Chapter 2 Plane Embeddings

Graphs can be represented in a variety of ways, for instance, as an adjacency matrix or using adjacency lists. In this chapter we explore another type of representations that are quite different in nature, namely *geometric* representations of graphs. Geometric representations are appealing because they allow to visualize a graph along with a variety of its properties in a succinct manner. There are many degrees of freedom in selecting the type of geometric objects and the details of their geometry. This freedom allows to tailor the representation to meet specific goals, such as emphasizing certain structural aspects of the graph at hand or reducing the complexity of the obtained representation.

The most common type of geometric graph representation is a *drawing*, where vertices are mapped to points and edges to curves. Making such a map injective by avoiding edge crossings is desirable, both from a mathematically aesthetic point of view and for the sake of the practical readability of the drawing. Those graphs that allow such an *embedding* into the Euclidean plane are known as *planar*. Our goal in the following is to study the interplay between abstract planar graphs and their plane embeddings. Specifically, we want to answer the following questions:

- What is the combinatorial complexity of planar graphs (number of edges and faces)?
- Under which conditions are plane embeddings unique (in a certain sense)?
- How can we represent plane embeddings (in a data structure)?
- What is the geometric complexity of plane embeddings, that is, can we bound the size of the coordinates used and the complexity of the geometric objects used to represent edges?

Most definitions we use directly extend to multigraphs. But for simplicity, we use the term "graph" throughout.

2.1 Drawings, Embeddings and Planarity

A curve is a set $C \subset \mathbb{R}^2$ that is of the form $\{\gamma(t) : 0 \leq t \leq 1\}$, where $\gamma : [0, 1] \to \mathbb{R}^2$ is a continuous function. The function γ is called a *parameterization* of C. The points $\gamma(0)$

and $\gamma(1)$ are the *endpoints* of the curve. For a *closed* curve, we have $\gamma(0) = \gamma(1)$. A curve is *simple*, if it admits a parameterization γ that is injective on [0, 1]. For a closed simple curve we allow as an exception that $\gamma(0) = \gamma(1)$. The following famous theorem describes an important property of the plane. A proof can, for instance, be found in the book of Mohar and Thomassen [24].

Theorem 2.1 (Jordan). Any simple closed curve C partitions the plane into exactly two regions (connected open sets), each bounded by C.



Figure 2.1: A Jordan curve and two points in one of its faces (left); a simple closed curve that does not disconnect the torus (right).

Observe that, for instance, on the torus there are closed curves that do not disconnect the surface (and so the theorem does not hold there).

Drawings. As a first criterion for a reasonable geometric representation of a graph, we would like to have a clear separation between different vertices and also between a vertex and nonincident edges. Formally, a *drawing* of a graph G = (V, E) in the plane is a function f that assigns

- a point $f(v) \in \mathbb{R}^2$ to every vertex $v \in V$ and
- a simple curve $f(\{u,v\}): [0,1] \to \mathbb{R}^2$ with endpoints f(u) and f(v) to every edge $\{u,v\} \in E$,

such that

- (1) f is injective on V and
- (2) $f({u,v}) \cap f(V) = {f(u), f(v)}, \text{ for every edge } {u,v} \in E.$

A common point $f(e) \cap f(e')$ between two curves that represent distinct edges $e, e' \in E$ is called a *crossing* if it is not a common endpoint of e and e'.

For simplicity, when discussing a drawing of a graph G = (V, E) it is common to treat vertices and edges as geometric objects. That is, a vertex $v \in V$ is treated as the point f(v) and an edge $e \in E$ is treated as the curve f(e). For instance, the last sentence of the previous paragraph may be phrased as "A common point of two edges that is not a common endpoint is called a crossing."

Often it is convenient to make additional assumptions about the interaction of edges in a drawing. For example, in a nondegenerate drawing one may demand that no three edges share a single crossing or that every pair of distinct edges intersects in at most finitely many points.

Planar vs. plane. A graph is *planar* if it admits a drawing in the plane without crossings. Such a drawing is also called a *crossing-free* drawing or a (plane) *embedding* of the graph. A planar graph together with a particular plane embedding is called a *plane* graph. Note the distinction between "planar" and "plane": the former indicates the possibility of an embedding, whereas the latter refers to a concrete embedding (Figure 2.2).



Figure 2.2: A planar graph (left) and a plane embedding of it (right).

A geometric graph is a graph together with a drawing, in which all edges are realized as straight-line segments. Note that such a drawing is completely defined by the mapping for the vertices. A plane geometric graph is also called a *plane straight-line graph* (PSLG). In contrast, a plane graph in which the edges may form arbitrary simple curves is called a *topological plane graph*.

The *faces* of a plane graph are the maximally connected regions of the plane that do not contain any point used by the embedding (as the image of a vertex or an edge). Each embedding of a finite graph has exactly one *unbounded face*, also called *outer* or *infinite* face. Using stereographic projection, it is not hard to show that the role of the unbounded face is not as special as it may seem at first glance.

Theorem 2.2. If a graph G has a plane embedding in which some face is bounded by the cycle (v_1, \ldots, v_k) , then G also has a plane embedding in which the unbounded face is bounded by the cycle (v_1, \ldots, v_k) .

Proof. (Sketch) Take a plane embedding Γ of G and map it to the sphere using stereographic projection: Imagine \mathbb{R}^2 being the x/y-plane in \mathbb{R}^3 and place a unit sphere S such that its south pole touches the origin. We obtain a bijective continuous map between \mathbb{R}^2 and $S \setminus \{n\}$, where n is the north pole of S, as follows: A point $p \in \mathbb{R}^2$ is mapped to the point p' that is the intersection of the line through p and n with S, see Figure 2.3.

Consider the resulting embedding Γ' of G on S: The infinite face of Γ corresponds to the face of Γ' that contains the north pole n of S. Now rotate the embedding Γ' on S



Figure 2.3: Stereographic projection.

such that the desired face contains n. Mapping back to the plane using stereographic projection results in an embedding in which the desired face is the outer face. \Box

Exercise 2.3. Consider the plane graphs depicted in Figure 2.4. For both graphs give a plane embedding in which the cycle 1,2,3 bounds the outer face.



Figure 2.4: Make 1, 2, 3 bound the outer face.

Duality. Every plane graph G has a *dual* G^* , whose vertices are the faces of G and two are connected by an edge in G^* , if and only if they have a common edge in G. In general, G^* is a multigraph (may contain loops and multiple edges) and it depends on the embedding. That is, an abstract planar graph G may have several nonisomorphic duals; see Figure 2.5 for an example. If G is a connected plane graph, then $(G^*)^* = G$. We will see later in Section 2.3 that the dual of a 3-connected planar graph is unique (up to isomorphism).

The Euler Formula and its ramifications. One of the most important tools for planar graphs (and more generally, graphs embedded on a surface) is the Euler–Poincaré Formula.

Theorem 2.4 (Euler's Formula). For every connected plane graph with n vertices, e edges, and f faces, we have n - e + f = 2.



Figure 2.5: Two plane drawings G_1 and G_2 of the same abstract planar graph and their duals G_1^* and G_2^* with $G_1^* \neq G_2^*$. (To see this, for instance, count the number of vertices of degree greater than three.)

In particular, this shows that for any planar graph the number of faces is the same in every plane embedding. In other words the number of faces is an invariant of an abstract planar graph. It also follows (stated below as a corollary) that planar graphs are *sparse*, that is, they have a linear number of edges (and faces) only. So the asymptotic complexity of a planar graph is already determined by its number of vertices.

Corollary 2.5. A simple planar graph on $n \ge 3$ vertices has at most 3n - 6 edges and at most 2n - 4 faces.

Proof. Without loss of generality we may assume that G is connected. (If not, add edges between components of G until the graph is connected. The number of faces remains unchanged and the number of edges only increases.) The statement is easily checked for n = 3, where G is either a triangle or a path and, therefore, has no more than $3 \cdot 3 - 6 = 3$ edges and no more than $2 \cdot 3 - 4 = 2$ faces. So consider a simple planar graph G on $n \ge 4$ vertices. Consider a plane drawing of G and denote by E the set of edges and by F the set of faces of G. Let

 $X = \{(e, f) \in E \times F : e \text{ bounds } f\}$

denote the set of incident edge-face pairs. We count X in two different ways.

First note that each edge bounds at most two faces and so $|X| \leq 2 \cdot |E|$.

Second note that in a simple connected planar graph on four or more vertices every face is bounded by at least three vertices: Every bounded face needs at least three edges to be enclosed and if there is no cycle on the boundary of the unbounded face, then—given that G is connected—G must be a tree on four or more vertices and so it has at least three edges, all of which bound the unbounded face. Therefore $|X| \ge 3 \cdot |F|$.

Using Euler's Formula we conclude that

$$\begin{array}{rcl} 4 & = & 2n-2|\mathsf{E}|+2|\mathsf{F}| \leqslant 2n-3|\mathsf{F}|+2|\mathsf{F}| = 2n-|\mathsf{F}| \ \text{and} \\ 6 & = & 3n-3|\mathsf{E}|+3|\mathsf{F}| \leqslant 3n-3|\mathsf{E}|+2|\mathsf{E}| = 3n-|\mathsf{E}| \ , \end{array}$$

which yields the claimed bounds.

It also follows that the degree of a "typical" vertex in a planar graph is a small constant. There exist several variations of this statement, a few more of which we will encounter during this course.

Corollary 2.6. The average vertex degree in a simple planar graph is less than six.

Exercise 2.7. Prove Corollary 2.6.

Exercise 2.8. Show that neither K_5 (the complete graph on five vertices) nor $K_{3,3}$ (the complete bipartite graph where both classes have three vertices) is planar.

