

anyt
[entry]anyt/global/

Chapter 13

Planar graphs and testing for planarity

By Sarel Har-Peled, April 8, 2018^①

Version: 0.1

This is an early draft of a new chapter. Read at your own peril.

At an archaeological site I saw fragments of precious vessels, well cleaned and groomed and oiled and spoiled. And beside it I saw a heap of discarded dust which wasn't even good for thorns and thistles to grow on. I asked: What is this gray dust which has been pushed around and sifted and tortured and then thrown away? I answered in my heart: This dust is people like us, who during their lifetime lived separated from copper and gold and marble stones and all other precious things - and they remained so in death. We are this heap of dust, our bodies, our souls, all the words in our mouths, all hopes.

In this chapter, we introduce planar graphs and review some standard results about them. We also present an algorithm for testing if a graph is planar. At an archaeological site, Yehuda Amichai

13.1. Definitions and some basic results

13.1.1. Background – What is a curve?

The notion of a curve drawn in the plane is quite natural, but it turns out to be surprisingly challenging to define formally. Indeed, the natural definition of a curve is a continuous one to one mapping from $[0, 1]$ to a set in the plane (i.e., the curve), but this definition also includes space filling curves, which are definitely do not capture our intuitive definition of a curve. Specifically, the Peano (or Hilbert) space-filling curve is a continuous(!) mapping from $[0, 1]$ to the unit square. See [Figure 13.1](#).

To avoid this pitfall, a *closed Jordan curve* is a closed curve, that does not self intersect, that can be continuously deformed into a circle. Similarly, a (regular) Jordan curve (or arc) is a curve that can be continuously deformed into a segment^② in the plane. More formally, a closed Jordan curve is a homeomorphism f from a circle to a set in the plane. A mapping f is a *homeomorphism*, if it is continuous and it has a continuous inverse function. As such, space filling curves are not Jordan curves, as one can find points that are arbitrarily close to each other in the image, that are far away from each other in the original range.

This leads to the following famous theorem, which is intuitively obvious, but proving it turns out to be challenging (see bibliographical notes for more details).

Theorem 13.1.1 (Jordan curve theorem). *A closed Jordan curve J partition the plane into two open connected components – the interior and the exterior. Any curve connecting a point in the interior, to a point in the exterior must intersect J .*

^①This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

^②A *segment* is the portion of a straight line connecting two points on it.




Figure 13.1: An inductive definition of Hilbert's space filling curve. The figure is taken from Wikipedia.

13.1.2. Planar graphs – a review

Notations. In the following, *component* refers to a connected component of a graph. A graph H is a *subgraph* of a given graph $G = (V, E)$, denoted by $H \subseteq G$, if $V(H) \subseteq V$, and $E(H) \subseteq E$. For a set of vertices $U \subseteq V(G)$, its *induced subgraph* of G is the graph $G_U = (U, E_U)$, formed by keeping only the edges of G with both their endpoints in U . Formally, we have $E_U = \{uv \in E(G) \mid u, v \in U\}$.

For a subgraph $H \subseteq G$, we denote by $G \setminus H$ the induced subgraph of G over $V \setminus V(H)$.

We abuse notations with impunity and without shame. For a vertex $v \in V(G)$, we denote by $G - v$ the graph resulting from removing v and the edges adjacent to it. Similarly, for a set X and an element x , we use $X - x$ as a shortcut for $X \setminus \{x\}$, and similarly $X + x = X \cup \{x\}$.

We provide a short introduction to planar graphs, and state some standard properties without proof (see bibliographical notes for relevant references).

A graph $G = (V, E)$ is *planar* if it can be drawn in the plane, such that every vertex is a point, and an edge is a Jordan curve connecting the corresponding points. Furthermore, the curves that corresponds to two different edges intersect only in their common endpoint (if they have one). Given a planar graph, we do not necessarily have this *embedding* of the graph (i.e., an explicit description of the curve forming each edge, and the location of each vertex). Instead, a (symbolic) representation of an embedding, would describe the faces, edges and vertices and their relationships. See Section 13.1.3 below for details.

Lemma 13.1.2 (Euler's formula). *In a connected planar graph with n vertices, m edges, and f faces, we have $n - m + f = 2$.*

Remark. Here, we usually deal with *simple graphs* that do not have self loops or parallel edges. However, Euler's formula holds even if one allows self-loops, and parallel edges.

