vertices of the original graph? If v was a vertex, with incident edges $e_0 = (v, w_0), \ldots, e_{k-1} = (v, w_{k-1})$ in clockwise order around it, then we now have dual edges $e_0^*, \ldots, e_{k-1}^*$. For any $i = 0, \ldots, k-1$, the dual edges e_i^* and e_{i+1}^* have a common endpoint, namely, the vertex v_{F_i} that corresponds to the face bounded by e_i and e_{i+1} (all additions modulo k). Also, e_i^* and e_{i+1}^* are consecutive at v_{F_i}. It follows that $e_0^*, \ldots, e_{k-1}^*$ form a face. Thus, the dual graph has a face F_v for every vertex v in the original graph. Using this observation, it is very easy to show:

Lemma 3.2 *If G is a connected plane graph, then the dual of the dual graph of G is again G, or more precisely, $(G^*)^* = G$.*

The appropriate definition of what the dual graph is for a graph that is not connected is somewhat unclear. The first approach is to let G^* be the union of the dual graphs of each of the connected components of G. This has a minor flaw: the outer-face is represented by as many vertices as there are connected components, and thus we cannot really talk of "the outer-face" anymore. The second approach is to take the union, as in the first approach, but then to merge all the vertices representing the outer-faces into one vertex. Now the outer-face is represented by one vertex, but it no longer holds that $(G^*)^* = G$, because G^* and therefore $(G^*)^*$ is connected while G is not. So neither solution is perfect.

Thus, keep in mind that the dual graph is properly defined only for connected graphs, and whenever graphs may become disconnected, a precise discussion of what the dual graph is must be given.

Operations in planar graphs

What happens if we perform an operation in a planar graph G? Is the resulting graph planar?

- If we delete a vertex or an edge from G to obtain G', then simply deleting the drawing of the vertex/edge in a planar drawing of G will give a planar drawing of G'. Thus, any subgraph of a planar graph is again planar. We will call this drawing of G' the *planar drawing of* G' *induced by the planar drawing of* G.
- If we add an edge $e = (v, w)$ to G to obtain G', then G' may or may not be planar. But, if we know that v and w are on one face, then G' is planar, because we can route the new edge by going through he face.
- If we contract v and w in G to obtain G', then G' may or may not be planar. But, if v and w are on one face, then G' is planar.
- If we contract edge (v, w) in G to obtain G', then G is planar, because v and w are on one face.

Another interesting question is what happens to the dual graph G^* when we perform an operation in the graph G (referred to in the following as the *primal graph*).

Lemma 3.3 *An edge deletion in the primal graph corresponds to an edge contraction in the dual graph, as long as all graphs are connected. More precisely, if G is a connected graph and e is an edge in G that is not a bridge, then $(G-e)^* = G^* \backslash e^*$.*

Proof: When we delete e, the two faces F, F' incident to e become one face. (These were different faces because e is not a bridge.) Thus in the dual graph, vertices v_F and $v_{F'}$ become one vertex, which is a contraction. See Figure 3.4 for an example. □

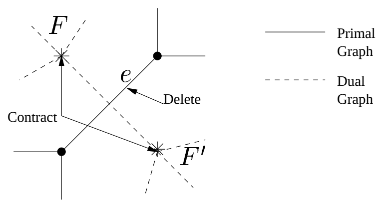

Figure 3.4: An edge deletion in the primal graph corresponds to an edge contraction in the dual graph.

3.2 Computing the dual graph

Given a plane connected graph, the dual graph can be computed in $O(n+m)$ time as follows:

Create a new graph (which will eventually be G^*), and add $2m$ vertices, two vertices $F[v,w]$ and $F[w,v]$ for each edge (v,w) of G. (These vertices represent the face "to the left" and "to the right" of (v,w).) Also, for any edge (v,w) of G, add an edge $(F[v,w].F[w,v])$ in the dual graph; this new edge is the dual edge of (v,w).

Now we want to collect all those vertices in G^* that represent the same face F. The idea to do so is to walk along the edges of F, gather all vertices representing the faces to the left of these edge, and contracting them into one vertex. The contraction part takes $O(\deg(F))$ time, because each of these vertices has only one incident edge at this moment.