Exercise 2.9. Let P be a set of $n \ge 3$ points in the plane such that the distance between any pair of points is at least one. Show that there are at most 3n-6 pairs of points in P at distance exactly one.

Characterizing planarity. The classical theorems of Kuratowski and Wagner provide a characterization of planar graphs in terms of forbidden substructures. A *subdivision* of a graph G = (V, E) is a graph that is obtained from G by replacing each edge with a path.

Theorem 2.10 (Kuratowski [22, 30]). A graph is planar if and only if it does not contain a subdivision of $K_{3,3}$ or K_5 .

A minor of a graph G = (V, E) is a graph that is obtained from G using zero or more edge contractions, edge deletions, and/or vertex deletions.

Theorem 2.11 (Wagner [33]). A graph is planar if and only if it does not contain $K_{3,3}$ or K_5 as a minor.

In some sense, Wagner's Theorem is a special instance¹ of a much more general theorem.

Theorem 2.12 (Graph Minor Theorem, Robertson/Seymour [27]). Every minor-closed family of graphs can be described in terms of a finite set of forbidden minors.

Being *minor-closed* means that for every graph from the family also all of its minors belong to the family. For instance, the family of planar graphs is minor-closed because planarity is preserved under removal of edges and vertices and under edge contractions. The Graph Minor Theorem is a celebrated result that was established by Robertson and Seymour in a series of twenty papers, see also the survey by Lovász [23]. They also describe an $O(n^3)$ algorithm (with horrendous constants, though) to decide whether a graph on n vertices contains a fixed (constant-size) minor. Later, Kawarabayashi et al. [20] showed that this problem can be solved in $O(n^2)$ time. As a consequence, every minorclosed property can be decided in polynomial time.

Unfortunately, the result is nonconstructive in the sense that in general we do not know how to obtain the set of forbidden minors for a given family/property. For instance,

¹It is more than just a special instance because it also specifies the forbidden minors explicitly.

for the family of toroidal graphs (graphs that can be embedded without crossings on the torus) more than 16'000 forbidden minors are known, and we do not know how many there are in total. So while we know that there exists a quadratic time algorithm to test membership for minor-closed families, we have no idea what such an algorithm looks like in general.

Graph families other than planar graphs for which the forbidden minors are known include forests (K_3) and outerplanar graphs $(K_{2,3} \text{ and } K_4)$. A graph is *outerplanar* if it admits a plane embedding such that all vertices appear on the outer face (Figure 2.6).



Figure 2.6: An outerplanar graph (left) and a plane embedding of it in which all vertices are incident to the outer face (right).

- Exercise 2.13. (a) Give an example of a 6-connected planar graph or argue that no such graph exists.
 - (b) Give an example of a 5-connected planar graph or argue that no such graph exists.
 - (c) Give an example of a 3-connected outerplanar graph or argue that no such graph exists.

Planarity testing. For planar graphs we do not have to contend ourselves with a cubictime algorithm, as there are several approaches to solve the problem in linear time. In fact, there is quite a number of papers that describe different linear time algorithms, all of which—from a very high-level point of view—can be regarded as an annotated depthfirst-search. The first such algorithm was described by Hopcroft and Tarjan [19], while the current state-of-the-art is probably among the "path searching" method by Boyer and Myrwold [6] and the "LR-partition" method by de Fraysseix et al. [14]. Although the overall idea in all these approaches is easy to convey, there are many technical details, which make an in-depth discussion rather painful to go through.

2.2 Graph Representations

There are two standard representations for an abstract graph G = (V, E) on n = |V| vertices. For the *adjacency matrix* representation we consider the vertices to be ordered as $V = \{v_1, \ldots, v_n\}$. The adjacency matrix of an undirected graph is a symmetric $n \times n$ -matrix $A = (a_{ij})_{1 \leq i,j \leq n}$ where $a_{ij} = a_{ji} = 1$, if $\{i, j\} \in E$, and $a_{ij} = a_{ji} = 0$, otherwise.

Storing such a matrix explicitly requires $\Omega(n^2)$ space, and allows to test in constant time whether or not two given vertices are adjacent.

In an *adjacency list* representation, we store for each vertex a list of its neighbors in G. This requires only O(n + |E|) storage, which is better than for the adjacency matrix in case that $|E| = o(n^2)$. On the other hand, the adjacency test for two given vertices is not a constant-time operation, because it requires a search in one of the lists. Depending on the representation of these lists, such a search takes O(d) time (unsorted list) or $O(\log d)$ time (sorted random-access representation, such as a balanced search tree), where d is the minimum degree of the two vertices.

Both representations have their merits. The choice of which one to use (if any) typically depends on what one wants to do with the graph. When dealing with embedded graphs, however, additional information concerning the embedding is needed beyond the pure incidence structure of the graph. The next section discusses a standard data structure to represent embedded graphs.

2.2.1 The Doubly-Connected Edge List

The *doubly-connected edge list* (DCEL) is a data structure to represent a plane graph in such a way that it is easy to traverse and to manipulate. In order to avoid unnecessary complications, let us discuss only connected graphs here that contain at least two vertices. It is not hard to extend the data structure to cover all plane graphs. For simplicity we also assume that we deal with a straight-line embedding and so the geometry of edges is defined by the mapping of their endpoints already. For more general embeddings, the geometric description of edges has to be stored in addition.

The main building block of a DCEL is a list of *halfedges*. Every actual edge is represented by two halfedges going in opposite direction, and these are called *twins*, see Figure 2.7. Along the boundary of each face, halfedges are oriented counterclockwise.



Figure 2.7: A halfedge in a DCEL.

A DCEL stores a list of halfedges, a list of vertices, and a list of faces. These lists are unordered but interconnected by various pointers. A vertex v stores a pointer halfedge(v)

to an arbitrary halfedge originating from v. Every vertex also knows its coordinates, that is, the point point(v) it is mapped to in the represented embedding. A face f stores a pointer halfedge(f) to an arbitrary halfedge within the face. A halfedge h stores *five* pointers:

- a pointer target(h) to its target vertex,
- a pointer face(h) to the incident face,
- a pointer twin(h) to its twin halfedge,
- a pointer next(h) to the halfedge following h along the boundary of face(h), and
- a pointer prev(h) to the halfedge preceding h along the boundary of face(h).

A constant amount of information is stored for every vertex, (half-)edge, and face of the graph. Therefore the whole DCEL needs storage proportional to |V| + |E| + |F|, which is O(n) for a plane graph with n vertices by Corollary 2.5.

This information is sufficient for most tasks. For example, traversing all edges around a face f can be done as follows:

```
\begin{array}{l} s \leftarrow \mathsf{halfedge}(f) \\ \mathsf{h} \leftarrow s \\ \textbf{do} \\ & \text{something with } \mathsf{h} \\ & \mathsf{h} \leftarrow \mathsf{next}(\mathsf{h}) \\ \textbf{while } \mathsf{h} \neq s \end{array}
```

Exercise 2.14. Give pseudocode to traverse all edges incident to a given vertex v of a DCEL.

Exercise 2.15. Why is the previous halfedge $prev(\cdot)$ stored explicitly and the source vertex of a halfedge is not?

2.2.2 Manipulating a DCEL

In many applications, plane graphs appear not just as static objects but rather they evolve over the course of an algorithm. Therefore the data structure used to represent the graph must allow for efficient update operations to change it.

First of all, we need to be able to generate new vertices, edges, and faces, to be added to the corresponding list within the DCEL and—symmetrically—the ability to delete an existing entity. Then it should be easy to add a new vertex v to the graph within some face f. As we maintain a connected graph, we better link the new vertex to somewhere, say, to an existing vertex u. For such a connection to be possible, we require that the open line segment uv lies completely in f. Of course, two halfedges are to be added connecting u and v. But where exactly? Given that from a vertex and from a face only some arbitrary halfedge is directly accessible, it turns out convenient to use a halfedge in the interface. Let h denote the halfedge incident to f for which target(h) = u. Our operation then becomes (see also Figure 2.8)

add-vertex-at(v, h)

Precondition: the open line segment point(v)point(u), where u := target(h), lies completely in f := face(h).

Postcondition: a new vertex v has been inserted into f, connected by an edge to u.



Figure 2.8: Add a new vertex connected to an existing vertex u.

and it can be realized by manipulating a constant number of pointers as follows.

```
add-vertex-at(v, h) {
      h_1 \gets a \text{ new halfedge}
       h_2 \leftarrow a \text{ new halfedge}
      halfedge(v) \leftarrow h<sub>2</sub>
       \mathsf{twin}(\mathsf{h}_1) \leftarrow \mathsf{h}_2
       \mathsf{twin}(\mathsf{h}_2) \leftarrow \mathsf{h}_1
      target(h_1) \leftarrow v
      target(h_2) \leftarrow u
      face(h_1) \leftarrow f
      face(h_2) \leftarrow f
       next(h_1) \leftarrow h_2
       next(h_2) \leftarrow next(h)
       prev(h_1) \leftarrow h
       prev(h_2) \leftarrow h_1
       next(h) \leftarrow h_1
       prev(next(h_2)) \leftarrow h_2
}
```

Similarly, it should be possible to add an edge between two existing vertices u and v, provided the open line segment uv lies completely within a face f of the graph, see

Figure 2.9. Since such an edge insertion splits f into two faces, the operation is called *split-face*. Again we use the halfedge h that is incident to f and for which target(h) = u.



Figure 2.9: Split a face by an edge uv.

Our operation becomes then

split-face(h, v)

```
Precondition: v is incident to f := face(h) but not adjacent to u := target(h).
```

The open line segment point(v)point(u) lies completely in f.

Postcondition: f has been split by a new edge uv.

