

Planar Graphs

Francis Lazarus and Arnaud de Mesmay

September 19, 2017

Contents

1	Topology	2
1.1	The Jordan Curve Theorem	3
1.2	Euler's Formula	5
2	Kuratowski's Theorem	7
2.1	The Subdivision Version	7
2.2	The Minor Version	11
3	Other Planarity Characterizations	12
4	Planarity Test	15
5	Drawing with Straight Lines	19

A graph is planar if it can be drawn on a sheet of paper so that no two edges intersect, except at common endpoints. This simple property not only allows to visualize planar graphs easily, but implies many nice properties. Planar graphs are sparse: they have a linear number of edges with respect to their number of vertices (specifically a simple planar graph with n vertices has at most $3n - 6$ edges), they are 4-colorable, they can be encoded efficiently, etc. Classical examples of planar graphs include the graphs formed by the vertices and edges of the five Platonic polyhedra, and in fact of any convex polyhedron. Although being planar is a topological property, planar graphs have purely combinatorial characterizations. Such characterizations may lead to efficient algorithms for planarity testing or, more surprisingly, for geometric embedding (=drawing).

In the first part of this lecture we shall deduce the combinatorial characterizations of planar graphs from their topological definition. That we can get rid of topological considerations should not be surprising. It is actually possible to develop a combinatorial theory of surfaces where a drawing of a graph is *defined* by a circular ordering of its edges around each vertex. The collection of these circular orderings is called a **rotation system**. A rotation system is thus described combinatorially by a single permutation over the (half-)edges of a graph; the cycle decomposition of the permutation induces the circular orderings around each vertex. The topology of the surface corresponding to a rotation system can be deduced from the computation of its Euler characteristic. Being planar then reduces to the existence of a rotation system with the appropriate Euler characteristic.

The following notes are largely inspired by the monographs of Mohar and Thomassen [MT01] and of Diestel [Die05].

1 Topology

A **graph** $G = (V, E)$ is defined by a set $V = V(G)$ of vertices and a set $E = E(G)$ of edges where each edge is associated one or two vertices, called its **endpoints**. A **loop** is an edge with a single endpoint. Edges sharing the same endpoints are said **parallel** and define a **multiple edge**. A graph without loops or multiple edges is said **simple** or **simplicial**. In a simple graph every edge is identified unambiguously with the pair of its endpoints. Edges should be formally considered as pairs of oppositely oriented **arcs**. A **path** is an alternating sequence of vertices and arcs such that every arc is preceded by its origin vertex and followed by the origin of its opposite arc. A path may have repeated vertices (beware that this is not standard, and usually called a **walk** in graph theory books). Two or more paths are **independent** if none contains an inner vertex of another. A **circuit** is a closed path, *i.e.* a path whose first and last vertex coincide. A **cycle** is a simple circuit (without repeated vertices). We will restrict to finite graphs for which V and E are finite sets.

The Euclidean distance in the plane \mathbb{R}^2 induces the usual topology where a subset $X \subset \mathbb{R}^2$ is **open** if every of its points is contained in a ball that is itself included in X . The **closure** \bar{X} of X is the set of limit points of sequences of points of X . The **interior** $\overset{\circ}{X}$ of X is the union of the open balls contained in X . An **embedding** of a non-loop edge in the plane is just a topological embedding (a homeomorphism onto its image) of the segment $[0, 1]$ into \mathbb{R}^2 . Likewise, an embedding of a loop-edge is an embedding of the circle $S^1 = \mathbb{R}/\mathbb{Z}$. An **embedding** of a finite graph $G = (V, E)$ in the plane is defined by a 1-1 map $V \hookrightarrow \mathbb{R}^2$ and, for each edge $e \in E$, by an embedding of e sending $\{0, 1\}$ to e 's endpoints such that the relative interior of e (the image of $]0, 1[$) is disjoint from other edge embeddings and vertices¹.

A graph is **planar** if it has an embedding into the plane. Thanks to the stereographic projection, the plane can be equivalently replaced by the sphere. A **plane graph** is a specific embedding of a planar graph. A connected plane graph in the plane has a single unbounded face. In contrast, all the faces play the same role in an embedding into the sphere and any face can be sent to the unbounded face of a plane embedding by a stereographic projection.

As far as planarity is concerned we can restrict to simple graphs. Indeed, it is easily seen that a graph has an embedding in the plane if and only if this is the case for the graph obtained by removing loop edges and replacing each multiple edge by a single edge. When each edge embedding $[0, 1] \rightarrow \mathbb{R}^2$ is piecewise linear the embedding is said **PL**, or **polygonal**. A **straight line embedding** corresponds to the case where each edge is a line segment.

Lemma 1.1. *A graph is planar if and only if it admits a PL embedding.*

¹In other words, this is a topological embedding of the quotient space $(V \sqcup [0, 1] \times E) / \sim$, where \sim identifies edge extremities $(0, e)$ and $(1, e)$ with the corresponding vertices.

The proof is left as an exercise. One can first show that a connected subset of the plane is connected by simple PL arcs.

1.1 The Jordan Curve Theorem

Most of the facts about planar graphs ultimately relies on the Jordan curve theorem, one of the most emblematic results in topology. Its statement is intuitively obvious: a simple closed curves cuts the plane into two connected parts. Its proof is nonetheless far from obvious, unless one appeals to more advanced arguments of algebraic topology. Camille Jordan (1838 – 1922) himself proposed a proof whose validity is still subject of debates [Hal07b]. A rather accessible proof was proposed by Helge Tverberg [Tve80] (see the course notes [Laz12] for a gentle introduction). Eventually, a formal proof was given by Thomas Hales (and other mathematicians) [Hal07b, Hal07a] and was automatically checked by a computer. Concerning the Jordan-Schoenflies theorem, the situation is even worse. This stronger version of the Jordan curve theorem asserts that a simple curve does not only cut a sphere into two pieces but that each piece is actually a topological disc. A nice proof by elementary means – but far from simple – and resorting to the fact that $K_{3,3}$ is not planar is due to Carsten Thomassen [Tho92].

The main source of difficulties in the proof of the Jordan curve theorem is that a continuous curve can be quite wild, e.g. fractal. When dealing with PL curves only, the theorem becomes much easier to prove.

Theorem 1.2 (Polygonal Jordan curve —). *Let C be a simple closed PL curve. Its complement $\mathbb{R}^2 \setminus C$ has two connected components, one of which is bounded and each of which has C as boundary.*

PROOF. Since C is contained in a compact ball, its complement has exactly one unbounded component. Define the *horizontal rightward* direction \vec{h} as some fixed direction transverse to the all the line segments of C . For every segment s of C we let \underline{s} be the lower half-open segment obtained from s by removing its upper endpoint. We also denote by h_p the ray with direction \vec{h} starting at a point $p \in \mathbb{R}^2$. We consider the *parity function* $\pi : \mathbb{R}^2 \setminus C \rightarrow \{\text{even, odd}\}$ that counts the parity of the number of lower half-open segments of C intersected by a ray:

$$\pi(p) := \text{parity of } |\{\text{segment } s \text{ of } C \mid h_p \cap \underline{s} \neq \emptyset\}|$$

Every $p \in \mathbb{R}^2 \setminus C$ is the center of small disk D_p over which π is constant. Indeed, let S_p be the set of segments (of C) that avoid h_p , let S'_p be the set of segments whose interior crosses h_p and let S''_p be the set of segments whose lower endpoint lies on h_p . If D_p is sufficiently small, then for every $q \in D_p$ we have $S_q = S_p$, $S'_q = S'_p$ and the parity of $|S''_q|$ and $|S''_p|$ is the same. See Figure 1. It follows that $\pi(q) = \pi(p)$. Since π is locally constant, it must be constant over each connected component of $\mathbb{R}^2 \setminus C$. Moreover, the parity function must take distinct values on points close to C that lie on a same horizontal but on each side of C . It follows that $\mathbb{R}^2 \setminus C$ has at least two components. To see that $\mathbb{R}^2 \setminus C$ has at most two components consider a small disk D centered at a point interior to a segment s of C . Then $D \setminus C = D \setminus s$ has two components. Moreover, any point in $\mathbb{R}^2 \setminus C$ can be joined to one of these components by a polygonal path that avoids C : first come

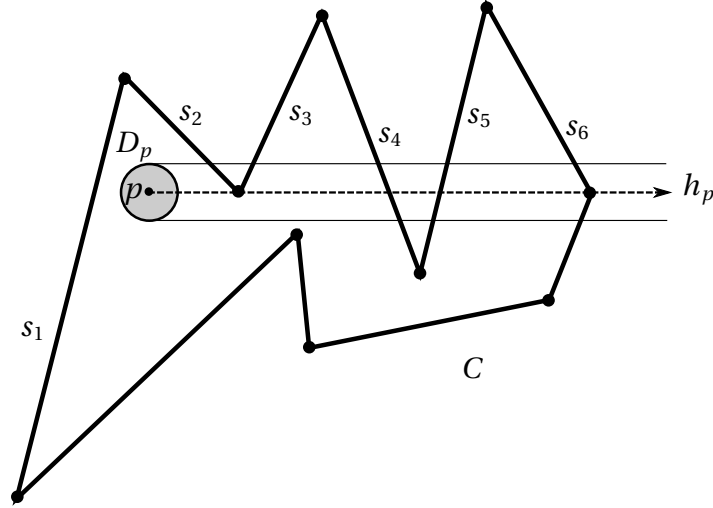


Figure 1: The horizontal ray through p cuts the five lower half-open segments $\underline{s}_2, \underline{s}_3, \underline{s}_4, \underline{s}_5, \underline{s}_6$. Here, we have $s_1 \in S_p$, $s_4, s_5 \in S'_p$ and $s_2, s_3, s_6 \in S''_p$.

