# Algorithms and Combinatorics of Embedded graphs Maths-CS track M2, UGE 

Arnaud de Mesmay

These are the lecture notes for the first half of the course on geometry, focusing on Algorithms and Combinatorics of Embedded Graphs.

Some practicalities:

- The course is on Mondays, 13:30 to 16:30 in room 2B092 of the Copernic building.

This first half of the course focuses on embedded graphs: we start with planar graphs and then move on to surface-embedded graphs. These are graphs that can be drawn without crossings in the plane or in more complicated surfaces, see Figure 1. Therefore they form a combinatorial object with a topological constraint. The objective of the course is to explore how topology interacts with combinatorics and algorithms on this very natural class of objects. Some questions we will explore are:

- What are the combinatorial consequences of being planar? Are there combinatorial characterizations?
- How to test algorithmically whether a graph is planar? If yes, how to draw the graph?
- Can one exploit planarity to design better algorithms for planar graphs than for general graphs?
- How do the previous answers generalize to graphs embedded in these more complicated surfaces?
- How to solve certain topological questions algorithmically on surfaces?

While the focus of the course will stay very theoretical (e.g., mostly with a theorem, lemma, proof structure), embedded graphs are of great interest for the practically-oriented mind, as they appear everywhere, for example in road networks (where underpasses and bridges can be modeled using additional topological features), chip design or the meshes that are ubiquitous in computer graphics or computer aided design. In all these applications, there is a strong need for a theoretical understanding of embedded graphs, as well as algorithmic primitives related to their topological features. Additionally, embedded graphs are an important lens to study graphs in general, since any graph can be embedded on some surfaces. This is especially the case in graph minor theory, where embedded graphs play an absolutely central role (but we will barely touch on this topic).

These lecture notes follow quite closely the material taught in class. This is actually their point, as in my experience lecture notes with too much content can easily get overwhelming (especially when one misses a class). Thus we will refrain from (too many) digressions and heavy references. The tone will be conversational. These lecture notes are mostly taken from a course I taught on the same topic in 2019-2020 at MPRI, see here, which are themselves heavily


Figure 1: Graphs embedded in the plane, the torus, and the three-holed torus.
inspired by the lecture notes of a previous iteration of the class(by Éric Colin de Verdière), and those for a course on Computational Topology I co-taught with Francis Lazarus in 2016-2018, and we refer to, and strongly recommend those for the missing digressions and references.

## 1 Planar Graphs

A graph is planar if it can be drawn on the plane so that no two edges intersect, except at their common endpoints. This very simple property has a very strong impact on the combinatorics of the graph: for example they are sparse (their number of edges is linear in the number of vertices), they can be colored with few colors; as well as on the algorithms that one can design for these graphs.

### 1.1 Point-set topology, graphs and embeddings

We will use basic notions of point set topology. For the unfamiliar reader, this is like metric topology, but without a metric (!). This is very relevant for because in most cases, we will not care about the Euclidean metric on the plane (or any other), and a vague notion of neighborhood will be enough. A topological space is a set $X$ with a collection of subsets of $X$, called open sets, such that:

- the empty set and $X$ are open.
- any union of open sets is open.
- Any finite intersection of open sets is open.

Any metric space gives rise to a topological space by taking as open sets the smallest family of sets containing the open (in the metric sense) balls and satisfying the above axioms. A map $f$ between two topological spaces is continuous if the inverse image of any open set by $f$ is open. A map $f$ is a homeomorphism if it is continuous, bijective, and if its inverse $f^{-1}$ is also continuous. Homeomorphic spaces throughout these notes will generally be considered to be the same, and thus identified without precaution. Therefore, when we talk about the sphere, we talk about any space homeomorphic to the 2 -sphere $\mathbb{S}^{2}=\left\{x, y, z \mid x^{2}+y^{2}+z^{2}=1\right\}$ with its standard 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 called simple. In a simple graph every edge is identified unambiguously with the pair of its endpoints. 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 $\dot{X}$ of $X$ is the union of the open balls contained in $X$. The boundary $\partial X$ is the closure minus the interior. An embedding of a non-loop edge in the plane is a topological embedding (i.e., an injective map) of the segment $[0,1]$ into $\mathbb{R}^{2}$. Likewise, an embedding of a loop-edge is an embedding of the circle $S^{1}=\mathbb{R} / \mathbb{Z}$. Such an embedding of the circle is called a simple closed curve in the plane. An embedding of a finite graph $G=(V, E)$ in the plane is defined by a bijective map $V \rightarrow \mathbb{R}^{2}$ 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

a.

b.

c.