The implementation is slightly more complicated compared to add-vertex-at above, because the face f is destroyed and so we have to update the face information of all incident halfedges. In particular, this is not a constant time operation, but its time complexity is proportional to the size of f.

```
split-face(h, v) {
       f_1 \leftarrow a \text{ new face}
       f_2 \leftarrow a \text{ new face}
       h_1 \leftarrow a \text{ new halfedge}
       h_2 \leftarrow a \text{ new halfedge}
       halfedge(f_1) \leftarrow h_1
       halfedge(f_2) \leftarrow h_2
       \mathsf{twin}(\mathsf{h}_1) \leftarrow \mathsf{h}_2
       \mathsf{twin}(\mathsf{h}_2) \leftarrow \mathsf{h}_1
       target(h_1) \leftarrow v
       target(h_2) \leftarrow u
       next(h_2) \leftarrow next(h)
       prev(next(h_2)) \leftarrow h_2
       prev(h_1) \leftarrow h
       next(h) \leftarrow h_1
       i \leftarrow h_2
       loop
```

}

```
\begin{array}{l} \mbox{face}(i) \leftarrow f_2 \\ \mbox{if target}(i) = \nu \mbox{ break the loop} \\ i \leftarrow next(i) \\ \mbox{endloop} \\ next(h_1) \leftarrow next(i) \\ prev(next(h_1)) \leftarrow h_1 \\ next(i) \leftarrow h_2 \\ prev(h_2) \leftarrow i \\ i \leftarrow h_1 \\ \mbox{do} \\ face(i) \leftarrow f_1 \\ i \leftarrow next(i) \\ \mbox{until target}(i) = u \\ \mbox{delete the face f} \end{array}
```

In a similar fashion one can realize the inverse operation join-face(h) that removes the edge (represented by the halfedge) h, thereby joining the faces face(h) and face(twin(h)).

It is easy to see that every connected plane graph on at least two vertices can be constructed using the operations add-vertex-at and split-face, starting from an embedding of K_2 (two vertices connected by an edge).

Exercise 2.16. Give pseudocode for the operation join-face(h). Also specify preconditions, if needed.

Exercise 2.17. Give pseudocode for the operation split-edge(h), that splits the edge (represented by the halfedge) h into two by a new vertex w, see Figure 2.10.



Figure 2.10: Split an edge by a new vertex.

2.2.3 Graphs with Unbounded Edges

In some cases it is convenient to consider plane graphs, in which some edges are not mapped to a line segment but to an unbounded curve, such as a ray. This setting is not really much different from the one we studied before, except that one vertex is placed "at infinity". One way to think of it is in terms of *stereographic projection* (see the proof of Theorem 2.2). The further away a point in \mathbb{R}^2 is from the origin, the closer its image on the sphere S gets to the north pole n of S. But there is no way to reach n except in the limit. Therefore, we can imagine drawing the graph on S instead of in \mathbb{R}^2 and putting the "infinite vertex" at n.



Figure 2.11: A DCEL with unbounded edges. Usually, we will not show the infinite vertex and draw all edges as straight-line segments. This yields a geometric drawing, like the one within the gray box.

All this is just for the sake of a proper geometric interpretation. As far as a DCEL representation of such a graph is concerned, there is no need to consider spheres or, in fact, anything beyond what we have discussed before. The only difference to the case with all finite edges is that there is this special infinite vertex, which does not have any point/coordinates associated to it. But other than that, the infinite vertex is treated in exactly the same way as the finite vertices: it has in- and outgoing halfedges along which the unbounded faces can be traversed (Figure 2.11).

Remarks. It is actually not so easy to point exactly to where the DCEL data structure originates from. Often Muller and Preparata [25] are credited, but while they use the term DCEL, the data structure they describe is different from what we discussed above and from what people usually consider a DCEL nowadays. Overall, there are a large number of variants of this data structure, which appear under the names *winged edge* data structure [3], *halfedge* data structure [34], or *quad-edge* data structure [16]. Kettner [21] provides a comparison of all these and some additional references.

2.2.4 Combinatorial Embeddings

The DCEL data structure discussed in the previous section provides a fully fleshed-out representation of what is called a *combinatorial embedding*. From a mathematical point of view this can be regarded an equivalence relation on embeddings: Two embeddings are equivalent if their face boundaries—regarded as circular sequences of edges (or vertices) in counterclockwise order—are the same (as sets) up to a global change of orientation (reversing the order of all sequences simultaneously). For instance, the faces of the plane graphs shown in Figure 2.12a are (each face is described as a circular sequence of vertices)

- (a) : $\{(1,2,3), (1,3,6,4,5,4), (1,4,6,3,2)\},\$
- (b) : $\{(1,2,3,6,4,5,4), (1,3,2), (1,4,6,3)\}$, and
- (c) : $\{(1,4,5,4,6,3), (1,3,2), (1,2,3,6,4)\}.$

Note that a vertex can appear several times along the boundary of a face (if it is a cut-vertex). Clearly (b) is not equivalent to (a) nor (c), because it is the only graph that contains a face bounded by seven vertices. However, (a) and (c) turn out to be equivalent: after reverting orientations f_1 takes the role of h_2 , f_2 takes the role of h_1 , and f_3 takes the role of h_3 .



Figure 2.12: Equivalent embeddings?

Exercise 2.18. Let G be a planar graph with vertex set $\{1, \ldots, 9\}$. Try to find an embedding corresponding to the following list of circular sequences of faces:

- (a) $\{(1,4,5,6,3), (1,3,6,2), (1,2,6,7,8,9,7,6,5), (7,9,8), (1,5,4)\}$
- $(b) \{(1,4,5,6,3), (1,3,6,2), (1,2,6,7,8,9,7,6,5), (7,9,8), (1,4,5)\}$

In a dual interpretation one can just as well define equivalence in terms of the cyclic order of neighbors around all vertices. In this form, a compact way to describe a combinatorial embedding is as a so-called *rotation system* that consists of a permutation π and an involution ρ , both of which are defined on the set of halfedges (in this context often called *darts* or *flags*) of the embedding. The orbits of π correspond to the vertices, as they iterate over the incident halfedges. The involution ρ maps each halfedge to its twin. Many people prefer this dual view because one does not have to discuss the issue of vertices or edges that appear several times on the boundary of a face. The following lemma shows that such an issue does not arise when dealing with biconnected graphs.

Lemma 2.19. In a biconnected plane graph every face is bounded by a cycle.

We leave the proof as an exercise. Intuitively the statement is probably clear. But we believe it is instructive to think about how to make a formal argument. An easy consequence is the following corollary, whose proof we also leave as an exercise.

Corollary 2.20. In a 3-connected plane graph the neighbors of a vertex lie on a cycle.

Note that the statement does not read "form a cycle" but rather "lie on a cycle".

Exercise 2.21. Prove Lemma 2.19 and Corollary 2.20.

2.3 Unique Embeddings

We have seen in Lemma 2.19 that all faces in biconnected plane graphs are bounded by cycles. Conversely one might wonder which cycles of a planar graph G bound a face in *some* plane embedding of G. Such a cycle is called a *facial cycle* (Figure 2.13).



Figure 2.13: The cycle (1,2,3) is facial and we can show that (2,3,4) is not.

Exercise 2.22. Describe a linear time algorithm that for a given (abstract) planar graph G and a cycle C in G tests whether C is a facial cycle of G. (You may use planarity testing can be done in linear time.)

In fact, we will look at a slightly different class of cycles, namely those that bound a face in *every* plane embedding of G. The lemma below provides a complete characterization of those cycles. In order to state it, let us introduce a bit more terminology. A *chord* of a cycle C in a graph G is an edge that connects two vertices of C but is not an edge of C. A cycle C in a graph G is an *induced cycle*, if C = G[V(C)], that is, C does not have any chord in G.

Lemma 2.23. Let C be a cycle in a planar graph G such that $G \neq C$ and G is not C plus a single chord of C. Then C bounds a face in every plane embedding of G if and only if C is an induced cycle and it is not separating (i.e., $G \setminus C$ is connected).

Proof. " \Leftarrow ": Consider any plane embedding Γ of G. As $G \setminus C$ is connected, by the Jordan Curve Theorem it is contained either in the interior of C or in the exterior of C in Γ . In either case, the other component of the plane is bounded by C, because there are no edges among the vertices of C.

" \Rightarrow ": Using contraposition, suppose that C is not induced or $G \setminus C$ is disconnected. We have to show that there exists a plane embedding of G in which C does not bound a face.

If C is not induced, then there is a chord c of C in G. As $G \neq C \cup c$, either G has a vertex v that is not in C or G contains another chord $d \neq c$ of C. In either case, consider any plane embedding Γ of G in which C bounds a face. (If such an embedding does not exist, there is nothing to show.) We can modify Γ by drawing the chord c in the face bounded by C to obtain an embedding Γ' of G in which C does not bound a face: one of the two regions bounded by C according to the Jordan Curve Theorem contains c and the other contains either the vertex v or the other chord d.

If $G \setminus C$ contains two components A and B, then consider a plane embedding Γ of G. If C is not a face in Γ , there is nothing to show. Hence suppose that C is a face of Γ (Figure 2.14a). From Γ we obtain induced plane embeddings Γ_A of $G \setminus B \supseteq A \cup C$ and Γ_B of $B \cup C$. Using Theorem 2.2 we may suppose that C bounds the outer face in Γ_A and it does not bound the outer face in Γ_B . Then we can glue both embeddings at C, that is, extend Γ_B to an embedding of G by adding Γ_A within the face bounded by C (Figure 2.14b). The resulting embedding is a plane drawing of G in which C does not bound a face.



Figure 2.14: Construct a plane embedding of G in which C does not bound a face.