close to C with a straight line and then follow C in parallel until D is reached. Finally, it is easily seen by similar arguments as above that every point of C is in the closure of both components of $\mathbb{R}^2 \setminus C$. \square

Corollary 1.3 (θ 's lemma). *Let C_1, C_2, C_3 be three simple PL paths with the same endpoints p, q and otherwise disjoint. The graph $G = C_1 \cup C_2 \cup C_3$ has three faces bounded by $C_1 \cup C_2$, $C_2 \cup C_3$ and $C_3 \cup C_1$, respectively.*

PROOF. From the Jordan curve theorem the three simple closed curves $G_k = C_i \cup C_j$, $\{i, j, k\} = \{1, 2, 3\}$, cut the plane into two components bounded by G_k . We let X_k and Y_k be respectively the bounded and unbounded component. We also denote by $\overset{\circ}{C}_i := C_i \setminus \{p, q\}$ the relative interior of C_i . We first remark that a simple PL path cuts an open connected subset of the plane into at most two components: as in the proof of the Jordan curve theorem we can first come close to the path and follow it until a small fixed disk is reached. Since $\overset{\circ}{C}_3$ is included in a face of G_3 , we deduce that $G = G_3 \cup \overset{\circ}{C}_3$ has at most three faces.

We claim that $\overset{\circ}{C}_i \subset X_i$ for at least one index $i \in \{1, 2, 3\}$. Otherwise we would have $C_i \subset \mathcal{L}X_i$, whence $G \subset \mathcal{L}X_i$, or equivalently $X_i \subset \mathcal{L}G$. So, X_i would be a face of G . Since the X_i 's are pairwise distinct (note that $C_i \subset \bar{X}_j$ while $\overset{\circ}{C}_i \not\subset \bar{X}_i$), we would infer that G has at least three bounded faces, hence at least four faces. This would contradict the first part of the proof. Without loss of generality we now assume $\overset{\circ}{C}_3 \subset X_3$.

From $G = G_1 \cup G_2$ we get that each face of G is a component of the intersection of a face of G_1 with a face of G_2 . From $G_3 \subset G \subset \mathcal{L}Y_3$ we get that Y_3 is a face of G . Since Y_3 is unbounded we must have $Y_3 \subset Y_1 \cap Y_2$.

Now, $C_1 \subset \bar{Y}_3 \subset \bar{Y}_1 = \mathcal{L}X_1$ implies $G = G_1 \cup C_1 \subset \mathcal{L}X_1$. It follows that X_1 is a face of G . Likewise, X_2 is a face of G . Moreover, these two faces are distinct (C_1 bounds X_1 but not X_2). We conclude that Y_3, X_1 and X_2 are the three faces of G . \square

1.2 Euler's Formula

The famous formula relating the number of vertices, edges and faces of a plane graph is credited to Leonhard Euler (1707-1783) although René Descartes had already deduced very close relations for the graph of a convex polyhedron. See the historical account of R. J. Wilson in [Jam99, Sec. 17.3] and in J. Erickson's course notes <http://jeffe.cs.illinois.edu/teaching/topology17/chapters/02-planar-graphs.pdf>

Recall that a graph G is **2-connected** if it contains at least three vertices and if removing any one of its vertices leaves a connected graph. If G is 2-connected, it can be constructed by iteratively adding paths to a cycle. In other words, there must be a sequence of graphs $G_0, G_1, \dots, G_k = G$ such that G_0 is a cycle and G_i is deduced from G_{i-1} by attaching a simple path between two distinct vertices of G_{i-1} .

Proposition 1.4. *Each face of a 2-connected PL plane graph is bounded by a cycle of the graph. Moreover, each edge is incident to (= is in the closure of) exactly two faces.*

PROOF. Let G be a 2-connected PL plane graph. Consider the sequence $G_0, G_1, \dots, G_k = G$ as above. We prove the proposition by induction on k . If $k = 0$, then G is a cycle and the proposition reduces to the Jordan curve theorem 1.2. Otherwise, by the induction hypothesis G_{k-1} satisfies the proposition. Let P be the attached path such that $G = G_{k-1} \cup P$. The relative interior of P must be contained in a face f of G_{k-1} . This face is bounded by a cycle C of G_{k-1} . We can now apply θ 's lemma 1.3 to $C \cup P$ and conclude that f is cut by G into two faces bounded by the cycles $C_1 \cup P$ and $C_2 \cup P$, where C_1, C_2 are the subpaths of C cut by the endpoints of P . Moreover all the other faces of G_{k-1} are faces of G bounded by the same cycles. It easily follows that the edges of G are each incident to exactly two faces. \square

Lemma 1.5. *Let G be a PL plane graph. If v is a vertex of degree one in G then $G - v$ and G have the same number of faces.*

PROOF. We denote by e the edge incident to v in G . Every face of G is contained in a face of $G - v$. Moreover, the relative interior of (the embedding of) e is contained in a face f of $G - v$. Hence, every other face of $G - v$ is also a face of G . It remains to count the number of faces of G in f . Let p, p' be two points in $f \setminus e$. There is a PL path in f connecting p and p' . This path may intersect e , but we may avoid this intersection by considering a detour in a small neighborhood N_e of e in f (indeed, $N_e \setminus e$ is connected). It follows that p and p' belong to a same component of $f \setminus e$. We conclude that G has only one face in f , so that G and $G - v$ have the same number of faces. \square

Theorem 1.6 (Euler's formula). *Let $|V|, |E|$ and $|F|$ be the number of vertices, edges and faces of a connected plane graph G . Then,*

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

PROOF. We argue by induction on $|E|$. If G has no edges then it has a single vertex and the above formula is trivial. Otherwise, suppose that G has a vertex v of degree one.

Then by Lemma 1.5, G has the same number of faces as $G - v$. Note that G has one vertex more and one edge more than $G - v$. By the induction hypothesis we can apply Euler's formula to $G - v$, from which we immediately infer the validity of Euler's formula for G . If every vertex of G has degree at least two, then G contains a cycle C . Let e be an edge of C . We claim that G has one face more than $G - e$. This will allow to conclude the theorem by applying Euler's formula to $G - e$, noting that G has the same number of vertices but one edge less than $G - e$. By the Jordan curve theorem 1.2, C cuts the plane into two faces (components) bounded by C . Since $G = C \cup (G - e)$, every face of G is included in the intersection of a face of C and a face of $G - e$. Let f be the face of $G - e$ containing the relative interior of e . Every other face of $G - e$ does not meet C , hence is also a face of G . Since f intersects the two faces of C (both bounded by e), G has at least one face more than $G - e$. By considering a small tubular neighborhood of e in f , one shows by an already seen argument that $f \setminus e$ has at most two components. It follows that f contains exactly two faces of G , which concludes the claim. \square

Application. Two old puzzles that go back at least to the nineteenth century are related to planarity and can be solved using Euler's formula. The first asks whether it is possible to divide a kingdom into five regions so that each region shares a frontier line with each of the four other regions. The second puzzle, sometimes called the *gaz-water-electricity problem* requires to join three houses to three gaz, water and electricity facilities using pipes so that no two pipes cross. By duality, the first puzzle translates to the question of the planarity of the **complete graph** K_5 obtained by connecting five vertices in all possible ways. The second problem reduces to the planarity of the **complete bipartite graph** $K_{3,3}$ obtained by connecting each of three independent vertices to each of three other independent vertices. It appears that these two puzzles are unfeasible.

Theorem 1.7. K_5 and $K_{3,3}$ are not planar.

PROOF. We give two proofs. The first one is based on Euler's formula.

1. Suppose by way of contradiction that $K_{3,3}$ has a plane embedding. Euler's formula directly implies that the embedding has $n = 2 - 6 + 9 = 5$ faces. Since $K_{3,3}$ is 2-connected, it follows from Proposition 1.4 that every edge is incident to two distinct faces. By the same proposition, each face is bounded by a cycle, hence by at least 4 edges (cycles in a bipartite graph have even lengths). It follows from the *handshaking lemma* that twice the number of edges is larger than four times the number of faces, *i.e.* $18 \geq 20$. A contradiction.

An analogous argument for K_5 implies that an embedding must have 7 faces. Since every face is incident to at least 3 edges, we infer that $2 \times 10 \geq 3 \times 7$. Another contradiction.

2. Let $\{1, 3, 5\}$ and $\{2, 4, 6\}$ be the two vertex parts of $K_{3,3}$. The cycle $(1, 2, 3, 4, 5, 6)$ separates the plane into two components in any plane embedding of $K_{3,3}$. By θ 's lemma the edges $(1, 4)$ and $(2, 5)$ must lie in the face that does not contain $(3, 6)$. Then $(1, 4)$ and $(2, 5)$ intersect, a contradiction. A similar argument applies for the non-planarity of K_5 .

□

Exercise 1.8. Every simple planar graph G with $n \geq 3$ vertices has at most $3n - 6$ edges and at most $2n - 4$ faces.