A planar graph is *maximal* if no edge can be added to it without violating its planarity or simplicity (i.e., we are not allowing parallel edges or self-loops). In such a maximal planar graph every face has exactly three boundary edges (i.e., the graph is *triangulated*).

Definition 13.1.3. A maximal planar graph is a *triangulation*. In any embedding of a triangulation, all its faces, including the outer face, are triangles (i.e., the boundary of a face is a cycle with three edges).

Lemma 13.1.4. *Given a (simple) planar graph $G = (V, E)$, one can add edges to it so that it becomes a triangulation (i.e., all its faces are triangles).*

Proof: If G is not connected, then add edges to it between different connected components till it becomes connected. Clearly, this can not violate planarity.

Next, fix an embedding of G . Consider a face f of G that is not a triangle. By connectivity, the face f has only a single boundary cycle C (assume it is the outer boundary), and C has more than three edges. Let $C = \langle u_1, \dots, u_k \rangle$ be the vertices encountered when following C in counterclockwise

direction (vertices might appear several times in this cyclical sequence). If $u_1 \neq u_3$, then connect them by an edge (they cannot be adjacent in C since $k \geq 4$). Otherwise $u_1 = u_3$. If $u_2 = u_4$ then the edge $u_1u_2 = u_2u_3 = u_3u_4$, which is impossible as a boundary cycle can use an edge at most twice. As such $u_2 \neq u_4$. This implies that one can add the edge u_2u_4 to the graph.

We repeat this process till all faces of G are triangles, and it is thus a triangulation. ■

Lemma 13.1.5. *A simple planar graph G with n vertices has at most $3n - 6$ edges and at most $2n - 4$ faces. A triangulation has exactly $3n - 6$ edges and $2n - 4$ faces.*

Proof: We add edges to G till it becomes a triangulation, see Lemma 13.1.4. Now, every face is a triangle, and as such the number of edges incident to faces is $3f$, where f is the number of faces of G . Similarly, every edge is incident to 2 faces, and as such the number of edges incident to faces is $2m$, where $m = |E(G)|$. We conclude that $2m = 3f$, and furthermore, Euler's formula states that $f - m + n = 2$ (recall, that we also count the outer face of the planar graph as a face). This implies that $(2/3)m - m + n = 2$, which implies that $m = 3n - 6$, and this is the maximum number of edges that any planar graph with n vertices might have. This also implies that in this case $f = (2/3)m = 2n - 4$. ■

Lemma 13.1.6 (Planar graphs are degenerate). *Every planar graph G is 5-degenerate – that is, it has a vertex of degree at most 5. Furthermore, this is a hereditary property that holds for any subgraph of G .*

Furthermore, if G is a triangulation and it has more than three vertices, then there is a vertex of degree at most 5 in G that is not on the outer face of G .

Proof: Consider a planar graph G with n vertices and m faces. Clearly, by Lemma 13.1.5, $\sum_{v \in V(G)} \deg(v) = 2m \leq 6n - 12$, which implies that there is a vertex v of degree ≤ 5 in G .

The second part requires slightly more work. The number of vertices $n > 3$, and as such, each one of the three vertices r, g, b of the outer face is of degree at least three. As such, for $V' = V(G) \setminus \{r, g, b\}$, we have $\alpha = \sum_{v \in V'} \deg(v) = 2m - \deg(r) - \deg(g) - \deg(b) \leq 2(3n - 6) - 9 = 6n - 21$. As such, the average degree of vertices in V' is $\alpha/(n - 3) \leq (6n - 21)/(n - 3) < 6$, which implies that there must be at least one vertex in V' of degree at most 5. ■

Remark 13.1.7. A nice application of Lemma 13.1.6 is showing that a planar graph can be colored using 6 colors. Indeed, let v be the vertex of degree at most 5 in G . Color recursively $G - v$ using 6 colors, and extend it to a coloring of G by assigning v a color that is not used (out of the six available colors) among its (at most five) neighbors. This results in a valid coloring of G .

Showing that planar graphs are five colorable requires some additional work. The celebrated **four color theorem** states that planar graphs can be colored using four colors – the only known proofs requires computers to check hundreds of special cases.