Finally, consider the case that $G \setminus C = \emptyset$ (which is not a connected graph according to our definition). As we considered above the case that C is not an induced cycle, the only remaining case is G = C, which is excluded explicitly.

For both special cases for G that are excluded in Lemma 2.23 it is easy to see that all cycles in G bound a face in every plane embedding. This completes the characterization. Also observe that in these special cases G is not 3-connected.

Corollary 2.24. A cycle C of a 3-connected planar graph G bounds a face in every plane embedding of G if and only if C is an induced cycle and it is not separating. \Box

The following theorem tells us that for a wide range of graphs we have little choice as far as a plane embedding is concerned, at least from a combinatorial point of view. Geometrically, there is still a lot of freedom, though.

Theorem 2.25 (Whitney [35]). A 3-connected planar graph has a unique combinatorial plane embedding (up to equivalence).

Proof. Let G be a 3-connected planar graph and suppose there exist two embeddings Φ_1 and Φ_2 of G that are not equivalent. That is, there is a cycle $C = (v_1, \ldots, v_k), k \ge 3$, in G that bounds a face in, say, Φ_1 but C does not bound a face in Φ_2 . By Corollary 2.24 such a cycle has a chord or it is separating. We consider both options.

Case 1: C has a chord $\{v_i, v_j\}$, with $j \ge i + 2$. Denote $A = \{v_x : i < x < j\}$ and $B = \{v_x : x < i \lor j < x\}$ and observe that both A and B are nonempty (because $\{v_i, v_j\}$ is a chord and so v_i and v_j are not adjacent in C). Given that G is 3-connected, there is at least one path P from A to B that does not use either of v_i or v_j . Let a denote the last vertex of P that is in A, and let b denote the first vertex of P that is in B. As C bounds a face f in Φ_1 , we can add a new vertex v inside the face bounded by C and connect v by four pairwise internally disjoint curves to each of v_i, v_j, a , and b. The result is a plane graph $G' \supset G$ that contains a subdivision of K₅ with branch vertices v, v_i, v_j, a , and b. By Kuratowski's Theorem (Theorem 2.10) this contradicts the planarity of G'.



Figure 2.15: Illustration of the two cases in Theorem 2.25.

Case 2: C is induced and separating. Then $G \setminus C$ contains two distinct components A and B. (We have $V(G) \neq V(C)$ and, in particular, $G \setminus C \neq \emptyset$ because C is induced and G is 3-connected.) Consider now the embedding Φ_1 in which C bounds a face, without loss of generality (Theorem 2.2) a bounded face f. Hence both A and B are embedded in the exterior of f.

Choose vertices $a \in A$ and $b \in B$ arbitrarily. As G is 3-connected, by Menger's Theorem (Theorem 1.2), there are at least three pairwise internally vertex-disjoint paths from a to b. Fix three such paths $\alpha_1, \alpha_2, \alpha_3$ and denote by c_i the first point of α_i that is on C, for $1 \leq i \leq 3$. Note that c_1, c_2, c_3 are well defined because C separates A and B, and they are pairwise distinct. Therefore, $\{a, b\}$ and $\{c_1, c_2, c_3\}$ are branch vertices of a $K_{2,3}$ subdivision in G. We can add a new vertex ν inside the face bounded by C and

connect v by three pairwise internally disjoint curves to each of c_1 , c_2 , and c_3 . The result is a plane graph $G' \supset G$ that contains a $K_{3,3}$ subdivision. By Kuratowski's Theorem (Theorem 2.10) this contradicts the planarity of G'.

In both cases we arrived at a contradiction and so there does not exist such a cycle C. Thus Φ_1 and Φ_2 are equivalent.

Whitney's Theorem does not provide a characterization of unique embeddability in general because there are both biconnected graphs that have a unique plane embedding (such as cycles) and biconnected graphs that admit several nonequivalent plane embeddings (for instance, a triangulated pentagon).

Exercise 2.26. Describe a family of biconnected planar graphs with exponentially many (combinatorial) plane embeddings. That is, show that there exists a constant $c \in \mathbb{R}$ such that for every $n \in \mathbb{N}$ there exists a biconnected planar graph on n vertices that has at least c^n pairwise distinct (combinatorial) plane embeddings.

2.4 Triangulating a Plane Graph

We like to study worst case scenarios not so much to dwell on "how bad things could get" but rather—phrased positively—because worst case examples provide universal bounds of the form "things are always at least this good". Most questions related to embeddings get harder the more edges the graph has because every additional edge needs to avoid potential crossings with other edges. Therefore, let us study the class of maximal planar graphs. A graph is *maximal planar* if no edge can be added so that the resulting graph is still planar. Corollary 2.5 tells us that a (maximal) planar graph on n vertices cannot have more than 3n - 6 edges. Yet we would like to learn a bit more about how these graphs look like.

Lemma 2.27. A maximal planar graph on $n \ge 3$ vertices is biconnected.

Proof. Consider a maximal planar graph G = (V, E). We may suppose that G is connected because adding an edge between two distinct components of a planar graph maintains planarity. Therefore, if G is not biconnected, then it has a cut-vertex v. Take a plane drawing Γ of G. As $G \setminus v$ is disconnected, removal of v also splits $N_G(v)$ into at least two components. Therefore, there are two vertices $a, b \in N_G(v)$ that are adjacent in the circular order of vertices around v in Γ and are in different components of $G \setminus v$. In particular, $\{a, b\} \notin E$ and we can add this edge to G (routing it very close to the path (a, v, b) in Γ) without violating planarity. This is in contradiction to G being maximal planar and so G is biconnected.

Lemma 2.28. In a maximal planar graph on $n \ge 3$ vertices, all faces are topological triangles, that is, every face is bounded by exactly three edges.

Proof. Consider a maximal planar graph G = (V, E) and a plane drawing Γ of G. By Lemma 2.27 we know that G is biconnected and so by Lemma 2.19 every face of Γ is

bounded by a cycle. Suppose that there is a face f in Γ that is bounded by a cycle ν_0, \ldots, ν_{k-1} of $k \ge 4$ vertices. We claim that at least one of the edges $\{\nu_0, \nu_2\}$ or $\{\nu_1, \nu_3\}$ is not present in G.

Suppose to the contrary that $\{\{v_0, v_2\}, \{v_1, v_3\}\} \subseteq E$. Then we can add a new vertex v' in the interior of f and connect v' inside f to all of v_0, v_1, v_2, v_3 by an edge (curve) without introducing a crossing. In other words, given that G is planar, also the graph $G' = (V \cup \{v'\}, E \cup \{\{v_i, v'\} : i \in \{0, 1, 2, 3\}\})$ is planar. However, v_0, v_1, v_2, v_3, v' are branch vertices of a K₅ subdivision in G': v' is connected to all other vertices within f, along the boundary ∂f of f each vertex v_i is connected to both $v_{(i-1) \mod 4}$ and $v_{(i+1) \mod 4}$ and the missing two connections are provided by the edges $\{v_0, v_2\}$ and $\{v_1, v_3\}$ (Figure 2.16a). By Kuratowski's Theorem this is in contradiction to G' being planar. Therefore, one of the edges $\{v_0, v_2\}$ or $\{v_1, v_3\}$ is not present in G, as claimed.



Figure 2.16: Every face of a maximal planar graph is a topological triangle.

So suppose without loss of generality that $\{v_1, v_3\} \notin E$. But then we can add this edge (curve) within f to Γ without introducing a crossing (Figure 2.16b). It follows that the edge $\{v_1, v_3\}$ can be added to G without sacrificing planarity, which is in contradiction to G being maximal planar. Therefore, there is no such face f bounded by four or more vertices.

Exercise 2.29. Does every minimal nonplanar graph G (that is, every nonplanar graph G whose proper subgraphs are all planar) contain an edge e such that $G \setminus e$ is maximal planar?

Many questions for graphs are formulated for connected graphs only because it is easy to add edges to a disconnected graph to make it connected. For similar reasons many questions about planar embeddings are formulated for maximal planar graphs only because it is easy to add edges to a planar graph so as to make it maximal planar. Well, this last statement is not entirely obvious. Let us look at it in more detail.

An augmentation of a given planar graph G = (V, E) to a maximal planar graph G' = (V, E') with $E' \supseteq E$ is also called a *topological triangulation*. The proof of Lemma 2.28 already contains the basic idea for an algorithm to topologically triangulate a plane graph.

Theorem 2.30. For a given connected plane graph G = (V, E) on n vertices one can compute in O(n) time and space a maximal plane graph G' = (V, E') with $E \subseteq E'$.

Proof. Suppose, for instance, that G is represented as a DCEL², from which one can easily extract the face boundaries. If some vertex v appears several times along the boundary of a single face, it is a cut-vertex. We fix this by adding an edge between the two neighbors of all but the first occurrence of v. This can easily be done in linear time by maintaining a counter for each vertex on the face boundary. The total number of edges and vertices along the boundary of all faces is proportional to the number of edges in G, which by Corollary 2.5 is linear. Hence we may suppose that all faces of G are bounded by a cycle.

For every face f that is bounded by more than three vertices, select a vertex v_f on its boundary and store with every vertex all faces that select it. Then process every vertex v as follows: First mark all neighbors of v in G. Then process all faces that selected v. For each such face f with $v_f = v$ iterate over the boundary $\partial f = (v, v_1, \ldots, v_k)$, where $k \ge 3$, of f to test whether there is any marked vertex other than the two neighbors v_1 and v_k of v along ∂f .

If there is no such vertex, we can safely triangulate f using a star from ν , that is, by adding the edges $\{\nu, \nu_i\}$, for $i \in \{2, \ldots, k-1\}$ (Figure 2.17a).