Exercise 1.9. Every simple planar graph with at least six vertices has a vertex with degree less than 6.

To conclude, we prove a very strong generalization of Exercise 1.8, which allows to quantify how non-planar dense graphs are. Here, a **drawing** of a graph is just a continuous map $f : G \rightarrow \mathbb{R}^2$, that is, a drawing of the graph on the plane where crossings are allowed. The **crossing number** $cr(G)$ of a graph is the minimal number of crossings over all possible drawings of G . For instance, $cr(G) = 0$ if and only if G is planar. The **crossing number inequality** [ACNS82, Lei84] provides the following lower bound on the crossing number.

Theorem 1.10. $cr(G) \geq \frac{|E|^3}{64|V|^2}$ if $|E| \geq 4|V|$.

The proof is a surprising application of (basic) probabilistic tools.

PROOF. Starting with a drawing of G with the minimal number of crossings, define a new graph G' obtained by removing one edge for each crossing. This graph is planar since we removed all the crossings, and it has at least $|E| - cr(G)$ edges (removing one edge may remove more than one crossing), so we obtain that $|E| - cr(G) \leq 3|V|$. (Note that we removed the -6 to obtain an inequality valid for any number of vertices.) This gives in turn

$$cr(G) \geq |E| - 3|V|.$$

This can be amplified in the following way. Starting from G , define another graph by removing vertices (and the edges adjacent to them) at random with some probability $1 - p < 1$, and denote by G'' the obtained graph. Taking the previous inequality with expectations, we obtain $\mathbb{E}(cr(G'')) \geq \mathbb{E}(|E''|) - 3\mathbb{E}(|V''|)$. Since vertices are removed with probability $1 - p$, we have $\mathbb{E}(|V''|) = p|V|$. An edge survives if and only if both its endpoints survive, and a crossing survives if and only if the four adjacent vertices survive (there may be less than four adjacent vertices in general, but not in the drawing minimizing the crossing number, we leave this as an exercise to check), so we get $\mathbb{E}(|E''|) = p^2|E|$ and $\mathbb{E}(cr(G'')) = p^4 cr(G)$. So we obtain

$$cr(G) \geq p^{-2}|E| - 3p^{-3}|V|,$$

and taking $p = 4|V|/|E|$ – which is less than 1 if $|E| \geq 4|V|$ – gives the result. □

2 Kuratowski's Theorem

2.1 The Subdivision Version

We say that H is **subdivision** of G if H is obtained by replacing the edges of G by independent simple paths of one or more edges. Obviously, a subdivision of a non-planar graph is also non-planar. It follows from Theorem 1.7 that a planar graph cannot

have a subdivision of K_5 or $K_{3,3}$ as a subgraph. In 1929, Kazimierz Kuratowski (1896 – 1980) succeeded to prove that this condition is actually sufficient for a graph to be planar. For this reason K_5 and $K_{3,3}$ are called the *Kuratowski graphs*, or the *forbidden graphs*.

Theorem 2.1 (Kuratowski, 1929). *A graph is planar if and only if it does not contain a subdivision of K_5 or $K_{3,3}$ as a subgraph.*

As just noted, we only need to show that a graph without any subdivision of a forbidden graph is planar. We follow the proof of Thomassen [MT01]. Recall that a graph is **3-connected** if it contains at least four vertices and if removing any two of its vertices leaves a connected graph. By Menger's theorem [Wil96, cor. 28.4], a graph is 3-connected if and only if any two distinct vertices can be connected by at least three independent paths. If e is an edge of a graph G we denote by $G//e$ the graph obtained by the **contraction** of e , *i.e.* by deleting e , identifying its endpoints, and merging each resulting multiple edge, if any, into a single edge. The proof of Kuratowski's theorem first restricts to 3-connected graphs. By Lemma 2.2 below we can repeatedly contract edges while maintaining the 3-connectivity until the graph is small enough so that it can be trivially embedded into the plane. We then undo the edge contractions one by one and construct corresponding embeddings. In the end, the existence of an embedding attests the planarity of the graph. In a second phase we extend the theorem to any graph, not necessarily 3-connected, that does not contain any subdivision of K_5 or $K_{3,3}$. This is done by adding as many edges as possible to the graph without introducing a (subdivision of a) forbidden graph. By Lemma 2.5 below the resulting graph is 3-connected and we may conclude with the first part of the proof.

Lemma 2.2. *Any 3-connected graph G with at least five vertices contains an edge e such that $G//e$ is 3-connected.*

PROOF. Suppose for the sake of contradiction that for any edge $e = xy$, the graph $G//e$ is not 3-connected. Denote by v_e the vertex of $G//e$ resulting from the identification of x and y . Then we can find a vertex $z \in V(G//e)$ such that $\{z, v_e\}$ disconnects $G//e$. In other words, for any edge $e = xy$ of G we can find a vertex $z \in V(G)$ such that $G - \{x, y, z\}$ is not connected. We choose e and z such that the number of vertices of the largest component, say H , of $G - \{x, y, z\}$ is maximal. Let u be adjacent to z in a component of $G - \{x, y, z\}$ other than H . See figure 2. By the above reformulation, we can find a vertex $v \in V(G)$ such that $G - \{z, u, v\}$ is not connected. We claim that the subgraph H' induced by $(V(H) \cup \{x, y\}) \setminus \{v\}$ is connected. Since H' is contained in $G - \{z, u, v\}$ and since H' has more vertices than H , this contradicts the choice of H , hence concludes the proof. To see that H' is connected we just need to check that every $t \in V(H)$ can be connected to x or y (themselves connected by e) by a path in H' . Since G is 3-connected, there is a path $p : t \rightsquigarrow x$ in G avoiding z and v . Replacing x by y if necessary, we can assume that $p - x$ does not contain y . It follows that $p - x$ is contained in $G - \{x, y, z, v\}$, hence in $H - v$. So p is in H' . \square

Exercise 2.3. Let e be an edge of G such that $G//e$ contains a subdivision of a forbidden graph. Show that G already contains such a subdivision. (Hint: $G//e$ and G need not contain subdivisions of the same forbidden graph.)

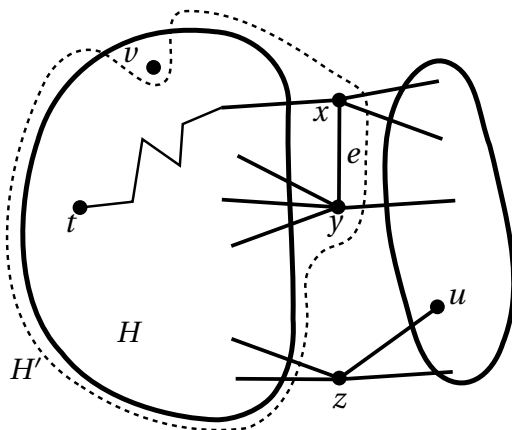


Figure 2: H' has more vertices than H even when $v \in \{x, y\}$.

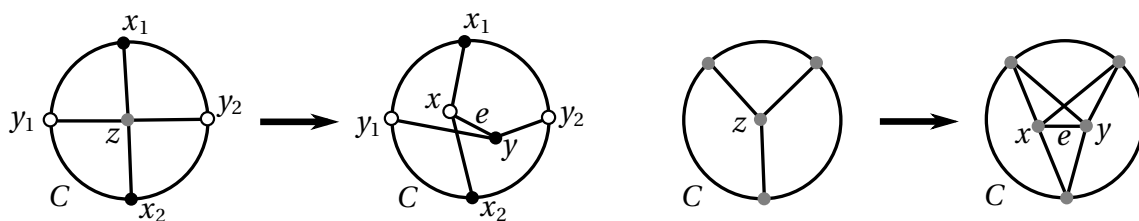


Figure 3: Left, two vertices $x_1, x_2 \in X$ and two vertices $y_1, y_2 \in Y$ appear in an alternate way along C . We infer the existence of a subdivision of $K_{3,3}$ in G . Right, X and Y share three vertices. We infer the existence of a subdivision of K_5 in G .

A straight line embedding is said **convex** if all its faces are bounded by convex polygons.

Proposition 2.4 (Kuratowski's theorem for 3-connected graphs). *A 3-connected graph G without any subgraph isomorphic to a subdivision of a forbidden graph admits a convex embedding.*

PROOF. We use induction on the number of vertices of G . The proposition is easily checked by hand if G has four vertices. Otherwise, G has at least five vertices, and by Lemma 2.2 we may choose an edge $e = xy$ such that $G' := G//e$ is 3-connected. Moreover, G' contains no subdivision of a forbidden graph. See Exercise 2.3. By induction, G' has a convex embedding. Let z be the vertex of G' resulting from the identification of x and y . Since $G' - z$ is 2-connected, we know by Proposition 1.4 that the face of $G' - z$ that contains z is bounded by a cycle C of G . Let $X = \{u \in V(G) \mid ux \in E(G)\}$ and let $Y = \{u \in V(G) \mid uy \in E(G)\}$. We claim that X and Y are *not interleaved* in C , i.e. $|X \cap Y| \leq 2$ and we cannot find two vertices in X and two vertices in Y that alternate along C . Otherwise, G would contain a subdivision of a forbidden graph as illustrated in Figure 3. We can obtain a convex embedding of G from the convex embedding of G' as follows: place x at the position of z and insert y close to x in the face of $G' - E_y$ incident to z and Y , where $E_y := \{zv \mid v \in Y\}$. We next connect x and y with line segments to their respective neighbors in X and Y , and finally x to y . The previous claim implies that the resulting straight line drawing of G is an embedding. It can easily be made convex using the fact that small perturbations of the vertices of a convex polygon leave the polygon convex. \square

The next lemma allows to extend the proposition to graphs that are not necessarily 3-connected and thus concludes the proof of Kuratowski's theorem.