But how do we find the edges of face F? We scan the list of vertices of G^*. If the vertex under consideration has been already gathered into a new face (which we presume to be recognizable, for example by keeping a flag at each vertex), we ignore this vertex and proceed with the next vertex. So assume vertex $F[v_0,v_1]$ has not been gathered into a new face yet. Then we will next gather the vertices of the face that is to the left of edge (v_0,v_1). To find them, we proceed as follows:

- Let (v_1,v_2) be the clockwise next edge after (v_0,v_1) in the incidence list of v_1.
- Let (v_2,v_3) be the clockwise next edge after (v_1,v_2) in the incidence list of v_2.
- Iterate until, for the first time, we see an edge again, so $(v_l,v_{l+1})=(v_k,v_{k+1})$, $k<l$.
- The counter-clockwise next edge after (v_l,v_{l+1}) in the incidence list of v_l is (v_{l-1},v_l) by definition. But, if $k>0$ then by $v_l=v_k$ this edge is (v_{k-1},v_k), which contradicts the choice of l. Therefore $k=0$.
- The edges $(v_0,v_1),(v_1,v_2),\ldots,(v_{l-1},v_l)$ form the boundary of a face. Contract the vertices $F[v_0,v_1],\ldots,F[v_{l-1},v_l]$.

The time spent to find the edges is thus proportional to the degree of the face that we are eventually creating. Therefore, the total time to compute the dual graph is $O(m+\sum_{\text{faces } F}\deg(F))=O(m)$.

Two remarks are in order:

- If we apply this algorithm to a graph that is not connected, then it will compute the dual where the outer-face is represented by many vertices.
- Note that nowhere in the algorithms have we used that the combinatorial embedding defined a drawing without crossing. Hence, it is possible to compute something like a dual graph for any graph where we have fixed some order of edges around each vertex. However, this dual graph has no nice geometrical interpretation unless the combinatorial embedding came from a drawing in the plane or some surface of higher genus (such as the torus).

3.3 Euler's formula

Theorem 3.4 (Euler's formula) *Let G be a connected, not necessarily simple, plane graph. Then $n-m+f=2$ where f is the number of faces.*

Proof: We will use induction on the number of faces. In the base case there is only one face. This implies that G contains no cycle, because each cycle divides the plane into an inner and an outer part which cannot belong to the same face. Graph G is therefore a forest, and because it is connected, it is a tree. A tree with n vertices has exactly $n-1$ edges, so $n - m + f = n - (n-1) + 1 = 2$.

Now assume that the graph has at least 2 faces and that the theorem holds for all graphs with fewer faces. Because there are at least two faces, there must be distinct faces that have a common boundary, say edge e is incident to two different faces F and F'. See Figure 3.5.

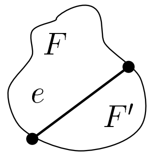

Figure 3.5: There must be an edge e incident to two different faces.

Let G' be the graph obtained by deleting the edge e from graph G. Graph G' is still connected because an edge of a planar graph is a bridge if and only if it is incident to only one face. Moreover G' has only $f-1$ faces because faces F and F' became one face. We have obtained a graph with $n' = n$ vertices, $m' = m-1$ edges and $f' = f-1$ faces. Using the induction hypothesis for G' we get $2 = n' - m' + f' = n - (m-1) + (f-1) = n - m + f$, which proves the claim. □

A number of useful results follow from Euler's formula. We start with a simple observation.

Corollary 3.5 *Let G be a planar graph. Then any planar drawing of G has the same number of faces.*

Proof: The number of faces is $m - n + 2$ by Euler's formula. Since m and n are both independent of the planar drawing, the result follows. □

The next corollary is quite important: any simple planar graph has $O(n)$ edges. The proof of this corollary is interesting in its own right, because it uses the combinatorial technique of double-counting which may be useful in other situations.

Lemma 3.6 *Every simple connected planar graph with at least 3 vertices has at most $3n-6$ edges.*

Proof: We will prove this statement by double-counting the edge-face incidences in some planar embedding of the graph. We will make a list of faces of a graph in one column and a list of edges in the other column. Then we draw a line from an edge to a face if they are incident. We draw two lines between one pair if the edge is incident to this face twice. See Figure 3.6.