Otherwise, let v_x be the first marked vertex in the sequence v_2, \ldots, v_{k-1} . The edge $\{v, v_x\}$ that is embedded as a curve in the exterior of f prevents any vertex from v_1, \ldots, v_{x-1} from being connected by an edge in G to any vertex from v_{x+1}, \ldots, v_k . (This is exactly the argument that we made in the proof of Lemma 2.28 above for the edges $\{v_0, v_2\}$ and $\{v_1, v_3\}$, see Figure 2.16a.) In particular, we can safely triangulate f using a bi-star from v_1 and v_{x+1} , that is, by adding the edges $\{v_1, v_i\}$, for $i \in \{x + 1, \ldots, k\}$, and $\{v_j, v_{x+1}\}$, for $j \in \{2, \ldots, x - 1\}$ (Figure 2.17b).



(a) Case 1: ν does not have any neighbor on ∂f other than ν_1 and ν_k .



(b) Case 2: v has a neighbor v_x on ∂f other than v_1 and v_k .

Figure 2.17: Topologically triangulating a plane graph.

Finally, conclude the processing of v by removing all marks on its neighbors.

²If you wonder how the—possibly complicated—curves that correspond to edges are represented: they do not need to be, since here we need a representation of the combinatorial embedding only.

Regarding the runtime bound, note that every face is traversed a constant number of times. In this way, each edge is touched a constant number of times, which by Corollary 2.5 uses linear time overall. Similarly, marking the neighbors of a chosen vertex is done at most twice (mark und unmark) per vertex. Therefore, the overall time needed can be bounded by $\sum_{\nu \in V} \deg_G(\nu) = 2|\mathsf{E}| = O(n)$ by the Handshaking Lemma and Corollary 2.5.

Theorem 2.31. A maximal planar graph on $n \ge 4$ vertices is 3-connected.

Exercise 2.32. Prove Theorem 2.31.

Using any of the standard planarity testing algorithms we can obtain a combinatorial embedding of a planar graph in linear time. Together with Theorem 2.30 this yields the following

Corollary 2.33. For a given planar graph G = (V, E) on n vertices one can compute in O(n) time and space a maximal planar graph G' = (V, E') with $E \subseteq E'$.

The results discussed in this section can serve as a tool to fix the combinatorial embedding for a given graph G: augment G using Theorem 2.30 to a maximal planar graph G', whose combinatorial embedding is unique by Theorem 2.25.

Being maximal planar is a property of an abstract graph. In contrast, a geometric graph to which no straight-line edge can be added without introducing a crossing is called a *triangulation*. Not every triangulation is maximal planar, as the example depicted to the right shows.



It is also possible to triangulate a geometric graph in linear time. But this problem is much more involved. Triangulating a single face of a geometric graph amounts to what is called "triangulating a simple polygon". This can be done in near-linear³ time using standard techniques, and in linear time using Chazelle's famous algorithm, whose description spans a fourty pages paper [9].

Exercise 2.34. We discussed the DCEL structure to represent plane graphs in Section 2.2.1. An alternative way to represent an embedding of a maximal planar graph is the following: For each triangle, store references to its three vertices and to its three neighboring triangles. Compare both approaches. Discuss different scenarios where you would prefer one over the other. In particular, analyze the space requirements of both.

Connectivity serves as an important indicator for properties of planar graphs. Already Wagner showed that a 4-connected graph is planar if and only if it does not contain K_5 as a minor. That is, assuming 4-connectivity the second forbidden minor $K_{3,3}$ becomes "irrelevant". For subdivisions this is a different story. Independently Kelmans and Semour conjectured in the 1970s that 5-connectivity allows to consider K_5 subdivisions only. This conjecture was proven only recently⁴ by Dawei He, Yan Wang, and Xingxing Yu.

 $^{^{3}}O(n \log n)$ or—using more elaborate tools— $O(n \log^{*} n)$ time

⁴The result was announced in 2015 and published in 2020.

Theorem 2.35 (He-Wang-Yu [18]). Every 5-connected nonplanar graph contains a subdivision of K_5 .

Exercise 2.36. Give a 4-connected nonplanar graph that does not contain a subdivision of K_5 .

Another example that illustrates the importance of the parameter connectivity is the following famous theorem of Tutte that provides a sufficient condition for Hamiltonicity.

Theorem 2.37 (Tutte [31]). Every 4-connected planar graph is Hamiltonian.

Moreover, for a given 4-connected planar graph a Hamiltonian cycle can also be computed in linear time [10].

2.5 Compact Straight-Line Drawings

As a next step we consider plane embeddings in the geometric setting, where every edge is drawn as a straight-line segment. A classical theorem of Wagner and Fáry states that this is not a restriction in terms of plane embeddability.

Theorem 2.38 (Fáry [13], Wagner [32]). Every planar graph has a plane straight-line embedding.

This statement is quite surprising, considering how much more freedom arbitrarily complex Jordan arcs allow compared to line segments, which are completely determined by their endpoints. In order to further increase the level of appreciation, let us note that a similar "straightening" is not possible when fixing the point set on which the vertices are to be embedded: Pach and Wenger [26] showed that for a given planar graph G on n vertices v_1, \ldots, v_n and a given set $\{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ of n points, one can always find a plane embedding of G such that $v_i \mapsto p_i$, for all $i \in \{1, \ldots, n\}$. However, this is not possible in general with a plane straight-line embedding. For instance, K_4 does not admit a plane straight-line embedding on a set of points that form a convex quadrilateral, such as a rectangle. In fact, it is NP-hard to decide whether a given planar graph admits a plane straight-line embedding on a given point set [7].

Exercise 2.39. a) Show that for every natural number $n \ge 4$ there exist a planar graph G on n vertices and a set $P \subset \mathbb{R}^2$ of n points in general position (no three points are collinear) so that G does not admit a plane straight-line embedding on P.

b) Show that for every natural number $n \ge 6$ there exist a planar graph G on n vertices and a set $P \subset \mathbb{R}^2$ of n points so that (1) P is in general position (no three points are collinear); (2) P has a triangular convex hull (that is, there are three points in P that form a triangle that contains all other points from P); and (3) G does not admit a plane straight-line embedding on P.

Exercise 2.40. Show that for every set $P \subset \mathbb{R}^2$ of $n \ge 3$ in general position (no three points are collinear) the cycle C_n on n vertices admits a plane straight-line embedding on P.

Although Fáry-Wagner's theorem has a nice inductive proof, we will not discuss it here. Instead we will prove a stronger statement that implies Theorem 2.38.

A very nice property of straight-line embeddings is that they are easy to represent: We need to store points/coordinates for the vertices only. From an algorithmic and complexity point of view the space needed by such a representation is important because it appears in the input and output size of algorithms that work on embedded graphs. While the Fáry-Wagner Theorem guarantees the existence of a plane straight-line embedding for every planar graph, it does not provide bounds on the size of the coordinates used in the representation. But the following strengthening provides such bounds, by describing an algorithm that embeds (without crossings) a given planar graph on a linear size integer grid.

Theorem 2.41 (de Fraysseix, Pach, Pollack [15]). Every planar graph on $n \ge 3$ vertices has a plane straight-line drawing on the $(2n-3) \times (n-1)$ integer grid.

2.5.1 Canonical Orderings

The key concept behind the algorithm is the notion of a canonical ordering, which is a vertex order that allows to construct a plane drawing in a natural (hence canonical) way. Reading it backwards one may think of a shelling or peeling order that destructs the graph vertex by vertex from the outside. A canonical ordering also provides a succinct representation for the combinatorial embedding.

Definition 2.42. A plane graph is internally triangulated if it is biconnected and every bounded face is a (topological) triangle. Let G be an internally triangulated plane graph and $C_o(G)$ its outer cycle. A permutation $\pi = (\nu_1, \nu_2, \dots, \nu_n)$ of V(G) is a canonical ordering for G, if

(CO1) G_k is internally triangulated, for all $k \in \{3, \ldots, n\}$;

(CO2) v_1v_2 is on the outer cycle $C_o(G_k)$ of G_k , for all $k \in \{3, \ldots, n\}$;

(CO3) v_{k+1} is located in the outer face of G_k , for all $k \in \{3, \ldots, n-1\}$;

where G_k is the subgraph of G induced by v_1, \ldots, v_k .

Figure 2.18 shows an example. Note that there are permutations that do not correspond to a canonical order: for instance, when choosing the vertex 4 as the next vertex to be removed in Figure 2.18b, the resulting graph $G'_7 = G[\{1,2,3,5,6,7,8\}]$ is not biconnected (because 1 is a cut-vertex).



Figure 2.18: An internally triangulated plane graph with a canonical ordering.

Theorem 2.43. For every internally triangulated plane graph G and every edge $\{v_1, v_2\}$ on its outer face, there exists a canonical ordering for G that starts with v_1, v_2 . Moreover, such an ordering can be computed in linear time.

Proof. Induction on n, the number of vertices. For a triangle, any order suffices and so the statement holds. Hence consider an internally triangulated plane graph G = (V, E) on $n \ge 4$ vertices. We claim that it is enough to select a vertex $\nu_n \notin \{\nu_1, \nu_2\}$ on $C_{\circ}(G)$ that is not incident to a chord of $C_{\circ}(G)$ and then apply induction on $G \setminus \{\nu_n\}$.

We will show later that such a vertex v_n always exists. First let us prove the claim. We need to argue that if v_n is selected as described

- (i) the plane graph $G_{n-1} := G \setminus \{v_n\}$ is internally triangulated,
- (ii) the given edge $\{\nu_1,\nu_2\}$ is on the outer face $C_\circ(G_{n-1})$ of $G_{n-1},$ and
- (iii) we can extend the inductively obtained canonical ordering for G_{n-1} with v_n to obtain a canonical ordering for G.