Lemma 2.5. *Let G be a graph with at least four vertices, containing no subdivision of K_5 or $K_{3,3}$ and such that the addition of any edge between non-adjacent vertices creates such a subdivision. Then G is 3-connected.*

PROOF. We argue by induction on the number n of vertices of G . Note that for $n = 4$ the lemma just says that K_4 is 3-connected. Since removing an edge in K_5 leaves a planar graph, the lemma is also true for $n = 5$. We now assume $n \geq 6$. We claim that G is 2-connected. Otherwise, we could write $G = G_1 \cup G_2$ where G_1 and G_2 have a single common vertex x . Let $y_i \in G_i, i = 1, 2$, be adjacent to x . Adding the edge y_1y_2 creates a subdivision K of a forbidden graph. Since K_5 and $K_{3,3}$ are 3-connected and since x and y_1y_2 are the only connections between G_1 and G_2 , the vertices of K of degree ≥ 3 must lie all in G_1 or all in G_2 . Moreover, K must contain a path using both x and the edge y_1y_2 . The subpath between x and the edge y_1y_2 can be replaced by one of the two edges xy_1 or xy_2 to produce another subdivision of the same forbidden graph that does not use y_1y_2 , hence contained in G . This last contradiction proves the claim.

Suppose that G has two vertices x, y such that $G - \{x, y\}$ is not connected. We claim that xy is an edge of G . Otherwise, we could write $G = G_1 \cup G_2$ where G_1 and G_2 are connected and only have the vertices x, y in common. $G \cup xy$ must contain a subdivision K of a forbidden graph. As above, the degree three vertices of K must all lie in the same subgraph, say G_1 . We could then replace the edge xy in K with a path connecting x and y in G_2 to produce a subdivision of a forbidden graph contained in G_1 . We again reach a contradiction.

We now assume for a contradiction that G is not 3-connected and we let x, y be two vertices disconnecting G . By the previous claim, we may write $G = G_1 \cup G_2$ where $G_1 \cap G_2$ reduces to the edge xy . By the same type of arguments used in the above claims we see that adding an edge to G_i ($i = 1, 2$) creates a subdivision of a forbidden graph in the same G_i . We can thus apply the induction hypothesis and assume that each G_i is 3-connected, or has at most three vertices. By Proposition 2.4, both graphs are planar and we can choose a convex embedding for each of them. Let $z_i \neq x, y$ be a vertex of a face F_i of G_i bounded by xy . Note that F_i must be equal to the triangle z_ixy . (Otherwise, we could add an edge to G_i inside F_i to obtain a larger planar graph.) Adding the edge z_1z_2 to G creates a subdivision K of a forbidden graph. We shall show that some planar modification of G_1 or G_2 contains a subdivision of a forbidden graph, leading to a contradiction.

If all the vertices of K of degree ≥ 3 were in G_1 , we could replace the path of K in $G_2 + z_1z_2$ that uses z_1z_2 by one of the two edges z_1x or z_1y . We would get another subdivision of the same forbidden graph in G_1 . Likewise, G_2 cannot contain all the vertices of K of degree ≥ 3 . Furthermore, $V(G_1) \setminus \{x, y\}$ and $V(G_2) \setminus \{x, y\}$ cannot both contain two vertices of degree ≥ 3 in K since there would be four independent paths between them, although G_1 and G_2 are only connected through x, y and z_1z_2 in $G + z_1z_2$. For the same reason, K cannot be a subdivision of K_5 . Hence, K is a subdivision of $K_{3,3}$ and five of its degree three vertices are in the same G_i . Adding a point p inside F_i and drawing the three line segments px, py, pz_i we would obtain a planar embedding of

$G_i + \{px, py, pz_i\}$ that contains a subdivision of $K_{3,3}$. This last contradiction concludes the proof. \square

Corollary 2.6. *Every triangulation of the sphere with at least four vertices is 3-connected.*

PROOF. By Euler's formula it is seen that such a triangulation has a maximal number of edges. By the previous lemma, it must be 3-connected. \square

We end this section with a simple characterization of the faces of a 3-connected planar graph. A cycle of a graph G is **induced** if it is induced by its vertices, or equivalently if it has no chord in G . It is **separating** if the removal of its vertices disconnects G . The set of boundary edges of a face of a plane embedding is called a **facial cycle**.

Proposition 2.7. *The face boundaries of a 3-connected plane graph are its non-separating induced cycles.*

PROOF. Suppose that C is a non-separating induced cycle of a 3-connected plane graph G . By the Jordan curve theorem $\mathbb{R}^2 \setminus C$ has two components. Since C is non-separating one of the two components contains no vertices of G . This component is not cut by an edge since C has no chord. It is thus a face of G .

Conversely, consider a face f of G . By Proposition 1.4 this face is bounded by a cycle C . If C had a chord $e = xy$ then by the 3-connectivity of G there would be a path p connecting the two components of $C - \{x, y\}$. However, p and e being in the same component of $\mathbb{R}^2 \setminus C$ (other than f), they would cross by an application of θ 's lemma 1.3. Finally, consider two vertices x, y of $G - C$. They are connected by three independent paths. By θ 's lemma f is included in one of the three components cut by these paths and the boundary of this component is included in the corresponding two paths. Hence, C avoids the third path. It follows that $G - C$ is connected. \square

This proposition says that a planar 3-connected graph has essentially a unique plane embedding: if we realize the graph as a net of strings there are only two ways of dressing the sphere with this net; they correspond to the two orientations of the sphere.

2.2 The Minor Version

A **minor** of a graph G is any graph obtained from a subgraph of G by contracting a subset of its edges. In other words, a minor results from any sequence of contraction of edges, deletion of edges or deletion of vertices (in any order). Equivalently, H is a minor of G if the vertices of H can be put into correspondence with the trees of a forest in G and if every edge of H corresponds to a pair of trees connected by a (non-tree) edge (but all such pairs do not necessarily give rise to edges). Being a minor of another graph defines a partial order on the set of graphs. This partial order is the object of the famous graph minor theory developed by Robertson and Seymour and culminating in the proof of Wagner's conjecture that the minor relation is a **well-quasi-order**, *i.e.* that every infinite sequence of graphs contains two graphs such that the first appearing in the sequence is a minor of the other. As an easy consequence, every **minor closed**

family of graphs is characterized by a *finite* set of **excluded minors**. In other words, if a family of graphs contains all the minors of its graphs, then a graph is in the family if and only if none of its minors belongs to a certain finite set of graphs. The set of all planar graphs is the archetypal instance of a minor closed family. Its set of excluded minors happens to be precisely the two Kuratowski graphs.

Theorem 2.8 (Wagner, 1937). *A graph G is planar if and only if none of K_5 or $K_{3,3}$ is a minor of G .*

Remark that if G contains a subdivision of H , then H is a minor of G , but the converse is not true in general (think of a counter-example). We can nonetheless deduce Wagner's version from Kuratowski's theorem: the condition in Wagner's theorem is obviously necessary by noting that a minor of a planar graph is planar and by Theorem 1.7. The condition is also sufficient by the above remark and by Kuratowski's theorem. In fact, the equivalence between Wagner and Kuratowski's theorems can be shown by proving that a graph contains a subdivision of K_5 or $K_{3,3}$ if and only if K_5 or $K_{3,3}$ is a minor of this graph [Die05, Sec. 4.4].

3 Other Planarity Characterizations

We give some other planarity criteria demonstrating the fascinating interplay between Topology, Combinatorics and Algebra.

An **algebraic cycle** of a graph G is any subset of its edges that induces an **Eulerian subgraph**, *i.e.* a subgraph of G with vertices of even degrees². It is a simple exercise to prove that any algebraic cycle can be decomposed into a set of (simple) cycles in the usual acceptance. The set of (algebraic) cycles is given a group structure by defining the sum of two cycles as the symmetric difference of their edge sets. It can be considered as a vector space over the field $\mathbb{Z}/2\mathbb{Z}$ and is called the **cycle space**, denoted by $Z(G)$ (the letter Z is short for the German word for cycle, *Zyklus*). The cycle space of a tree is trivial. Also, the cycle space is the direct sum of the cycle spaces of the connected components of G . Given a spanning tree of G , each non-tree edge gives rise to a cycle by joining its endpoints by a path in the tree. It is not hard to prove that these cycles form a basis of the cycle space. Hence, when $G = (V, E)$ is connected,

$$\dim Z(G) = 1 - |V| + |E| \quad (1)$$

This number is sometimes called the **cyclomatic number** of G . A basis of the cycle space is a **2-basis** if every edge belongs to at most two cycles of the basis.

Theorem 3.1 (MacLane, 1936). *A graph G is planar if and only if $Z(G)$ admits a 2-basis.*

PROOF. It is not hard to prove that a graph that admits a 2-basis has a 2-basis composed of simple cycles only. See Exercise 3.2. Such a 2-basis must be the union of the

²An Eulerian subgraph in this sense is not necessarily connected.