Figure 2: a. A planar graph (!). b. The same planar graph, this time embedded in the plane, i.e., a plane graph. c. A non-planar graph
vertices $\mathbb{1}^{1}$. Note that this definition can be (and will be) generalized to other topological spaces than $\mathbb{R}^{2}$.

A graph is planar if it has an embedding into the plane, see Figure 2. The data of the planar graph and a specific embedding is called a plane graph, although it is customary to mix up everything and also call it a planar graph. We consider two plane graphs to be the same if they are homeomorphic, i.e., if there exists a homeomorphism of the plane sending one to the other.
Exercise 1.1. 1. Any graph can be embedded in $\mathbb{R}^{3}$.
2. A graph can be embedded in the plane if and only if it can be embedded in the sphere. (Hint: prove that the plane is homeomorphic to the sphere with one arbitrary point removed).

The second exercise allows us to jump at will between embeddings in the plane and on the sphere, which we will constantly do.

### 1.2 The Jordan curve theorem

The main additional structure that a planar graph has, compared to an abstract one, is that it has faces. A face of a graph embedding is a connected component of the complement of the image (via the embedding) of the vertices and edges of the graph. Then the basic properties of planar graphs follow from the fact that a simple closed curve in the plane is the boundary of exactly two components, one of which is a disk:

Theorem 1.2 (Jordan curve theorem). Let $C$ be a simple closed curve in the plane. Then its complement has two connected components, one of which is bounded, and each of which has C as a boundary.

This obvious fact is famously painful to prove. One reason is that simple closed curves can be very pathological, like fractals, see for example Figure 3. In the simpler case that $C$ is polygonal, i.e., a concatenation of segments, the proof is much simpler.

Proof in the polygonal case. 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 $C$. Rotate the plane so that the direction that we have

[^0]

Figure 3: Three simple closed curves in the plane
chosen is actually horizontal. 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 $\pi: \mathbb{R}^{2} \backslash C \rightarrow\{$ even, odd $\}$ that counts the parity of the number of lower half-open segments of $C$ intersected by a ray:

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

Every $p \in \mathbb{R}^{2} \backslash C$ is the center of small disk $D_{p}$ over which $\pi$ is constant. Indeed, let $S_{p}$ be the set of segments (of $C$ ) that avoids $h_{p}$, let $S_{p}^{\prime}$ be the set of segments whose interior crosses $h_{p}$ and let $S_{p}^{\prime \prime}$ 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 $S_{q}=S_{p}, S_{q}^{\prime}=S_{p}^{\prime}$ and the parity of $\left|S_{q}^{\prime \prime}\right|$ and $\left|S_{p}^{\prime \prime}\right|$ is the same, see Figure 4 It follows that $\pi(q)=\pi(p)$.


Figure 4: Polygonal Jordan curve theorem
Since $\pi$ is locally constant, it must be constant over each connected component of $\mathbb{R}^{2} \backslash C$. 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 $\mathbb{R}^{2} \backslash C$ has at least two components. To see that $\mathbb{R}^{2} \backslash C$ has at most two components consider a small disk $D$ centered at a point interior to a segment $s$ of $C$. Then $D \backslash C=D \backslash s$ has two components. Moreover, any point in $\mathbb{R}^{2} \backslash C$ can be joined to one of these components by a polygonal path that avoids $C$ : first come 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 $\mathbb{R}^{2} \backslash C$.

A stronger version of this theorem, which we will not prove either, shows that one of the two components is homeomorphic to an open disk.