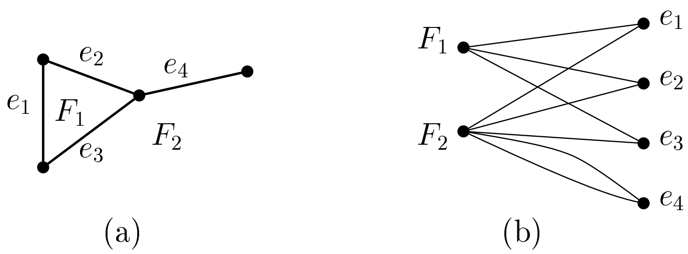

Figure 3.6: The method of double-counting.

Now let us count the number L of lines twice. We know that each edge is incident to two faces (not necessarily different) and thus there are two lines leading from each edge. Therefore $L = 2m$.

We can prove that each face is incident to at least three edges. If there were a face incident to only one edge, then it would be enclosed in this edge, i.e., the edge would be a loop, which contradicts that the graph is simple. If there were a face incident to only one edge twice, this edge would be a connected component by itself, which contradicts that the graph has at least 3 vertices and is connected. If a face were incident to exactly two different edges, then these edges would form a cycle, i.e., they are multiple edges, which contradicts that the graph is simple.

So each face is incident to at least three edges, therefore at least three lines leave from each face and $L \geq 3f$. We have counted the number of lines L twice and if we put the results together we get $3f \leq 2m$.

To finish the proof we multiply Euler's formula by 3 and plug the inequality $3f \leq 2m$ into it to obtain $6 = 3n - 3m + 3f \leq 3n - 3m + 2m = 3n - m$. □

Observe that the bound of $\leq 3n - 6$ edges does not hold for graphs with one or two vertices. On the other hand, the bound of $\leq 3n$ edges holds for any simple planar graph, even if it is not connected, which one can show easily by induction on the number of connected components. Thus, every planar simple graph has $O(n)$ edges. On the other hand, no such bound holds for planar graphs that are not simple: we could have 2 vertices, and arbitrarily many edges between them.

Corollary 3.7 *Every simple planar graph has a vertex of degree at most 5.*

Proof: We will prove this statement only for connected graphs. It then also holds for graphs that are not connected, since there must be one such vertex in each connected component. We also prove this statement only for $n \geq 3$; any graph with $n \leq 2$ vertices has a vertex of degree ≤ 1 by simplicity.

Now by Lemma 3.6 we have $6n - 12 \geq 2m = \sum_{v \in V} \deg(v)$. If each vertex in a graph had degree higher than 5, then the sum of vertex degrees would be at least $6n$ which is a contradiction. □

We can similar, but even stronger, results for planar bipartite graphs.

Lemma 3.8 *Any simple planar bipartite graph with at least 3 vertices has at most $2n - 4$ edges.*

Proof: The proof is almost identical to the proof of Lemma 3.6, except for the following observation: If a graph G is bipartite, then all cycles in G have even length, which means in particular that G has no triangle. Therefore, every face must be incident to at least four edges. Using again double-counting, we obtain $4f \leq 2m$, which in conjunction with Euler's formula yields the result. □

Lemmas 3.6 and 3.8 can be used to show that some graphs are not planar.

Corollary 3.9 *K_5 and $K_{3,3}$ are not planar graphs.*

Proof: Recall that K_5 is a complete graph with 5 vertices and $K_{3,3}$ is a complete bipartite graph with 3 vertices in each partition.

K_5 has $\binom{5}{2} = 10$ edges. However, according to Lemma 3.6 a planar graph with 5 vertices can have at most $3 \cdot 5 - 6 = 9$ edges. K_5 is therefore not planar.

$K_{3,3}$ has $3 \cdot 3 = 9$ edges. However, according to Lemma 3.8 a planar bipartite graph with 6 vertices can have at most $2 \cdot 6 - 4 = 8$ edges. $K_{3,3}$ is therefore not planar. □

In essence, these are the only two graphs that are not planar: any other graph that is not planar contains one of these graphs, or a subdivision of them.

Theorem 3.10 (Kuratowski) *Graph G is not planar if and only if it contains a subdivision of K_5 or $K_{3,3}$ as a subgraph.*

For a proof of this theorem, see for example [Gib85].

3.4 Triangulated graphs

We will finish the chapter with a brief discussion of triangulated graphs.

Definition 3.11 *A triangulated graph is a plane graph where every face is a triangle.*

In Figure 3.7 we see a graph that is not triangulated because its outer-face is not a triangle, and a graph that is triangulated but is not simple.

Figure 3.7: Examples concerning the definition of triangulated graphs.