2-bases of the blocks in the **block decomposition**³ of G . Moreover, G is planar if and only if its blocks are. We may thus assume that G has a single block, or equivalently that G is 2-connected.

Suppose that G is planar and consider the set B of boundaries of its bounded faces in a plane embedding. Every edge belongs to at most two such boundaries by Proposition 1.4. Furthermore, by the same proposition and the Jordan curve theorem, a simple cycle C of G is the sum of the boundaries of the faces included in the bounded region of C . Thus B generates $Z(G)$. Using Euler's formula, the number of bounded faces of G appears to be precisely $\dim Z(G)$. Hence, B is 2-basis.

For the reverse implication, suppose that G has a 2-basis. Note that it is equivalent that any subdivision of G admits a 2-basis. Moreover, $G - e$ has a 2-basis for any edge e : if e appears in two elements of the 2-basis replace these two elements by their sum, otherwise simply remove the basis element that contains e , if any. It follows that any subdivision of a subgraph of G has a 2-basis. We claim that none of the forbidden graphs can have a 2-basis, so that G is planar by Kuratowski's theorem. Indeed, assume the converse and let C_1, \dots, C_d be a 2-basis of a forbidden graph. The C_i 's being linearly independent, $\sum_i C_i$ is non-trivial hence contains at least 3 edges. It follows that $\sum_i |C_i| \leq 2|E| - 3$. From formula (1) we compute $\dim Z(K_{3,3}) = 4$. Since every cycle in a bipartite graph has length at least four, we have $\sum_{1 \leq i \leq 4} |C_i| \geq 4 \cdot 4 = 16$, in contradiction with $\sum_i |C_i| \leq 2 \cdot 9 - 3 = 15$. Similarly, we compute $\dim Z(K_5) = 6$, whence $\sum_{1 \leq i \leq 6} |C_i| \geq 6 \cdot 3 = 18$, in contradiction with $\sum_i |C_i| \leq 2 \cdot 10 - 3 = 17$. \square

Exercise 3.2. Show that a graph with a 2-basis admits a 2-basis whose elements are simple cycles. (Hint: Any algebraic cycle is a sum of edge-disjoint simple cycles. Try to minimize the total number of such simple cycles in the 2-basis.)

A **cut** in a graph $G = (V, E)$ is a partition of its vertices. A cut can be associated with the subset of edges with one endpoint in each part. Just as for the cycle space, the set of cuts can be given a vector space structure over $\mathbb{Z}/2\mathbb{Z}$ by defining the sum of two cuts as the symmetric difference of the associated edge sets. Equivalently, we observe that the sum of two cuts $\{V_1, V_2\}$ and $\{W_1, W_2\}$ is the cut $\{(V_1 \cap W_1) \cup (V_2 \cap W_2), (V_1 \cap W_2) \cup (V_2 \cap W_1)\}$. Remark that the cut space is generated by the **elementary cuts** of the form $\{v, V - v\}$, for $v \in V$. A cut is **minimal** if its edge set is not contained in the edge set of another cut. In a connected graph minimal cuts correspond to partitions both parts of which induce a connected subgraph. Such minimal cuts generate the cut space.

Given a plane graph G , we define its **geometric dual** G^* as the graph obtained by placing a vertex inside each face of G and connecting two such vertices if their faces share an edge in G . Note that distinct plane embeddings of a planar graph may give rise to non-isomorphic duals. When the plane graph G is connected, its vertex, edge and face sets are in 1-1 correspondence with the face, edge and vertex sets of G^* respectively. Note that the geometric dual of a plane tree has a single vertex, so that G^* may not be simple even if G is. It is not hard to prove that the set of edges of a (simple) cycle of G corresponds to a minimal cut in G^* . The converse is also true since the dual of the dual is the original graph.

³The blocks of G are its subgraphs induced by the classes of the following equivalence relation on its set of edges: $e \sim e'$ if there is a cycle in G that contains both e and e' .

For non-planar graphs the above construction is meaningless and we define an abstract notion of duality that applies in all cases. A graph G^* is an **abstract dual** of a graph G if the respective edge sets can be put in 1-1 correspondence so that every (simple) cycle in G corresponds to a minimal cut in G^* .

Theorem 3.3 (Whitney, 1933). *A graph is planar if and only if it has an abstract dual.*

The theorem can be proved by mimicking the proof of MacLane's theorem 3.1, first showing that if a graph has an abstract dual so does its subgraphs and subdivisions. We provide a shorter proof based on MacLane's theorem.

PROOF. The theorem can be easily reduced to the case of connected graphs. By the above discussion a geometric dual *is* an abstract dual, so that the condition is necessary. For the reverse implication, suppose that a graph G has an abstract dual G^* . The cycle space of G is generated by its simple cycles, hence by the dual edge sets of the minimal cuts of G^* . Those cuts are themselves generated by the elementary cuts. Clearly an edge appears in at most two elementary cuts (loop-edges do not appear in any cuts). It follows that the cycle space of G has a 2-basis, and we may conclude with MacLane's theorem. \square

We list below some other well-known characterizations of planarity without proof.

A strict **partial order** on a set S is a transitive, antisymmetric and irreflexive binary relation, usually denoted by $<$. Two distinct elements $x, y \in S$ such that either $x < y$ or $y < x$ are said **comparable**. A partial order is a **linear**, or **total**, order when all the elements are pairwise comparable. The **dimension** of a partial order is the minimum number of linear orders whose intersection (as binary relations) is the partial order. The **order complex** of a graph $G = (V, E)$ is the partial order on the set $S = V \cup E$ where the only relations are $v < e$ for v an endpoint of e .

Theorem 3.4 (Schnyder, 1989). *A graph is planar if and only if its order complex has dimension at most 3.*

See Mohar and Thomassen [MT01, p. 36] for more details. The **contact graph** of a family of interior disjoint disks in the plane is the graph whose vertices are the disks in the family and whose edges are the pairs of tangent disks.

Theorem 3.5 (Koebe-Andreev-Thurston). *A graph is planar if and only if it is the contact graph of a family of disks.*

Section 2.8 in [MT01] is devoted to this theorem and its extensions. A **3-polytope** is an intersection of half-spaces in \mathbb{R}^3 which is bounded and has non-empty interior. Its graph, or 1-skeleton, is the graph defined by its vertices and edges.

Theorem 3.6 (Steinitz, 1922). *A 3-connected graph is planar if and only if it is the graph of a 3-polytope.*

A proof can be found in the monograph by Ziegler [Zie95, Chap. 4]. We end this section with a nice and simple planarity criterion relying on a result by Hanani (1934) stating that any drawing of K_5 and of $K_{3,3}$ has a pair of independent edges with an odd number of crossings. (Recall that two edges are **independent** if they do not share any endpoint.) In fact, we have the stronger property that the number of pairs of independent edges crossing oddly is odd. This can be proved by first observing the property on a straight line drawing of K_5 (resp. $K_{3,3}$) and then deforming any other drawing to the given one using a sequence of elementary moves that preserve⁴ the parity of the number of oddly crossing pairs of independent edges. Together with Kuratowski's theorem, this proves the following

Theorem 3.7 (Hanani-Tutte). *A graph is planar if and only if it has a drawing in which every pair of independent edges crosses evenly.*

A weaker version of the theorem asks that every pair of edges, not necessarily independent, should cross evenly. See Mustafa's course notes for a geometric proof, not relying on Kuratowski's theorem.

4 Planarity Test

There is a long and fascinating story for the design of planarity tests, culminating with the first optimal linear time algorithm by Hopcroft and Tarjan [HT74] in 1974. Patrignani [Pat13] offers a nice and comprehensive survey on planarity testing. Although most of the linear time algorithms have actual implementations, they are rather complex and we only describe a simpler non optimal algorithm based on works of de Fraysseix and Rosenstiel [dFR85, Bra09]. We first recall that the block decomposition decomposes a connected graph into 2-connected subgraphs connected by trees in a tree structure. Hence, a graph is planar if and only if its blocks are planar. We can thus restrict the planarity test to 2-connected graphs. Note that the block decomposition of a graph can be computed in linear time using depth-first search. (See West [Wes01, p. 157].)

Also recall that a depth-first search in a graph discovers its vertices from a root vertex by following edges that form a spanning tree called a **depth-first search tree**. We say that a vertex v_1 of that tree is **higher** than another vertex v_2 if v_1 is a descendent of v_2 . The non-tree edges are called **back edges**. A back edge always connects a vertex to one that is lower in the depth-first search tree. The depth-first search induces an orientation of the tree edges directed from the root toward the leaves of the tree. The back edges are then directed from their highest toward their lowest vertex. Each back edge b defines an oriented **fundamental cycle**, $C(b)$, obtained by connecting its endpoints with the unique tree path between its target and source points. We write uv for an edge directed

⁴Those moves are of five types: (i) two edges locally (un)crossing and creating or canceling a bigon, (ii) an edge locally (un)crossing and creating or canceling a monogon, (iii) an edge passing over a crossing, (iv) an edge passing over a vertex, and (v) two consecutive edges around a vertex swapping their circular order. The three first moves are analogous to the Reidemeister moves performed on knot diagrams. (i), (ii), (iii) and (v) clearly preserve the number of oddly crossing pairs of independent edges. For (iv) we use the fact that for every vertex and every edge of K_5 or $K_{3,3}$ the edge is independent with an even number of the edges incident to the vertex.