Every face can be the outer face. Consider a planar graph G with an embedding \mathcal{G} of it in the plane. We can always turn any face of the embedding to be the outer face – one way to see that is via inversions. Indeed, consider a face f in \mathcal{G} , and draw a circle \circ , so that it is fully contained in f . Let r be the radius of \circ , and let p be its center. Consider the mapping $f(q) = \frac{r}{\|q-p\|}(q-p) + p$. This is an inversion that maps the outside of \circ to its interior, and vice versa. Applied to \mathcal{G} , it results in an “inverted” drawing, having f as the outer face. See Figure 13.2.





Figure 13.2: inversion of a planar graph drawing making an arbitrary face the outer face.

13.1.3. Representing an embedding of a planar graph

An (implicit) representation of a planar graph embedding, in addition to the regular information of vertices and edges, also lists the faces of the embedding. For each face, there is a list of its boundary cycles (with a special flag designating the outer cycle). For each edge uv , there is a pointer to its (at most two) adjacent faces.

A standard such representation is the *doubly connected edge list (DCEL)*. Every edge uv is associated with two directed edges $u \rightarrow v$ and $v \rightarrow u$ (called *half-edges*), that are *twins*. Here, one can think about a planar graph as a road map – with the convention of driving on the left. A vertex is a crossing, and an edge is a two lane road connecting two crossings (i.e., a lane is a half-edge). In particular, the outer boundary component of a face is a cycle of half-edges, oriented such that the cycle goes in counterclockwise direction as we traverse it. Similarly, an inner boundary cycle is oriented in a clockwise direction. An half-edge is as such adjacent to a single face, which lies to its “left”. As such, an half-edge belong naturally to one boundary cycle that is a part of. In particular, every half-edge stores pointers to (i) its twin, (ii) its adjacent face, (iii) next half-edge in the cycle it is on, and (iv) previous half-edge in the cycle.



For a vertex, there is a cyclic clockwise sorted list of half-edges that leave it. This list can be represented implicitly, as it is enough for a vertex to store a pointer to a single one half-edge that leaves it, and it is then easy, using the above pointers, to extract this cyclic list.




Figure 13.3: A planar graph, and a straight-line embedding of this graph.




Figure 13.4: Illustration of Lemma 13.1.8

13.1.4. A straight line drawing of a planar graph

Two embeddings of a planar graph are *homeomorphic*, or simply *equivalent*, if one can continuously deform one into the other. It is not hard to check that two embeddings of a planar graph are homeomorphic if (i) the same face is marked as the outer face, (ii) they have the same incidence structure between vertices, edges and faces, (iii) specifically, the order of edges around each vertex is the same in both embeddings, and (iv) the order of edges (and vertices) around each face is the same. That is, their DCEL description is the same. A *straight-line embedding* is a drawing of a planar graph where the edges are segments. And a straight-line embedding is in *general position* if no two segments are colinear. A straight-line embedding is depicted in Figure 13.3.

Lemma 13.1.8. *Let G be a given (simple) planar graph with an embedding \mathcal{G} . Then, there is an equivalent straight-line embedding of G in general position.*

Proof: We might as well assume G is a triangulation, by Lemma 13.1.4. The proof is by induction. For $n = |V(G)| = 3$ the claim is obvious as the graph is a triangle.

Otherwise, consider a vertex v of G that is not on the outer face of G . Consider removing v and all its adjacent edges in the given embedding of G (i.e., its an embedding of $G - v$). The removal of the vertex v created a hole – a face f with k edges. The face f is without holes, and furthermore, there is no edge that appears twice on its boundary, because this would imply that v has two parallel edges to some vertex on the boundary of f . See Figure 13.4.

Next, we pick an arbitrary vertex u on f , and connect it to all the other vertices of f not adjacent to it, and let e_1, \dots, e_t be these added diagonals (all drawn inside f). Let H be the resulting triangulated graph, together with the constructed embedding. The graph H has $n - 1$ vertices, and by induction, the current embedding can be realized by an equivalent straight line embedding (in general position). The face f is now a simple polygon.