Lemma 3.12 *Let G be a simple connected planar graph with at least 3 vertices. Then $m = 3n - 6$ if and only if G is a triangulated graph.*

Proof: Examine the proof of Lemma 3.6, where we proved that $m \leq 3n - 6$. In order to obtain equality $m = 3n - 6$, we need to have number of lines L in the double-counting method equal to $3m$ which means that each face is incident to exactly three edges. Conversely, if every face is incident to exactly three edges, then equality holds and $m = 3n - 6$. □

Also from Euler's formula, this means that every simple triangulated graphs with at least 3 vertices has exactly $2n - 4$ faces.

Finally, we need to show that triangulated graphs actually have a high connectivity.

Lemma 3.13 *Any simple triangulated graph is 3-connected.*

Proof: Assume that a graph G is simple and triangulated. The claim holds if $n \leq 2$, so assume $n \geq 3$. We will show that unless G is 3-connected, we can add an edge to G without destroying simplicity or planarity. The new graph then has $3n - 5$ edges by Lemma 3.12 and is planar and simple, which contradicts Lemma 3.6.

Assume G is not connected, say it has components $G_1, \ldots, G_k$. Then some face must contain vertices from two components, say the outer-face contains $v_1 \in G_1$ and $v_2 \in G_2$. Then we can add edge (v_1, v_2); this will not destroy planarity because both vertices are on one face, and this edge didn't exist before because the endpoints were in different connected components.

Assume next that G is connected, but not biconnected, say it has a cut-vertex v which belongs to the biconnected components $G_1, \ldots, G_k$, $k \geq 2$. Scan the edges incident to v in clockwise order.

At some point, there must be a transition point from edges in one biconnected component to edges in another biconnected component, say there are two consecutive edges (v, v_1) and (v, v_2) with $v_1 \in G_1$ and $v_2 \in G_2$. Then we can add edge (v_1, v_2); this will not destroy planarity because v_1 and v_2 are consecutive neighbors of v, and this edge didn't exist before because v_1 and v_2 belong to different biconnected components.

Finally assume that G is biconnected, but not triconnected, say it has a cutting pair $\{v, w\}$. Let $G'_1, \dots, G'_k$ be the connected components of $G - \{v, w\}$, and let G_i be the subgraph of $G - (v, w)$ induced by v, w and the vertices of G'_i. Any edge $\neq (v, w)$ belongs to one of the graphs $G_1, \dots, G_k$.

Scan the edges incident to v in clockwise order starting at (v, w) (if it exists) and at an arbitrary edge otherwise. The next edge belongs to one of the above subgraphs, say G_1. Keep scanning the edges until for the first time we encountered an edge (v, v_2) not in G_1. This edge cannot be (v, w), because we haven't seen any edge from G_2 yet, and we would be done scanning otherwise. So (v, v_2) belongs to some subgraph other than G_1, say G_2. Let (v, v_1) be the edge just before (v, v_2) in clockwise order.

Then we can add edge (v_1, v_2); this will not destroy planarity because v_1 and v_2 are consecutive neighbors of v. Also, this edge didn't exist before, because otherwise v_1 and v_2 would be in the same component of $G - \{v, w\}$, and thus in the same subgraph. □

See Figure 3.8 for an illustration of this proof.

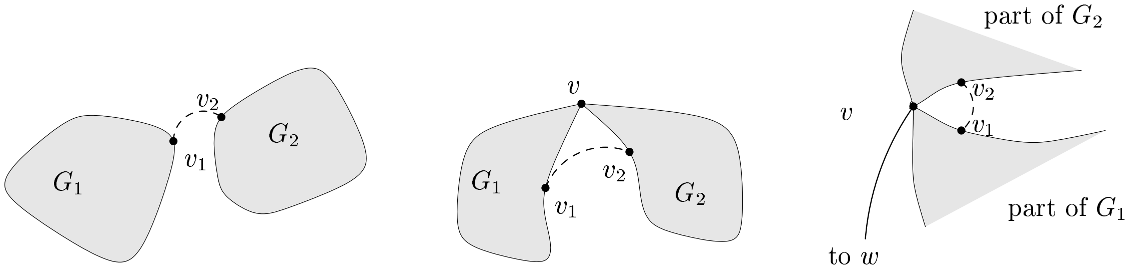

Figure 3.8: A triangulated graph is 3-connected, otherwise we can add an edge.