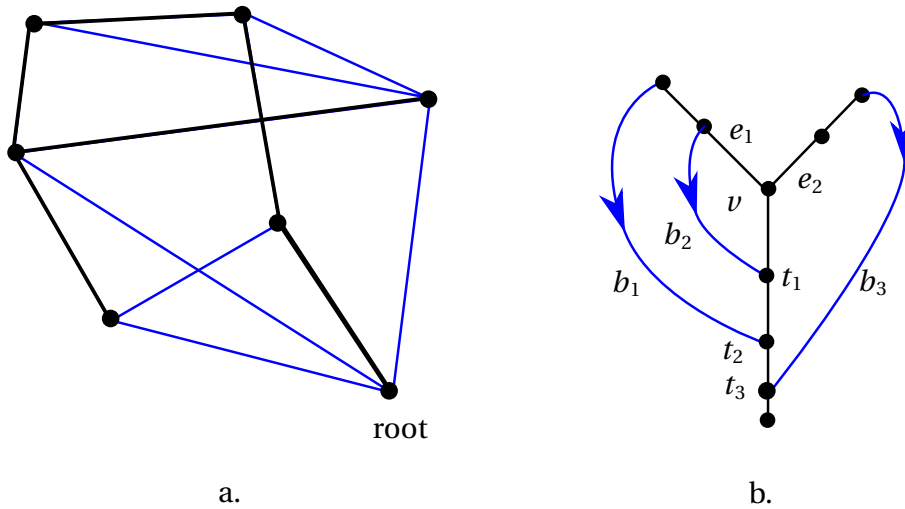


Figure 4: a. A graph (in blue) and a DFS tree in black. b. v is the branching point of a fork, b_1 and b_2 are two return edges for e_1 , b_3 is a return edge for e_2 and the lowpoint of e_1 is t_2 . The back edges b_1 and b_2 are left and the back edge b_3 is right.

from u to v . Two fundamental cycles may only intersect along a tree path, in which case the last edge uv along this path together with the **outgoing edges** vw_1 and vw_2 along the two cycles is called a **fork** with **branching point** v . A back edge vw is a **return edge** for itself and for every tree edge xy such that w is lower than x , and v is either higher than y or equal to⁵ y . The **return points** of an edge are the targets of its return edges. The **lowpoint** of the edge is its lowest return point, if any, or its source if none exists. The lowpoint of a back edge is thus its target point. We refer to Figure 4 for an illustration of all these concepts.

The idea of the planarity test is as follows. Suppose that a graph G has a plane embedding and consider a depth-first search tree of G . Without loss of generality, we may assume that the root is adjacent to the outer face of the plane embedding. The induced orientation of each fundamental cycle may appear clockwise or counterclockwise with respect to the embedding of G . A back edge is said **right** (with respect to the embedding) if its fundamental cycle is oriented clockwise, and **left** otherwise. Consider a fork with outgoing edges e_1, e_2 . They must have return edges since the graph is 2-connected. Then we have the following necessary conditions:

Fork condition:

1. All return edges of e_1 whose lowpoints are higher than the lowpoint of e_2 have fundamental cycles oriented the same way and
2. all return edges of e_2 whose lowpoints are higher than the lowpoint of e_1 have fundamental cycles oriented the other way.

⁵a vertex is neither higher nor lower than itself!

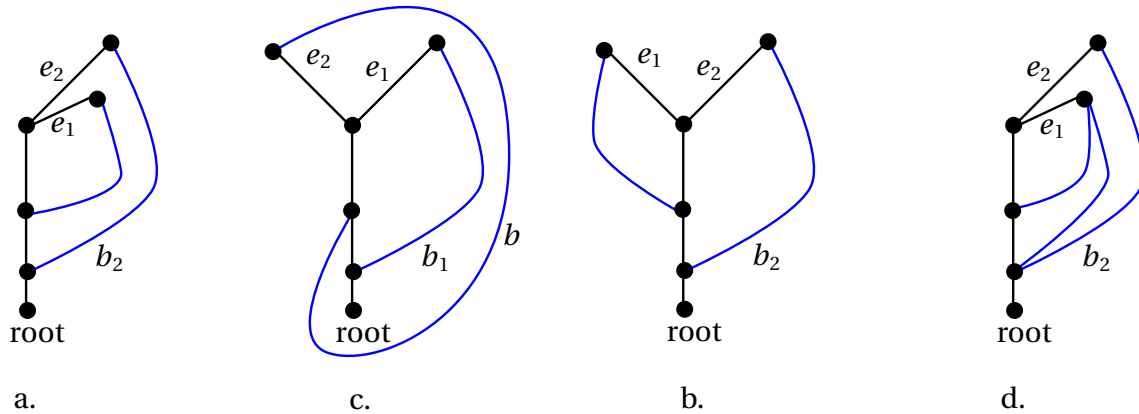


Figure 5: a. and c. The two cases occurring in the fork condition. b. A forbidden case. d. In this case, one chooses $e_2 \prec e_1$.

Lemma 4.1. *In a plane embedding, the orientations of the return edges satisfy the fork condition.*

PROOF. Let us denote by b_i the return edge having the same lowpoint as e_i . Then either the disks bounded by $C(b_1)$ and $C(b_2)$ have disjoint interior, or one is included in the other:

- In the latter case, swapping the indices 1 and 2 if necessary, we may assume that e_1 is inside $C(b_2)$. This is pictured in Figure 5a. Then any return edge of e_1 must also be inside the disk bounded by $C(b_2)$, and thus be oriented as b_2 . In particular, b_1 is oriented as b_2 and e_2 is outside $C(b_1)$. It follows that any return edge of e_2 must lie outside $C(b_1)$. Furthermore, a return edge b from e_2 having lowpoint higher than the one of b_1 must also lie outside $C(b_2)$, since otherwise $C(b)$ could not join its lowpoint without crossing $C(b_1)$. Now, the cycle $C(b)$ cannot contain the root in its interior as on Figure 5c, since the root is on the outer face. We infer that b is oriented oppositely to b_2 . The fork condition is thus satisfied.
- In the former case, any return edge b of e_1 must lie outside $C(b_2)$. See Figure 5b. If the lowpoint of b is higher than that of b_2 , then b must be oriented oppositely to b_2 , since $C(b)$ cannot contain the root in its interior. The mirror argument shows that a return edge from e_2 whose lowpoint is higher than the lowpoint of b_1 must be oriented oppositely to b_1 . We again conclude that the fork condition is satisfied.

□

An **LR partition** is a left-right assignment of the back edges such that the induced orientations of the fundamental cycles satisfy the fork condition for all the possible forks. The above lemma shows that a planar graph has an LR partition deduced from any particular plane embedding. As the following theorem shows, the existence of an LR partition happens to be sufficient for attesting planarity!

Theorem 4.2 (de Fraysseix and Rosenstiehl, 1985). *A connected graph G is planar if and only if it admits an LR partition with respect to some (and thus any) depth-first search tree.*

PROOF (SKETCH). Essentially, the proof starts by constructing a combinatorial embedding of G from the LR partition, *i.e.* a circular ordering of the edges around each vertex, then checking that this combinatorial embedding can indeed be realized in the plane without introducing crossings. Note that the fork conditions cannot involve back edges in different blocks in the block decomposition of G , so that we can assume G to be 2-connected by the above discussion. For each vertex v we define a total ordering \prec on its outgoing edges as follows. If v is the root, it can have only a single outgoing edge by the 2-connectivity of G and there is nothing to do. Otherwise, v has a unique incoming tree edge e and the total ordering will correspond to the circular clockwise ordering around v broken at e into a linear ordering. Let e_1, e_2 be two edges going out of v , and for $i = 1, 2$, let b_i be equal to e_i if e_i is a return edge, or a return edge of e_i with the lowest return point (there might be several ones) among its return edges. We need to decide if $e_1 \prec e_2$ or the opposite. The idea is that in any plane drawing of the graph, the ordering of e_1 and e_2 is enforced by the LR-assignment of b_1 and b_2 .

- If b_1 is a left back edge while b_2 is a right back edge, then we declare $e_1 \prec e_2$ since it must be the case in any plane drawing of G that respects the LR assignment (as in Figure 5b).
- If b_1 and b_2 are both right back edges we let $e_2 \prec e_1$ if either the lowpoint of b_2 is lower than the lowpoint b_1 (as in Figure 5a), or if e_1 has another right return edge towards another return point (as in Figure 5d). By the fork condition, it is impossible for both e_1 and e_2 to have another right return edge towards another return point, so this is well-defined.
- If b_1 and b_2 are both left back edges, the previous situation leads to the opposite decision.
- If none of this applies, we order them arbitrarily.