Property (ii) is an immediate consequence of $v_n \notin \{v_1, v_2\}$. Regarding (iii) note that (CO1)-(CO3) for k = n hold by assumption. Hence to be able to apply induction it suffices to show (i).

The way G_{n-1} is obtained from G, every bounded face f of G_{n-1} also appears as a bounded face of G. As G is internally triangulated, f is a triangle. It remains to show that G_{n-1} is biconnected. The outer cycle $C_o(G_{n-1})$ of G_{n-1} is obtained from $C_o(G)$ by removing v_n and replacing it with the (possibly empty) sequence u_2, \ldots, u_{m-1} . As v_n is not incident to a chord of $C_o(G)$ (and so neither of u_2, \ldots, u_{m-1} appeared along $C_o(G)$ already), the resulting sequence forms a cycle, indeed. Add a new vertex v in the outer face of G_{n-1} and connect v to every vertex of $C_o(G_{n-1})$ to obtain a maximal planar graph $H \supset G_{n-1}$. By Theorem 2.31 H is 3-connected and so G_{n-1} is biconnected, as desired. This also completes the proof of the claim. Next let us show that we can always find a vertex $v_n \notin \{v_1, v_2\}$ on $C_o(G)$ that is not incident to a chord of $C_o(G)$. If $C_o(G)$ does not have any chord, this is obvious because every cycle has at least three vertices, one of which is neither v_1 nor v_2 . So suppose that $C_o(G)$ has a chord c. The endpoints of c split $C_o(G)$ into two paths, one of which does not have v_1 nor v_2 as an internal vertex. Among all chords of $C_o(G)$ select c such that this path has minimal length. (It has always at least two edges because there is always at least one vertex "behind" a chord.) Then by definition of c this path is an induced path in G and none of its (at least one) interior vertices is incident to a chord of $C_o(G)$ because such a chord would cross c. So we can select v_n from these vertices. By the way the path is selected with respect to c, this procedure does not select v_1 nor v_2 .

Regarding the runtime bound, we maintain the following information for each vertex v: whether it has been chosen already, whether it is on the outer face of the current graph, and the number of incident chords with respect to the current outer cycle. Given a combinatorial embedding of G, it is straighforward to initialize this information in linear time. (Every edge is considered at most twice, once for each endpoint on the outer face.) We also maintain an unordered list of the *eligible* vertices, that is, those vertices that are on the outer face and not incident to any chord. This list is straightforward to maintain: Whenever a vertex information is updated, check before and after the update whether it is eligible and correspondingly add it to or remove it from the list of eligible vertices.

When removing a vertex, there are two cases: Either v_n has two neighbors u_1 and u_2 only (Figure 2.19a), in which case the edge u_1u_2 ceases to be a chord. Thus, the chord count for u_1 and u_2 has to be decremented by one. Otherwise, there are $m \ge 3$ neighbors u_1, \ldots, n_m (Figure 2.19b) and (1) all vertices u_2, \ldots, u_{m-1} are new on the outer cycle, and (2) every edge incident to u_i , for $i \in \{2, \ldots, m-1\}$, and some other vertex on the outer cycle other than u_{i-1} or u_{i+1} is a new chord (and the corresponding counters at both endpoints have to by incremented by one).



Figure 2.19: Processing a vertex when computing a canonical ordering.

During the course of the algorithm every vertex appears once as a new vertex on the outer face. At this point all incident edges are examined. Overall, every edge is inspected at most twice—once for each endpoint—which takes linear time by Corollary 2.5. \Box

Using one of the linear time planarity testing algorithms, we can obtain a combinatorial embedding for a given maximal planar graph G. As every maximal plane graph is internally triangulated, we can then use Theorem 2.43 to provide us with a canonical ordering for G, in overall linear time.

Corollary 2.44. Every maximal planar graph admits a canonical ordering. Moreover, such an ordering can be computed in linear time. \Box

Exercise 2.45. (a) Compute a canonical ordering for the following internally triangulated plane graphs:



- (b) Give an infinite family of internally triangulated plane graphs G_n on n = 2k vertices with at least k! canonical orderings.
- (c) Give an infinite family of internally triangulated plane graphs that have a unique canonical ordering for a specific choice of the starting edge v_1v_2 .
- **Exercise 2.46.** (a) Describe a plane graph G with n vertices that can be embedded (while preserving the outer face) on a grid of size $(2n/3 1) \times (2n/3 1)$ but not on a smaller grid.
 - (b) Can you draw G on a smaller grid if you are allowed to change the embedding?

As simple as they may appear, canonical orderings are a powerful and versatile tool to work with plane graphs. As an example, consider the following partitioning theorem.

Theorem 2.47 (Schnyder [29]). For every maximal planar graph G on at least three vertices and every fixed face f of G, the multigraph obtained from G by doubling the (three) edges of f can be partitioned into three spanning trees.

Exercise 2.48. Prove Theorem 2.47. Hint: Take a canonical ordering and build one tree by taking for every vertex v_k the edge to its first neighbor on the outer cycle $C_{\circ}(G_{k-1})$.

Of a similar flavor is the following open problem, for which only partial answers are known [1, 4, 5].

Problem 2.49 (In memoriam Ferran Hurtado (1951–2014)).

Can every complete geometric graph on n = 2k vertices (the complete straight line graph on a set of n points in general position) be partitioned into k plane spanning trees?

2.5.2 The Shift-Algorithm

Let (v_1, \ldots, v_n) be a canonical ordering. The general plan is to construct an embedding by inserting vertices in this order, starting from the triangle $P(v_1) = (0,0)$, $P(v_3) = (1,1)$, $P(v_2) = (2,0)$; see Figure 2.20.



Figure 2.20: Initialization of the shift algorithm.

At each step, some vertices will be shifted to the right, making room for the edges to the freshly inserted vertex. For each vertex v_i already embedded, maintain a set $L(v_i)$ of vertices that move rigidly together with v_i . Initially $L(v_i) = \{v_i\}$, for $1 \le i \le 3$.

Ensure that the following invariants hold after Step k (that is, after v_k has been inserted):

- (i) $P(v_1) = (0,0), P(v_2) = (2k-4,0);$
- (ii) The x-coordinates of the points on $C_{\circ}(G_k) = (w_1, \ldots, w_t)$, where $w_1 = v_1$ and $w_t = v_2$, are strictly increasing (in this order)⁵;
- (iii) each edge of $C_{\circ}(G_k)$ is drawn as a straight-line segment with slope ± 1 .

Clearly these invariants hold for G_3 , embedded as described above. Invariant (i) implies that after Step n we have $P(v_2) = (2n - 4, 0)$, while (iii) implies that the Manhattan distance⁶ between any two points on $C_{\circ}(G_k)$ is even.

Idea: put v_{k+1} at $\mu(w_p, w_q)$, where w_p, \ldots, w_q are its neighbors on $C_{\circ}(G_k)$ and

$$\mu((x_{p}, y_{p}), (x_{q}, y_{q})) = \frac{1}{2}(x_{p} - y_{p} + x_{q} + y_{q}, -x_{p} + y_{p} + x_{q} + y_{q})$$

is the point of intersection between the line ℓ_1 : $y = x - x_p + y_p$ of slope 1 through $w_p = (x_p, y_p)$ and the line ℓ_2 : $y = x_q - x + y_q$ of slope -1 through $w_q = (x_q, y_q)$.

⁵The notation is a bit sloppy here because both t and the w_i in general depend on k. So in principle we should write w_i^k instead of w_i . But as the k would just make a constant appearance throughout, we omit it to avoid index clutter.

⁶The Manhattan distance of two points (x_1, y_1) and (x_2, y_2) is $|x_2 - x_1| + |y_2 - y_1|$.

Proposition 2.50. If the Manhattan distance between w_p and w_q is even, then $\mu(w_p, w_q)$ is on the integer grid.

Proof. By Invariant (ii) we know that $x_p < x_q$. Suppose without loss of generality that $y_p \leq y_q$. The Manhattan distance d of w_p and w_q is $x_q - x_p + y_q - y_p$, which by assumption is an even number. Adding the even number $2x_p$ to d yields the even number $x_q + x_p + y_q - y_p$, half of which is the x-coordinate of $\mu((x_p, y_p), (x_q, y_q))$. Adding the even number $2y_p$ to d yields the even number $x_q - x_p + y_q + y_p$, half of which is the y-coordinate of $\mu((x_p, y_p), (x_q, y_q))$.

After Step n we have $P(v_n) = (n-2, n-2)$ because v_n is a neighbor of both v_1 and v_2 . However, $P(v_{k+1})$ may not "see" all of w_p, \ldots, w_q , in case that the slope of $w_p w_{p+1}$ is 1 and/or the slope of $w_{q-1} w_q$ is -1 (Figure 2.21).



Figure 2.21: (a) The new vertex v_{k+1} is adjacent to all of w_p, \ldots, w_q . If we place v_{k+1} at $\mu(w_p, w_q)$, then some edges may overlap, in case that w_{p+1} lies on the line of slope 1 through w_p or w_{q-1} lies on the line of slope -1 through w_q ; (b) shifting w_{p+1}, \ldots, w_{q-1} by one and w_q, \ldots, w_t by two units to the right solves the problem.

In order to resolve these problems we shift some points around so that after the shift w_{p+1} does not lie on the line of slope 1 through w_p and w_{q-1} does not lie on the line of slope -1 through w_q . The process of inserting v_{k+1} then looks as follows.

- 1. Shift $\bigcup_{i=p+1}^{q-1} L(w_i)$ to the right by one unit.
- 2. Shift $\bigcup_{i=q}^{t} L(w_i)$ to the right by two units.