In this embedding, create a copy of u and reassign e_1, \dots, e_t from u to the “new” vertex v (these diagonals are now segments). Next, move v slightly into the interior of f , so that the segments e_1, \dots, e_t have the same ordering around v as around u , and they do not intersect the boundary of f in their interior (here, implicitly, we are using the general position assumption to argue that such a movement is possible). In addition, connect v to u by a segment, and perturb v if needed to ensure the embedding is in general position. Clearly, the resulting straight-line embedding is equivalent to the given embedding of G . ■




Figure 13.5: Illustration of why $K_{3,3}$ is not planar.

13.1.5. Characterizing planarity by forbidden subdivisions

We remind the reader that K_n denotes the complete graph (i.e., *clique*) over n vertices. The graph $K_{n,m}$ denotes the *bipartite clique* with n vertices on one side, and m vertices on the other side, and edges connecting all possible pairs of edges that are on the two different sides.

Lemma 13.1.9. *The bipartite clique $K_{3,3}$ and the clique K_5 are not planar graphs.*

Proof: The graph K_5 is a graph with $n = 5$ vertices and $m = \binom{5}{2} = 10$ edges. By Lemma 13.1.5, $10 = m \leq 3n - 6 = 9$, which is impossible.

We provide two proofs that $K_{3,3}$ is not planar. To this end, let the two sets of the vertices of $K_{3,3}$ be $X = \{b, c, d\}$ and $Y = \{\beta, \gamma, \delta\}$. The set of edges of the graph are $E = \{xy \mid x \in X, y \in Y\}$. Consider a (fictional) planar embedding of $K_{3,3}$, and consider the cycle $C = \langle b, \beta, c, \gamma, d, \delta \rangle$ (in this order) – by planarity this is indeed a cycle. The edge $c\delta$ connects two antipodal vertices of C , and assume that $c\delta$ is contained in the interior of C (if it goes through the exterior of C , one can apply inversion, to make it an interior edge). The edge $b\gamma \in E$ must be in the exterior of C , as otherwise it would cross $c\delta$, see Figure 13.5. But then, the cycle $J = \langle b, \delta, c, \gamma \rangle$ contains d in its interior, and β is outside C . By the Jordan curve theorem (T13.1.1) the embedding of the edge $d\beta$ to intersect J , which implies that the drawing is not planar.

The second proof relies on extending Lemma 13.1.5 to bipartite graphs. In a planar graph the boundary of a face is a cycle. In $K_{3,3}$ such a cycle alternates between vertices of X and vertices of Y , which implies that it must be of either length four or six. So consider a planar graph with n vertices and m edges, where all the faces have exactly four edges (as usual, if there faces with more edges on the boundary, we insert a new edge to split the face into smaller faces, which with at least four edges on their boundaries. Let f be the number of faces in this graph. We have that $2m = 4f$ and $f - m + n = 2$ by Euler's formula. As such, we have $m/2 - m + n = 2 \implies m = 2n - 4$. As such, we conclude that for a bipartite planar graph, we have that $m \leq 2n - 4$. Getting back to $K_{3,3}$, we have that $n = 6$ and $m = 9$, but $9 = m \leq 2n - 4 = 8$, which is impossible. ■

We state, without proof, the following beautiful characterization of planar graphs. A *subdivision* of a graph is formed by subdividing its edges into paths of one or more edges. A graph H *contains a subdivision* of G , if there is a subgraph of $Q \subseteq H$, that is a subdivision of H – formally, Q is *isomorphic* to some subdivision of G . Here, two graphs G_1 and G_2 are *isomorphic* if up to renaming of vertices, they are the same graph.

Theorem 13.1.10 (Kuratowski's theorem). *A graph G is planar if and only if it does not contain $K_{3,3}$ and K_5 as a subdivision.*

The proof of Kuratowski's theorem is similar in spirit to the argument used in the planarity testing algorithm presented in Section 13.2. A closely related and equivalent result is *Wagner's theorem*,




Figure 13.6: (A) The graph G and its subgraph H . (B) All the H -fragments that are edges. (C) A bigger fragment. (D) The other big fragment.

which states that a graph is planar if and only if it does not contain $K_{3,3}$ and K_5 as a minor. A graph H is a *minor* of G if there is a sequence of edge deletions, vertex deletions, and edge contractions that transform G to H .