There remains to include the incoming return edges in this ordering. Let $e_1 \prec e_2 \prec \dots \prec e_\ell$ be the resulting ordering of the edges going out of v . We denote by $L(e_i)$ and $R(e_i)$ the left and right incoming back edges whose source points are in the subtree rooted at the target of e_i , or equivalently whose fundamental cycles contain e_i . We order the elements of $L(e_i)$ as follows: we let $b_1 \prec b_2$ if and only if the fork of their cycles $C(b_1)$ and $C(b_2)$ has outgoing edges $a_2 \prec a_1$. An analogous ordering is defined for $R(e_i)$. We finally concatenate all those orderings as follows, the rationale is pictured in Figure 6:

$$L(e_1) \prec e_1 \prec R(e_1) \prec L(e_2) \prec e_2 \prec \dots \prec L(e_\ell) \prec e_\ell \prec R(e_\ell)$$

For the root vertex we define the ordering $L(e) \prec e \prec R(e)$ where e is the unique outgoing edge of the root and $L(e), R(e)$ and their ordering are defined similarly as above. It remains to prove that the computed orderings define a planar combinatorial embedding. To this end, we first embed the depth-first search tree into the plane by

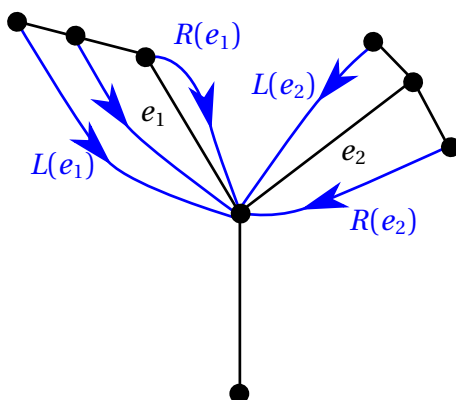


Figure 6: Ordering the incoming return edges.

respecting the computed orderings. This is obviously always possible. We then insert a small initial and final piece for each back edge in its place while respecting the circular orderings and without introducing crossings. Consider a simple closed curve C that goes along the embedding of the depth-first search tree, staying close to it. Each inserted back edge piece intersects C in a single point. Those points are paired according to the back edge to which they belong. We claim that the constructed orderings are such that the list of intersections along C is a well parenthesized expression. To see this we just need to prove that any two pairs of points appear in the good order (not interlaced) along C . There are two cases to consider: the pair corresponds to back edges, say b_1, b_2 , that are either on the same side, or on opposite sides. Suppose for instance that b_1 and b_2 are both right edges. If they have the same lowpoint then the constructed orderings implies that their initial and final pieces indeed appear in the good order along C . Similar arguments hold for the other cases. It follows from the claim that we can connect all the paired pieces without introducing crossings, thus proving that G has a plane embedding. \square

In order to test if G has an LR partition we can first compute a constraint graph whose nodes are the back edges and whose links are 2-colored constraints: the blue links connect nodes that must be on the same side and the red links connect nodes that must lie on opposite sides. All the links are obtained from the fork conditions. This graph can easily be constructed in quadratic time with respect to the number of edges of G . It remains to contract the blue links and check if the resulting constraint graph is bipartite to decide if G has an LR partition or not. This can clearly be done in quadratic time.

5 Drawing with Straight Lines

Proposition 2.4 together with Lemma 2.5 show that every planar graph has a straight line embedding. One of the oldest proof of existence of straight line embeddings is credited to Fáry [Fár48] (or Wagner, 1936) and does not rely on Kuratowski's theorem. By adding edges if necessary we can assume given a maximally planar graph G , so that adding any other edge yields a non-planar graph. Every embedding of G is thus a triangulation, since otherwise we could add more edges without breaking the planarity.

We show by induction on the number of vertices that any (topological) embedding of G can be realized with straight lines. Choose one embedding. By Euler's formula, G has a vertex v of degree at most 5 that is not a vertex of the unbounded face (triangle) of the embedding. Consider the plane triangulation H obtained from that of G by first deleting v and then adding edges (at most two) to triangulate the face of $G - v$ that contains v in its interior. By the induction hypothesis, H can be realized with straight lines. We now remove the at most two edges that were added and embed v in the resulting face. Since the face is composed of at most 5 edges, it must be star-shaped and we can put v in its center to join it with line segments to the vertices of the face. We obtain this way a straight line embedding of G .

There is another proof of Proposition 2.4 due to Tutte [Tut63] that actually provides an algorithm to explicitly compute a convex embedding of any 3-connected planar graph $G = (V, E)$. The algorithm can be interpreted by a physical spring-mass system. Consider a facial cycle C of G (recall that those are determined by Proposition 2.7) and nail its vertices in some strictly convex positions onto a plane. Connect every other vertex of G , considered as a punctual mass, to its neighbors by means of springs. Now, relax the system until it reaches the equilibrium. The final position provides a convex embedding! The system equilibrium corresponds to a state with minimal kinetic energy. By differentiating this energy one easily gets a linear system of equations where each *internal* vertex in $V_I := V \setminus V(C)$ is expressed as the barycenter of its neighbors. The barycentric coefficients are the stiffnesses of the springs. In practice, we associate with every edge e in $E \setminus E(C)$ a positive weight (stiffness) λ_e . In fact, if u and v are neighbor vertices it is not necessary that $\lambda_{uv} = \lambda_{vu}$. One may use "oriented" stiffness. Formally, we have

Theorem 5.1 (Tutte, 1963). *Every strictly convex embedding of the vertices of C extends to a unique map $\tau : V \rightarrow \mathbb{R}^2$ such that for every internal vertex v , its image $\tau(v)$ is the convex combination of the image of its neighbors $N(v)$ with weights λ_{vw} , for $w \in N(v)$:*

$$\forall v \in V_I, \quad \sum_{w \in N(v)} \lambda_{vw} (\tau(v) - \tau(w)) = 0. \quad (2)$$

Moreover, τ induces a convex embedding of G by connecting the images of every pair of neighbor vertices with line segments.

For conciseness, we number the vertices in V_I from 1 to k and the vertices of C from $k + 1$ to n (hence, $k = |V_I|$ and $n = |V|$). We also write λ_{ij} for the weight of edge ij and denote by $N(i)$ the set of neighbors of vertex i . We finally put $\lambda_{ij} = 0$ for $j \notin N(i)$. We follow the proof from [RG96] and from the course notes of Éric Colin de Verdière <http://www.di.ens.fr/~colin/cours/all-algo-embedded-graphs.pdf>.

Lemma 5.2. *If G is connected, the system (2) has a unique solution.*

PROOF. (2) can be written

$$\Lambda \begin{bmatrix} \tau_1 \\ \vdots \\ \tau_k \end{bmatrix} = \begin{bmatrix} \sum_{j>k} \lambda_{1j} \tau_j \\ \vdots \\ \sum_{j>k} \lambda_{kj} \tau_j \end{bmatrix}$$

where τ_i stands for $\tau(i)$ and

$$\Lambda = \begin{bmatrix} \sum_{j=1}^n \lambda_{1j} & -\lambda_{12} & \dots & -\lambda_{1k} \\ -\lambda_{21} & \sum_{j=1}^n \lambda_{2j} & \dots & -\lambda_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ -\lambda_{k1} & -\lambda_{k2} & \dots & \sum_{j=1}^n \lambda_{kj} \end{bmatrix}$$

We need to prove that Λ is invertible. Let $x \in \mathbb{R}^k$ such that $\Lambda x = 0$ and let x_i be one of its components with maximal absolute value. We set $x_{k+1} = x_{k+2} = \dots = x_n = 0$. Since $(\Lambda x)_i = \sum_{j \in N(i)} \lambda_{ij}(x_i - x_j) = 0$ and $\lambda_{ij} > 0$ for $j \in N(i)$ we infer that $x_j = x_i$ for $j \in N(i)$. By the connectivity of G , all the x_j , $j = 1, \dots, n$, are null. We conclude that Λ is non-singular. \square

In the sequel, we refer to τ as **Tutte's embedding**. We also assume once and for all that G is 3-connected.

Remark 5.3. Since the weights are positive the Tutte embedding of every internal vertex is in the relative interior of the convex hull of its neighbors. In particular, this remains true for the projection of the vertex and its neighbors on any affine line.

We shall derive a maximal principal from this simple remark. Let K be a cycle of G . By Proposition 2.7, G has a unique embedding on the sphere (up to change of orientation) and its faces can be partitioned into two families corresponding to the two connected components of the complement of K . The vertices of $G - K$ incident to a face in the part that does not contain C are said **interior to K** .

Lemma 5.4 (Maximum principle). *Let h be a non-constant affine form over \mathbb{R}^2 such that the Tutte embedding of K is included in the half-plane $\{h \leq 0\}$ and such that at most two vertices of K are on the line $\{h = 0\}$. Then each vertex v interior to K satisfies $h(\tau(v)) < 0$.*

PROOF. Consider a vertex v interior to K that maximizes h and suppose for a contradiction that $h(\tau(v)) \geq 0$. Let H be the subgraph of G induced by the vertices interior to K and let H_v be the component of v in H . By the above Remark 5.3, all the neighbors $w \in N(v)$, which are either interior to K or on K , must satisfy $h(\tau(w)) = h(\tau(v))$. Hence, the Tutte embedding of H_v is included in $\{h \geq 0\}$. Since G is 3-connected, H_v must be attached to K by at least three vertices. These attachment vertices are embedded in $\{h \leq 0\}$ and at least one of them, call it u , is embedded in $\{h < 0\}$ since at most two are on $\{h = 0\}$. Remark 5.3 applied to any vertex of H_v adjacent to u then leads to a contradiction. \square

Corollary 5.5. *The Tutte embedding of every internal vertex lies in the interior of the convex hull of the given strictly convex embedding of C .*

PROOF. By the maximum principle, every half-plane that contains C contains the interior vertices in its interior. \square

Let h be a nonzero linear form over \mathbb{R}^2 . A vertex of G whose Tutte's embedding is aligned with the Tutte embedding of its neighbors in the direction of the kernel of h is said h -passive, and h -active otherwise.