3.
$$P(v_{k+1}) := \mu(w_p, w_q)$$
.

4. $L(v_{k+1}) := \{v_{k+1}\} \cup \bigcup_{i=p+1}^{q-1} L(w_i).$

Observe that the Manhattan distance between w_p and w_q remains even because the shift increases their x-difference by two and leaves the y-coordinates unchanged. Therefore by Proposition 2.50 the vertex v_{k+1} is embedded on the integer grid.

The slopes of the edges $w_p w_{p+1}$ and $w_{q-1} w_q$ (might be just a single edge, in case that p+1=q) become < 1 in absolute value, whereas the slopes of all other edges along the outer cycle remain ± 1 . As all edges from v_{k+1} to w_{p+1}, \ldots, w_{q-1} have slope > 1 in

absolute value, and the edges $v_{k+1}w_p$ and $v_{k+1}w_q$ have slope ± 1 , each edge $v_{k+1}w_i$, for $i \in \{p, \ldots, q\}$ intersects the outer cycle in exactly one point, which is w_i . In other words, adding all edges from v_{k+1} to its neighbors in G_k as straight-line segments results in a plane drawing.

Next we argue that the invariants (i)-(iii) are maintained. For (i) note that we start shifting with w_{p+1} only so that even in case that v_1 is a neighbor of v_{k+1} , it is never shifted. On the other hand, v_2 is always shifted by two because we shift every vertex starting from (and including) w_q . Clearly both the shifts and the insertion of v_{k+1} maintain the strict order along the outer cycle, and so (ii) continues to hold. Finally, regarding (iii) note that the edges $w_p w_{p+1}$ and $w_{q-1} w_q$ (possibly this is just a single edge) are the only edges on the outer cycle whose slope is changed by the shift. But these edges do not appear on $C_o(G_{k+1})$ anymore. The two edges $v_{k+1}w_p$ and $v_{k+1}w_q$ incident to the new vertex v_{k+1} that appear on $C_o(G_{k+1})$ have slope 1 and -1, respectively. So all of (i)-(iii) are invariants of the algorithm, indeed.

So far we have argued about the shift with respect to vertices on the outer cycle of G_k only. To complete the proof of Theorem 2.41 it remains to show that the drawing remains plane under shifts also in its interior part.

Lemma 2.51. Let G_k , $k \ge 3$, be straight-line grid embedded as described, $C_o(G_k) = (w_1, \ldots, w_t)$, and let $\delta_1 \le \ldots \le \delta_t$ be nonnegative integers. If for each i, we shift $L(w_i)$ by δ_i to the right, then the resulting straight-line drawing is plane.

Proof. Induction on k. For G_3 this is obvious. Let $v_k = w_\ell$, for some $1 < \ell < t$. Construct a delta sequence Δ for G_{k-1} as follows. If v_k has only two neighbors in G_k , then $C_\circ(G_{k-1}) = (w_1, \ldots, w_{\ell-1}, w_{\ell+1}, \ldots, w_t)$ and we set $\Delta = \delta_1, \ldots, \delta_{\ell-1}, \delta_{\ell+1}, \ldots, \delta_t$. Otherwise, $C_\circ(G_{k-1}) = (w_1, \ldots, w_{\ell-1} = u_1, \ldots, u_m = w_{\ell+1}, \ldots, w_t)$, where u_1, \ldots, u_m are the $m \ge 3$ neighbors of v_k in G_k . In this case we set

$$\Delta = \delta_1, \dots, \delta_{\ell-1}, \underbrace{\delta_\ell, \dots, \delta_\ell}_{m-2 \text{ times}}, \delta_{\ell+1}, \dots, \delta_t$$
.

Clearly, Δ is monotonely increasing and by the inductive assumption a correspondingly shifted drawing of G_{k-1} is plane. When adding ν_k and its incident edges back, the drawing remains plane: All vertices u_2, \ldots, u_{m-1} (possibly none) move rigidly with (by exactly the same amount as) ν_k by construction. Stretching the edges of the chain $w_{\ell-1}, w_{\ell}, w_{\ell+1}$ by moving $w_{\ell-1}$ to the left and/or $w_{\ell+1}$ to the right cannot create any crossings.

Linear time. The challenge in implementing the shift algorithm efficiently lies in the eponymous shift operations, which modify the x-coordinates of potentially many vertices. In fact, it is not hard to see that a naive implementation—which keeps track of all coordinates explicitly—may use quadratic time. De Fraysseix et al. described an implementation of the shift algorithm that uses $O(n \log n)$ time. Then Chrobak and Payne [11] observed how to improve the runtime to linear, using the following ideas.

Recall that $P(v_{k+1}) = (x_{k+1}, y_{k+1})$, where

$$\begin{aligned} \mathbf{x}_{k+1} &= \frac{1}{2}(\mathbf{x}_{p} - \mathbf{y}_{p} + \mathbf{x}_{q} + \mathbf{y}_{q}) \text{ and} \\ \mathbf{y}_{k+1} &= \frac{1}{2}(-\mathbf{x}_{p} + \mathbf{y}_{p} + \mathbf{x}_{q} + \mathbf{y}_{q}) = \frac{1}{2}((\mathbf{x}_{q} - \mathbf{x}_{p}) + \mathbf{y}_{p} + \mathbf{y}_{q}). \end{aligned}$$
(2.52)

Thus,

$$x_{k+1} - x_p = \frac{1}{2} ((x_q - x_p) + y_q - y_p).$$
(2.53)

In other words, we need the y-coordinates of w_p and w_q together with the relative xposition (offset) of w_p and w_q only to determine the y-coordinate of v_{k+1} and its offset to w_p .

Maintain the outer cycle as a rooted binary tree T, with root v_1 . For each node v of T, the *left child* is the first vertex covered by insertion of v (if any), that is, w_{p+1} in the terminology from above (if $p + 1 \neq q$), whereas the *right child* of v is the next node along the outer cycle (if any; either along the current outer cycle or along the one at the point where both points were covered together). See Figure 2.22 for an example.



Figure 2.22: Maintaining a binary tree representation to keep track of the xcoordinates when inserting a new vertex v_{k+1} . Red (dashed) arrows point to left children, blue (solid) arrows point to right children.

At each node v of T we also store its x-offset dx(v) with respect to the parent node. For the root v_1 of the tree set $dx(v_1) = 0$. In this way, a whole subtree (and, thus, a whole set $L(\cdot)$) can be shifted by changing a single offset entry at its root.

Initially, $dx(v_1) = 0$, $dx(v_2) = dx(v_3) = 1$, $y(v_1) = y(v_2) = 0$, $y(v_3) = 1$, $left(v_1) = left(v_2) = left(v_3) = 0$, $right(v_1) = v_3$, $right(v_2) = 0$, and $right(v_3) = v_2$.

Inserting a vertex v_{k+1} works as follows. As before, let w_1, \ldots, w_t denote the vertices on the outer cycle $C_{\circ}(G_k)$ and w_p, \ldots, w_q be the neighbors of v_{k+1} .

- 1. Increment $dx(w_{p+1})$ and $dx(w_q)$ by one. (This implements the shift.)
- 2. Compute $\Delta_{pq} = \sum_{i=p+1}^{q} dx(w_i)$. (This is the total offset between w_p and w_q .)

- 3. Set $dx(v_{k+1}) \leftarrow \frac{1}{2}(\Delta_{pq} + y(w_q) y(w_p))$ and $y(v_{k+1}) \leftarrow \frac{1}{2}(\Delta_{pq} + y(w_q) + y(w_p))$. (This is exactly what we derived in (2.52) and (2.53).)
- 4. Set right(w_p) $\leftarrow v_{k+1}$ and right(v_{k+1}) $\leftarrow w_q$. (Update the current outer cycle.)
- 5. If p + 1 = q, then set left(v_{k+1}) ← 0; else set left(v_{k+1}) ← w_{p+1} and right(w_{q-1}) ← 0. (Update L(v_{k+1}), the part that is covered by insertion of v_{k+1}.)
- 6. Set dx(w_q) ← Δ_{pq} dx(v_{k+1});
 if p + 1 ≠ q, then set dx(w_{p+1}) ← dx(w_{p+1}) dx(v_{k+1}).
 (Update the offsets according to the changes in the previous two steps.)

Observe that the only step that possibly cannot be executed in constant time is Step 2. To analyze Step 2, note that all vertices but the last vertex w_q for which we sum the offsets are covered by the insertion of v_{k+1} . As every vertex can be covered at most once, the overall complexity of this step during the algorithm is linear. Therefore, this first phase of the algorithm can be completed in linear time.

In a second phase, the final x-coordinates can be computed from the offsets by a single recursive pre-order traversal of the tree. The (pseudo-)code given below is to be called with the root vertex v_1 and an offset of zero. Clearly this yields a linear time algorithm overall.

```
compute_coordinate(Vertex v, Offset d) {
  if (v == 0) return;
   x(v) = dx(v) + d;
   compute_coordinate(left(v), x(v));
   compute_coordinate(right(v), x(v));
}
```

2.5.3 Remarks and Open Problems

From a geometric complexity point of view, Theorem 2.41 provides very good news for planar graphs in a similar way that the Euler Formula does from a combinatorial complexity point of view. Euler's Formula tells us that we can obtain a combinatorial representation (for instance, as a DCEL) of any plane graph using O(n) space, where n is the number of vertices.

Now the shift algorithm tells us that for any planar graph we can even find a geometric plane (straight-line) representation using O(n) space. In addition to the combinatorial information, we only have to store 2n numbers from the range $\{0, 1, \ldots, 2n - 4\}$.