13.2. Planarity testing

13.2.1. Fragments and conflicts

For a graph $J \subseteq G$, its *cut* in G is the set of edges $\text{cut}(J) = \{uv \in E(G) \mid u \in V(J) \text{ and } v \in V \setminus V(J)\}$. For a set of edges E' , and a graph H , we denote by $H \cup E'$ the graph formed by adding the edges of E' to the graph G . Formally, we have $H \cup E' = (V(H) \cup V(E'), E(H) \cup E')$.

Definition 13.2.1. For a graph $H \subseteq G$, an *H -fragment* X of G is either

- (A) An edge $e \in E(G_{V(H)}) \setminus E(H)$ (i.e., an edge of G missing in H with both endpoints in $V(H)$).
- (B) A subgraph X of G , formed by taking a connected component C of $G \setminus H$, together with its cut. Formally, $X = C \cup \text{cut}(C)$.

For an H -fragment X , its *interface* is the set of vertices $\partial X = V(X) \cap V(H)$.

See Figure 13.6 for examples of fragments.

Lemma 13.2.2. *Let C be a cycle in G , X a C -fragment of G , and let $x, y, z \in C$ be three distinct vertices that are also in ∂X . Then, there exists a vertex $u \in V(X)$, and paths $\pi_{ux}, \pi_{uy}, \pi_{uz}$ that connects u to these three interface vertices, respectively, and furthermore, these paths are interior disjoint.*

Proof: Let \mathcal{T} any spanning tree of X , and consider the shortest path σ_y between x and y in \mathcal{T} , and the shortest path σ_z between x and z in \mathcal{T} . These two paths share a prefix, and then they diverge, and never meet again. The vertex of divergence is u , and the desired paths are the respective portion from the interface vertices to u . ■

Lemma 13.2.3. *Let C be a cycle in a planar graph G . Given two C -fragments X and Y , such that there are four vertices v_1, v_2, v_3, v_4 in cyclic order along C , such that $v_1, v_3 \in \partial X$ and $v_2, v_4 \in \partial Y$. Then, X and Y are **conflicting** – in any planar drawing of G these two fragments are on different sides of C (i.e., one of the fragments would be inside the close Jordan curve formed by the cycle C , and the other fragment would be outside this cycle).*

Similarly, X and Y conflict if there are three boundary vertices v_1, v_2, v_3 on C , such that $v_1, v_2, v_3 \in \partial X, \partial Y$.




Figure 13.7




Figure 13.8

Proof: Assume, for the sake of contradiction, that this is false, and there is a drawing of $C \cup X \cup Y$ having both fragments inside the cycle C (the case that they are both outside is handled in a similar fashion).

But then, we can add an additional vertex u outside C , connect it to v_1, v_2, v_3, v_4 by edges outside C , and extract two paths $\pi \subseteq X$ and $\pi' \subseteq Y$, where π connects v_1 to v_3 , and π' connects v_2 to v_4 , and these two paths do not intersect, see Figure 13.7. This is a planar drawing of a subdivision of K_5 , which contradicts Kuratowski's Theorem (T13.1.10).

The other case follows by a similar argument, if one could draw the two fragments in the same side of C , with both having three common interface vertices v_1, v_2, v_3 , then one can draw $K_{3,3}$. Indeed, by extracting two center vertices in the respective fragment (using Lemma 13.2.2), connecting these center vertices to v_1, v_2, v_3 , and now adding an external vertex outside C and connecting it to v_1, v_2, v_3 by edges outside C , we get the desired drawing. See Figure 13.8. A contradiction.

Observation 13.2.4. *In the settings of Lemma 13.2.3, the C -fragments X and Y can not be in the same face of a planar drawing, if ∂X and ∂Y share three vertices.*

13.2.2. Algorithm

13.2.2.1. Bridges and 2-connected graphs

Definition 13.2.5. A **bridge** is an edge in a graph whose removal disconnects the graph. Similarly, a **cut vertex** is a vertex whose removal increases the number of connected components in the graph.

A graph G is **k -connected** if the smallest set of vertices whose removal disconnects G is of size at least k .