Lemma 5.6. *Let h be a non-trivial linear form and let v be an h -active interior vertex. G contains two paths $U(v, h)$ and $D(v, h)$ such that*

1. $U(v, h) := v_0, v_1, \dots, v_b$ joins $v = v_0$ to a vertex v_b of C and h is strictly increasing along $U(v, h)$, i.e. $h(\tau(v_{j+1})) > h(\tau(v_j))$ for $1 \leq j < b$.
2. $D(s, h)$ joins v to a vertex of C and h is strictly decreasing along $D(s, h)$.

PROOF. Since v is h -active, Remark 5.3 implies the existence of some neighbor w with $h(\tau(w)) > h(\tau(v))$. If this neighbor is on C then we may set $U(v, h) = vw$. Otherwise, w is itself h -active and we can repeat the process until we reach a vertex of C , thus defining the path $U(v, h)$. An analogous construction holds for the downward path $D(s, h)$. \square

Lemma 5.7. *For every non-trivial linear form h , all the interior vertices are h -active.*

PROOF. By way of contradiction, suppose that some interior vertex v is h -passive. By Lemma 5.5, some vertex w of C satisfies $h(\tau(w)) > h(\tau(v))$. Since G is 3-connected, we can choose three independent paths P_1, P_2, P_3 from v to w . For $i = 1, 2, 3$, let Q_i be the initial segment of P_i from v to the first h -active vertex w_i along P_i . Remark that Q_i has at least one edge and that it is contained in the line $\{h = h(\tau(v))\}$. By Lemma 5.6, we can choose two paths $U(w_i, h)$ and $D(w_i, h)$ from w_i to vertices on C . By the preceding remark, the three paths $Q_i, U(w_i, h)$ and $D(w_i, h)$ are pairwise disjoint except at w_i . Using that Q_1, Q_2, Q_3 only share their initial vertex v , it is easily seen that $C \cup_{i=1,2,3} (Q_i \cup U(w_i, h) \cup D(w_i, h))$ contains a subdivision of $K_{3,3}$, in contradiction with Kuratowski's theorem. \square

Recall that G is supposed to have a plane embedding with facial cycle C . Using the stereographic projection if necessary, we may assume that C is the facial cycle of the unbounded face. We temporarily assume that all facial cycles of G , except possibly C , are triangles.

Lemma 5.8. *Let uvx and uvy be the two facial triangles incident to an edge uv of G not in C , then $\tau(x)$ and $\tau(y)$ are on either sides of any line through $\tau(u)$ and $\tau(v)$.*

PROOF. Note that we do not assume $\tau(v) \neq \tau(u)$ in the lemma. Let h be a linear form whose kernel has the direction of a line ℓ through $\tau(u)$ and $\tau(v)$. By Lemma 5.6, there are two paths $U(u, h)$ and $U(v, h)$ embedded strictly above ℓ connecting respectively u and v to C . We can extract from $\{uv\} \cup U(u, h) \cup U(v, h)$ a cycle K above ℓ with only u and v on ℓ . By the maximum principle, all the vertices interior to K are embedded strictly above ℓ . One of the faces uvx and uvy must be contained in the interior of K , so that either $\tau(x)$ or $\tau(y)$ is strictly above ℓ . An analogous argument using $D(u, h)$ and $D(v, h)$ shows that one of $\tau(x)$ or $\tau(y)$ is strictly below ℓ . \square

Corollary 5.9. *All facial triangles uvw are non-degenerate, i.e. $\tau(u)$, $\tau(v)$ and $\tau(w)$ are pairwise distinct.*

PROOF. By Corollary 5.5 all the triangles with an edge in C are non-degenerate. By Lemma 5.8 all their adjacent triangles are themselves non-degenerate and by connectivity of the dual graph, all the triangles are non-degenerate. \square

Corollary 5.10. *If all the facial triangles other than C are triangles the Tutte embedding indeed induces a straight line embedding of G .*

PROOF. Since all the facial triangles are non-degenerate, it is enough to prove that their embeddings have pairwise disjoint interiors. Let p be a point contained in the interior of the embedding of some triangle t . Consider a ray r issued from p that avoids all the embeddings $\tau(V)$ of the vertices of G . This half-line crosses some edge e_0 of $t_0 := t$. By Lemma 5.9 the other triangle t_1 incident to e_0 crosses r on the other side of t_0 , away from p . In turn, r crosses another edge e_1 of t_1 and we define t_2 as the other incident triangle. This way we define a sequence of interior disjoint triangles t_0, t_1, \dots, t_i and edges e_0, e_1, \dots, e_i all crossed by r , each time further away from p until we hit C , i.e. e_i belongs to C . Remark that t_i only depends on r as it is the unique triangle incident to the intersection of r and C . Let t' be another triangle that contains p in its interior. It gives rise to another sequence $t'_0 = t', t'_1, \dots, t'_j$ of triangles crossed by r . By the preceding remark, $t_i = t_j$. Since the preceding triangles are defined unambiguously, we conclude that the two sequences are equal. In particular $t = t'$. \square

PROOF OF TUTTE'S THEOREM. This last corollary concludes the proof of Tutte's theorem 5.1 when all facial cycles other than C are triangles. When this is not the case, we can triangulate the faces other than C , adding $m - 3$ edges in each face of length m to obtain a planar graph G' with the above property. It is possible to put weights on the edges of G' , including those of G , so that the solution for system (2) written for G' is the same as for the initial system for G . This is a consequence of Remark 5.3 and of the next exercise. By Corollary 5.10, Tutte's embedding provides a straight line embedding of G' . Removing the extra edges, we obtain a straight line embedding of G . It remains to observe that each face of this embedding is convex since by Lemma 5.7 and Remark 5.3, the angle at every vertex of a face is smaller than π . \square

Exercise 5.11. Let p be a point interior to the convex hull of a finite point set P . Show that p is a convex combination of the points of P with strictly positive coefficients only. (Hint: the convex hull of $P \cup \{p\}$ is star-shaped with respect to p .)

One may wonder whether the barycentric method of Tutte could be extended in three dimensions in order to embed a triangulated 3-ball given a convex embedding of its boundary. However, É. Colin de Verdière et al. gave counterexamples to such an extension showing that expressing each interior vertex as the barycenter of its neighbors does not always yield an embedding [CdVPV03].

References

- [ACNS82] Miklós Ajtai, Vašek Chvátal, Monroe M Newborn, and Endre Szemerédi. Crossing-free subgraphs. *North-Holland Mathematics Studies*, 60:9–12, 1982.
- [Bra09] Ulrik Brandes. The left-right planarity test. <http://www.inf.uni-konstanz.de/algo/publications/b-lrpt-sub.pdf>, 2009.
- [CdVPV03] Éric Colin de Verdière, Michel Pocchiola, and Gert Vegter. Tutte’s barycenter method applied to isotopies. *Computational Geometry: Theory and Applications*, 26(1):81–97, 2003.
- [dFR85] Hubert de Fraysseix and Pierre Rosenstiehl. A characterization of planar graphs by trémaux orders. *Combinatorica*, 5(2):127–135, 1985.
- [Die05] Reinhard Diestel. *Graph theory*. Springer Verlag, 2005.
- [Fár48] István Fáry. On straight line representation of planar graphs. *Acta sci. math. (Szeged)*, 1:229–233, 1948.
- [Hal07a] Thomas C. Hales. The Jordan Curve Theorem, Formally and Informally. *The American Mathematical Monthly*, 114:882–894, dec 2007.
- [Hal07b] Thomas C. Hales. Jordan’s Proof of the Jordan Curve Theorem. *STUDIES IN LOGIC, GRAMMAR AND RHETORIC*, 10(23):45–60, 2007.
- [HT74] J. Hopcroft and R. Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974.
- [Jam99] I. M. James, editor. *History of topology*. Elsevier, 1999.
- [Laz12] Francis Lazarus. Géométrie algorithmique. Notes de cours. <http://www.gipsa-lab.fr/~francis.lazarus/Enseignement/geoAlgo.html>, 2012.
- [Lei84] Frank Thomson Leighton. New lower bound techniques for vlsi. *Mathematical Systems Theory*, 17(1):47–70, 1984.
- [MT01] Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. John Hopkins University Press, 2001.
- [Pat13] Maurizio Patrignani. *Handbook of graph drawing and visualization*, Roberto Tamassia editor, chapter Planarity testing and embedding, pages 1–42. CRC press, 2013. <https://cs.brown.edu/rt/gdhandbook/>.
- [RG96] Jürgen Richter-Gebert. *Realization spaces of Polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer, Berlin, 1996.
- [Tho92] Carsten Thomassen. The Jordan-Schönflies Theorem and the Classification of Surfaces. *American Mathematical Monthly*, pages 116–130, feb 1992.

- [Tut63] William T. Tutte. How to Draw a Graph. *Proc. London Mathematical Society*, 13:743–768, 1963.
- [Tve80] Helge Tverberg. A Proof of the Jordan Curve Theorem. *Bull. London Math. Soc.*, 12:34–38, 1980.
- [Wes01] Douglas B. West. *Introduction to Graph Theory*. Prentice Hall, second edition, 2001.
- [Wil96] Robin J. Wilson. *Introduction to Graph Theory*. Prentice Hall, fourth edition, 1996.
- [Zie95] Günter M. Ziegler. *Lectures on polytopes*. Number 152 in Graduate texts in mathematics. Springer, rev. first ed edition, 1995.