When we make such claims regarding space complexity we implicitly assume the socalled *word RAM model*. In this model each address in memory contains a *word* of b bits, which means that it can be used to represent any integer from $\{0, \ldots, 2^b - 1\}$. One also assumes that b is sufficiently large, for instance, in our case $b \ge \log n$. There are also different models such as the *bit complexity model*, where one is charged for every bit used to store information. In our case that would already incur an additional factor of log n for the combinatorial representation: for instance, for each halfedge we store its endpoint, which is an index from $\{1, ..., n\}$.

Edge lengths. Theorem 2.41 shows that planar graphs admit a plane straight-line drawing where all vertices have integer coordinates. It is an open problem whether a similar statement can be made for edge lengths.

Problem 2.54 (Harborth's Conjecture [17]). Every planar graph admits a plane straightline drawing where all Euclidean edge lengths are integral.

Without the planarity restriction such a drawing is possible because for every $n \in \mathbb{N}$ one can find a set of n points in the plane, not all collinear, such that their distances are all integral. In fact, such a set of points can be constructed to lie on a circle of integral radius [2]. When mapping the vertices of K_n onto such a point set, all edge lengths are integral. In the same paper it is also shown that there exists no infinite set of points in the plane so that all distances are integral, unless all of these points are collinear. Unfortunately, collinear point sets are not very useful for drawing graphs. The existence of a dense subset of the plane where all distances are rational would resolve Harborth's Conjecture. However, it is not known whether such a set exists, and in fact the suspected answer is "no".

Problem 2.55 (Erdős–Ulam Conjecture [12]). There is no dense set of points in the plane whose Euclidean distances are all rational.

Generalizing the Fáry-Wagner Theorem. As discussed above, not every planar graph on n vertices admits a plane straight-line embedding on every set of n points. But Theorem 2.38 states that for every planar graph G on n vertices there *exists* a set P of n points in the plane so that G admits a plane straight-line embedding on P (that is, so that the vertices of G are mapped bijectively to the points in P). It is an open problem whether this statement can be generalized to hold for several graphs, in the following sense.

Problem 2.56. What is the largest number $k \in \mathbb{N}$ for which the following statement holds? For every collection of k planar graphs G_1, \ldots, G_k on n vertices each, there exists a set P of n points so that G_i admits a plane straight-line embedding on P, for every $i \in \{1, \ldots, k\}$.

By Theorem 2.38 we know that the statement holds for k = 1. Already for k = 2 it is not known whether the statement holds. However, it is known that k is finite [8]. Specifically, there exists a collection of 49 planar graphs on 11 vertices each so that for every set P of 11 points in the plane at least one of these graphs does not admit a plane straight-line embedding on P [28]. Therefore we have $k \leq 49$.

Questions

- 1. What is an embedding? What is a planar/plane graph? Give the definitions and explain the difference between planar and plane.
- 2. How many edges can a planar graph have? What is the average vertex degree in a planar graph? Explain Euler's formula and derive your answers from it.
- 3. How can plane graphs be represented on a computer? Explain the DCEL data structure and how to work with it.
- 4. How can a given plane graph be (topologically) triangulated efficiently? Explain what it is, including the difference between topological and geometric triangulation. Give a linear time algorithm, for instance, as in Theorem 2.30.
- 5. What is a combinatorial embedding? When are two combinatorial embeddings equivalent? Which graphs have a unique combinatorial plane embedding? Give the definitions, explain and prove Whitney's Theorem.
- 6. What is a canonical ordering and which graphs admit such an ordering? For a given graph, how can one find a canonical ordering efficiently? Give the definition. State and prove Theorem 2.43.
- 7. Which graphs admit a plane embedding using straight line edges? Can one bound the size of the coordinates in such a representation? State and prove Theorem 2.41.

References

- Oswin Aichholzer, Thomas Hackl, Matias Korman, Marc van Kreveld, Maarten Löffler, Alexander Pilz, Bettina Speckmann, and Emo Welzl, Packing plane spanning trees and paths in complete geometric graphs. *Inform. Process. Lett.*, 124, (2017), 35-41.
- [2] Norman H. Anning and Paul Erdős, Integral distances. Bull. Amer. Math. Soc., 51/8, (1945), 598-600.
- [3] Bruce G. Baumgart, A polyhedron representation for computer vision. In Proc. AFIPS Natl. Comput. Conf., vol. 44, pp. 589–596, AFIPS Press, Alrington, Va., 1975.
- [4] Ahmad Biniaz and Alfredo García, Packing plane spanning trees into a point set. Comput. Geom. Theory Appl., 90, (2020), 101653.
- [5] Prosenjit Bose, Ferran Hurtado, Eduardo Rivera-Campo, and David R. Wood, Partitions of complete geometric graphs into plane trees. Comput. Geom. Theory Appl., 34/2, (2006), 116-125.

- [6] John M. Boyer and Wendy J. Myrvold, On the cutting edge: simplified O(n) planarity by edge addition. J. Graph Algorithms Appl., 8/3, (2004), 241-273.
- [7] Sergio Cabello, Planar embeddability of the vertices of a graph using a fixed point set is NP-hard. J. Graph Algorithms Appl., 10/2, (2006), 353-363.
- [8] Jean Cardinal, Michael Hoffmann, and Vincent Kusters, On universal point sets for planar graphs. J. Graph Algorithms Appl., 19/1, (2015), 529-547.
- Bernard Chazelle, Triangulating a simple polygon in linear time. Discrete Comput. Geom., 6/5, (1991), 485-524.
- [10] Norishige Chiba and Takao Nishizeki, The Hamiltonian cycle problem is linear-time solvable for 4-connected planar graphs. J. Algorithms, 10/2, (1989), 187–211.
- [11] Marek Chrobak and Thomas H. Payne, A linear-time algorithm for drawing a planar graph on a grid. Inform. Process. Lett., 54, (1995), 241-246.
- [12] Paul Erdős, Ulam, the man and the mathematician. J. Graph Theory, 9/4, (1985), 445-449.
- [13] István Fáry, On straight lines representation of planar graphs. Acta Sci. Math. Szeged, 11/4, (1948), 229–233.
- [14] Hubert de Fraysseix, Patrice Ossona de Mendez, and Pierre Rosenstiehl, Trémaux trees and planarity. Internat. J. Found. Comput. Sci., 17/5, (2006), 1017-1030.
- [15] Hubert de Fraysseix, János Pach, and Richard Pollack, How to draw a planar graph on a grid. Combinatorica, 10/1, (1990), 41-51.
- [16] Leonidas J. Guibas and Jorge Stolfi, Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. ACM Trans. Graph., 4/2, (1985), 74–123.
- [17] Heiko Harborth and Arnfried Kemnitz, Plane integral drawings of planar graphs. Discrete Math., 236/1-3, (2001), 191-195.
- [18] Dawei He, Yan Wang, and Xingxing Yu, The Kelmans-Seymour conjecture IV: A proof. J. Combin. Theory Ser. B, 144, (2020), 309–358.
- [19] John Hopcroft and Robert E. Tarjan, Efficient planarity testing. J. ACM, 21/4, (1974), 549-568.
- [20] Ken ichi Kawarabayashi, Yusuke Kobayashi, and Bruce Reed, The disjoint paths problem in quadratic time. J. Combin. Theory Ser. B, 102/2, (2012), 424–435.

- [21] Lutz Kettner, Software design in computational geometry and contour-edge based polyhedron visualization. Ph.D. thesis, ETH Zürich, Zürich, Switzerland, 1999.
- [22] Kazimierz Kuratowski, Sur le problème des courbes gauches en topologie. Fund. Math., 15/1, (1930), 271–283.
- [23] László Lovász, Graph minor theory. Bull. Amer. Math. Soc., 43/1, (2006), 75–86.
- [24] Bojan Mohar and Carsten Thomassen, Graphs on surfaces, Johns Hopkins University Press, Baltimore, 2001.
- [25] David E. Muller and Franco P. Preparata, Finding the intersection of two convex polyhedra. Theoret. Comput. Sci., 7, (1978), 217-236.
- [26] János Pach and Rephael Wenger, Embedding planar graphs at fixed vertex locations. Graphs Combin., 17, (2001), 717–728.
- [27] Neil Robertson and Paul Seymour, Graph Minors. XX. Wagner's Conjecture. J. Combin. Theory Ser. B, 92/2, (2004), 325–357.
- [28] Manfred Scheucher, Hendrik Schrezenmaier, and Raphael Steiner, A Note on Universal Point Sets for Planar Graphs. In Proc. 27th Int. Sympos. Graph Drawing Network Visualization (GD 2019), vol. 11904 of Lecture Notes Comput. Sci., pp. 350-362, Springer, 2019.
- [29] Walter Schnyder, Planar graphs and poset dimension. Order, 5, (1989), 323–343.
- [30] Carsten Thomassen, Kuratowski's Theorem. J. Graph Theory, 5/3, (1981), 225– 241.
- [31] William T. Tutte, A theorem on planar graphs. Trans. Amer. Math. Soc., 82/1, (1956), 99–116.
- [32] Klaus Wagner, Bemerkungen zum Vierfarbenproblem. Jahresbericht der Deutschen Mathematiker-Vereinigung, 46, (1936), 26-32.
- [33] Klaus Wagner, Uber eine Eigenschaft der ebenen Komplexe. Math. Ann., 114/1, (1937), 570–590.
- [34] Kevin Weiler, Edge-based data structures for solid modeling in a curved surface environment. *IEEE Comput. Graph. Appl.*, 5/1, (1985), 21–40.
- [35] Hassler Whitney, Congruent graphs and the connectivity of graphs. Amer. J. Math., 54/1, (1932), 150–168.