If a planar graph G has a bridge e , we can embed the two connected components of $G \setminus e$ independently, and then combine them to get an embedding for G . Similarly, if the planar graph has a vertex v whose removal disconnects it, then embed every v -fragment of G separately, make sure that v is on the outer face in each of the embeddings, and then glue all the copies of v together to get the desired embedding.

13.2.2.2. Description of the algorithm

Initial checks. The algorithm scans the graph and removes parallel edges and self loops if they exist. Next, the algorithm counts the number of edges m in the graph, if $m > 3n - 6$, then by Lemma 13.1.5, the graph is not planar, and the algorithm stops.

Next, it computes the bridges and cut vertices in the given graph using **DFS** in linear time. It removes the bridges, computes the embedding for each component separately as described below, and

then glue them back together for an embedding of the whole graph, and this takes linear time. It handles the cut vertices in a similar fashion.

Embedding a component. Since the algorithm broke the graph at cut vertices, one can assume the given graph is 2-connected. The above suggests a natural algorithm for a planarity testing – start with a cycle G_0 in the given graph G . In the i th iteration, find a path that connects two vertices of G_{i-1} , and its internal vertices are fully contained in a G_{i-1} -fragment. (Observe, that any fragment has at least two interface vertices, since there are no bridges.) Add this path to the graph G_{i-1} to form G_i , and repeat till all of G is laid out, or until the algorithm get stuck.

The problem, of course, is how to choose which face in the current embedding should contain the new added path. Potentially, a path (or a fragment) can be placed in many faces, see Figure 13.9. To this end, let \mathcal{G}_{i-1} denote the computed embedding of G_{i-1} in the start of the i th iteration.

For a G_{i-1} -fragment H , let $F(i, H)$ be the set of faces that contain (on their boundary) all the interface vertices of H (i.e., the vertices of ∂H); that is,

$$F(i, H) = \left\{ f \in \text{faces}(\mathcal{G}_{i-1}) \mid \partial H \subseteq V(f) \right\}.$$

If there is a fragment H such that $n(i, H) = |F(i, H)|$ is zero (i.e., no face contains all the interface boundaries of H), then the algorithm had failed^③ and the graph G is not planar.

Similarly, if $n(i, H) = 1$ for some fragment H , then we have a single face f of \mathcal{G}_{i-1} that may contain this fragment, we compute a path π_i (in H) between two interface vertices of H , add π to G_{i-1} to form G_i , and add π to \mathcal{G}_{i-1} to form \mathcal{G}_i by splitting f into two new faces.

The interesting case is when $n(i, H) > 1$ for all the G_{i-1} -fragments. Surprisingly, in this case, Yogi Berra was right^④ – pick an arbitrary fragment, and an arbitrary face that might contain it, and perform the same path embedding described above. We repeat this till G is fully embedded, or till failure.

13.2.2.3. Correctness

Lemma 13.2.6. *The above algorithm computes a planar embedding for a graph. If it fails, then the graph is not planar.*

Proof: The proof is by induction, arguing that for any partial embedding \mathcal{G}_j computed, there is an extension of it so that it embeds the whole graph (if the given graph is planar). The claim is obvious for \mathcal{G}_0 , as any cycle in a planar graph has to be drawn as a cycle in any embedding. So, assume this is true for \mathcal{G}_{i-1} .

For a G_{i-1} -fragment H , if $n(i, H) = |F(i, H)| = 1$, then there is only one face of \mathcal{G}_{i-1} that might contain H in the extension, and the algorithm adds a path in H to \mathcal{G}_{i-1} , to get \mathcal{G}_i preserving feasibility of the planar embedding of the whole graph.

So, consider the case that, for all G_{i-1} -fragment H , we have $n(i, H) > 1$. The algorithm had embedded a path π that belongs to some G_{i-1} -fragment H in a face f of \mathcal{G}_{i-1} , and assume that this was a mistake, and the algorithm should have embedded H (and thus π) in a face f' of \mathcal{G}_{i-1} , and let \mathcal{G} be an embedding




Figure 13.9

^③Not us! We can never fail.

^④When you come to a fork in the road, take it.




Figure 13.10: illustration of the proof of Lemma 13.2.6.

of the whole graph under this choice. The idea is to modify \mathcal{G} into an embedding of G , where π is inside f .

To this end, let $B = V(f) \cap V(f')$ be the set of vertices that appear in both faces, and observe that $\partial H \subseteq B$. We take all the fragments whose interface vertices are contained in B , and we flip them in the embedding \mathcal{G} between f and f' , let \mathcal{G}' be the resulting embedding. See Figure 13.10.

If this slight of hand succeeded, then we are done, as we found a feasible embedding of the graph that is in agreement with the choices the algorithm made so far. However, potentially, this failed because some G_{i-1} -fragment Q that is embedded in f' , say, in \mathcal{G} , had conflicted with another fragment J that is already in f and did not change its face. See Figure 13.10.

It must be that $\partial J \not\subseteq B$, as otherwise J would have happily flipped and would not have collided with Q . Let u be an arbitrary vertex in $\partial J \setminus B$. Trace the cycle boundary of f clockwise (resp. counterclockwise) till encountering a vertex x (resp. y) of ∂Q . These two vertices must exist (since the interface of a fragment always have at least two vertices). Let Z be the set of vertices of the boundary of f between x and y (including x and y). If Z contain all the vertices of ∂J then there is no collision between Q and J . As such, there must be an additional vertex, say $v \in \partial J \setminus Z$, see Figure 13.11.





Figure 13.11

The key observation is that there is no other face g of \mathcal{G}_{i-1} (in addition to f), such that $\partial J \subseteq V(g)$. Assume, for the sake of contradiction that there is such a face g . First, observe $g \neq f'$, since $u \notin B \subseteq V(g)$ and $u \in \partial J$.

We claim that this implies a planar drawing of K_5 . Indeed, starting with the embedding \mathcal{G}_{i-1} , connect u to v by a path σ that is fully contained in the mystical face g . Similarly connect x and y through f' . The edges xu, uy, yv and vx can be drawn inside f by tracing them along the boundary of f . Finally, add a surprise vertex z in the interior of f , and connect it to the vertices x, y, u, v inside f in the natural way. This results in the desired drawing of K_5 , which is impossible.



As such, $n(i, J) = 1$. But then, the algorithm would have used this fragment at this iteration, and not H . A contradiction to the choice of H . We conclude that the flipping succeeded, and the resulting embedding \mathcal{G}' is a valid embedding of G that agrees with the choice of the algorithm in the i th iteration, implying the algorithm can not get stuck, unless the graph is not planar. ■

13.2.2.4. Efficient implementation

Clearly the above planarity algorithm has polynomial running time. Getting quadratic running time requires some care.

During the execution the algorithm, the faces of the embedding \mathcal{G}_i are all simple – every face has a single boundary component, which is a cycle. Every vertex keeps a list of the fragments attached to it, and every fragment keeps a list of its interface vertices. Furthermore, every fragment X keeps a count $\alpha(X)$ of the number of faces that it might be drawn in (i.e., it is a variable holding the value of $n(i, X)$).

Given a face f , one can mark all the vertices of $V(f)$ (as a preprocessing step). Then, given a fragment X one can decide, in $O(|\partial X|)$ time, if $\partial X \subseteq V(f)$. Furthermore, the total number of fragments, in any point in time, is at most $O(n)$, since every fragment contains at least one edge that belongs only to this fragment, where $n = |V(G)|$ (here, we use $|E(G)| = O(n)$).

In the i th iteration, when the algorithm adds the path π_i , it splits a face f into two faces f_1, f_2 . This splits the fragment X_i that contains π_i into two or more new fragments. For such a new fragment $Y \subseteq X_i$ we temporarily set $\alpha(Y) = \alpha(X_i)$.

Next, scan the current set of fragments (including the new fragments). For each fragment X , check whether it could be embedded in f , and let $\beta(X, f) = 1$ if so, and zero otherwise. If $\beta(X, f) = 1$ then check if X can be embedded in f_1 and f_2 . Computing this information for all fragments takes $O(\sum_X |\partial X|) = O(n)$ time, since every interface vertex can be uniquely charged to an edge that is not yet embedded, and there are $O(n)$ such edges. Now, set $\alpha(X) \leftarrow \alpha(X) - \beta(X, f) + \beta(X, f_1) + \beta(X, f_2)$. This updates $\alpha(\cdot)$ for all fragments. If during the execution any of these fragment counts become zero, the algorithm stops, and outputs that the graph is not planar.

Now, in the beginning of each iteration, the algorithm scans the list of fragments looking for a fragment with face count of one. If such a fragment is found it is embedded, otherwise the algorithm picks an arbitrary fragment to embed.

Clearly, the total amount of work done in each iteration is $O(n)$, as desired.

Lemma 13.2.7. *Given a graph G with n vertices, the above algorithm checks, in $O(n^2)$ time, if it is a planar graph, and if so it outputs a planar embedding of G .*

It is clear that this algorithm is not that efficient, and one should be able to do planarity embedding faster. And indeed, linear time algorithms for this problem are known. Intuitively, one can do the embedding in a more systematic fashion, keeping the invariant that the embedded part is a tight cluster in the graph, such that the paths added as the algorithm progresses are on the outside of the parts that were already embedded. Nailing the details down and getting a linear time algorithm proved to be surprisingly challenging, and the technical details are quite subtle. We state one such (relatively recent) result [bm-oces-04] without proof – the algorithm is relatively simple.

Theorem 13.2.8. *Given a graph G with n vertices, an algorithm can check if it is a planar graph in $O(n)$ time, and if so it outputs a planar embedding of G .*

13.3. Bibliographical notes

A good book on graphs containing the basic results on planar graphs is the classical text by Bondy and Murty [bm-gta-76], which was at some point available online for free (legally) on the web. There is also a more updated version of this book [bm-gt-11]. Another good textbook is by West [w-igt-01].

Jordan curve theorem. The Jordan curve theorem has a long history – it was first observed that it is not obvious, and a Bolzano was the first to state that it requires a formal proof. Jordan provided a proof in his book (1887), but it was somewhat sketchy – and it was claimed to be incorrect by Veblen who provided a more elaborate formal proof (however, the original proof by Jordan is believed to be

correct). Many proofs of the theorems were provided later on. A nice short proof is provided by Maehara [m-jctbf-84] – the proof is not completely elementary, and uses Brouwer’s fixed point theorem, and that identity mapping on a curve can be extended to a disk. A nice discussion of the Jordan curve theorem is provided by Wikipedia.

Other criteria for planarity. The Hanani-Tutte theorem states that a graph is planar if and only if there is a drawing of a graph in the plane such that every pair of edges intersect an even number of times. This leads to an algebraic approach to testing planarity, that is less efficient than the approach shown here. A good survey is provided by Schaefer [s-htrr-13].

A rant against planar graphs. While planar graphs are quite common and have many beautiful properties, they are fragile – add a single edge to a planar graph and it may no longer be planar. There are good reasons for this – constant degree expanders are far from being planar graphs. In particular, a union of three random perfect matchings over n vertices (i.e., perfect matchings in the complete graph) form a graph that is a constant degree expanders with good probability. It is not hard to show that such an expander requires $\Omega(n^2)$ edge crossings when drawn in the plane. On the other hand, the union of the first two matchings form a collection of cycles, which is definitely a planar graph. That is – injecting n edges into this graph, completely ruins its planarity. Furthermore, many problems are computationally harder on expanders than on planar graphs.

It is thus natural to look to other criteria (than planarity) if one wants a robust family of graphs (that can withstand a moderate number of insertions/deletions/edit operations), which is computationally tractable. One such family is *low-density* graphs, which are related to representation of graphs as intersection graph of geometric objects (it includes planar graphs by the circle packing theorem), see [hq-aaldg-15-arxiv].

Other relevant chapters:

- (A) The *crossing lemma* – how many crossings must a drawing of a non-planar graph have in described in the book [h-gaa-11].
- (B) The grid embedding chapter, available [here](#)^⑤, present results on how to draw a planar graph as a straight line embedding on a small grid.
- (C) The circle packing theorem chapter, which shows that planar graphs can be represented as the intersection graph of disks, is available [here](#)^⑥.
- (D) The chapter on the planar separator theorem and its variants is available from [here](#)^⑦.

13.4. Exercises

13.5. From previous lectures

^⑤http://sarielhp.org/book/chapters/planar_embed_grid.pdf.

^⑥http://sarielhp.org/book/chapters/planar_circle_packing.pdf.

^⑦http://sarielhp.org/book/chapters/planar_separator.pdf.