Theorem 1.3 (Jordan-Schoenflies Theorem). Let $C$ be a simple closed curve in the plane. Then its complement has two connected components, one of which is a topological disk, and each of which has $C$ as a boundary.

One (not that immediate) consequence of this theorem is that in an embedding of a connected graph, all of the faces except one are open disks (or equivalently, on a sphere, all of the faces
are open disks). Such a plane graph is called a celullar embedding. This definition is not very interesting for now but will get more important when embedding graphs on other surfaces. Each face can be described by a facial walk, which is the circuit on the graph describing the boundary of the face. Note that a facial walk may have repeated vertices and edges (see Figure 5, left), but this does not happen when the graph is connected enough. A graph $G$ is 2 -connected is for any vertex $x, G \backslash x$ is connected.

Lemma 1.4. In a 2-connected planar graph, facial walks are cycles of the graphs (i.e., without repeated vertices and edges).

Proof. Assume otherwise, let $F$ be an open disk bounded by a facial walk that uses the same vertex $x$ twice. Then $F \cup\{x\}$ contains a Jordan curve separating the graph, contradicting the 2 -connectedness assumption.

The dual graph of a cellular graph embedding $G=(V, E)$ on $\mathbb{S}^{2}$ is a plane graph $G^{*}$ defined as follows: put one vertex $f^{*}$ of $G^{*}$ in the interior of each face of $G$, and for each edge $e$ of $G$, create a dual edge in $G^{*}$ crossing $e$ and no other edge of $G$. The dual graph is also cellular, and is, up to homeomorphism, uniquely defined by the above. Note that the dual graph crucially depends on the embedding, and that duality does not preserve simplicity. As the name suggests, duality is an involution: $G^{* *}=G$. See Figure 5, right).


Figure 5: A cellularly embedded graph and its dual
This allows us to think of a plane graph as a collection of open disks glued together with a combinatorial data, which is convenient both for combinatorial and algorithmic purposes.

### 1.3 Actual planar graphs

When dealing with cellularly embedded graphs in an algorithmic fashion, we have to worry about how the embedding graph is specified to the algorithm. One approach could be to embed each edge as a polygonal curve (which is always doable (somewhat easy exercise)), and specify the coordinates of all the vertices and breaking points of the polygonal edges. This suffers from many issues: it is a very wasteful encoding since at least naively, there is no easy bound on the number of polygonal bends needed on each edge (we will come back to this question later). And it is not at all obvious when two encodings describe the same graph, i.e., two homeomorphic graphs.

Instead, one customary way to describe an embedding of a graph on a sphere is to use a combinatorial map. A combinatorial map associated to a celullar graph embedding $G$ is the set of closed walks in $G$ obtained from walking around the boundary of each face of $G$. Note that these walks may have repeated vertices or edges. By the previous subsection, each face is a disk. Thus the combinatorial map contains enough information to build an embedding of $G$ on a sphere abstractly, by taking the abstract graph and gluing a disk on each face. For graphs embedded in the plane, it suffices to mark one face as the unbounded face.

Another approach is to use a rotation system: the idea is to use as an encoding the cyclic permutation describing the order of the edges around each vertex in the embedding. See Figure 6 for examples of combinatorial maps and rotations systems.


Figure 6: A cellularly embedded graph, its facial walks, and a rotation system describing it.
Exercise 1.5. Think about naively how to compute a rotation system from a combinatorial map and vice versa. Is the complexity any good?

Both approaches are a bit unwieldy for actual implementations, as basic operations like moving around a vertex or a face are overly costly. One more refined data structure is to use flags. A flag is an incidence between a vertex, an edge and a face of the embedding. Figure 7 shows a set of flags for a planar graph. Flags come with three involutions allowing us to move to a nearby flag:


Figure 7: Flags, and the three involutions moving one of them.

1. $v i$ moves to the flag with the same edge-face incidence, but with a different vertex incidence;
2. ei moves to the flag with the same vertex-face incidence, but with a different edge incidence;
3. $f i$ moves to the flag with the same vertex edge incidence, but with a different face incidence;

Note that a flag is not uniquely defined by a triple ( $v, e, f$ ) where $v$ is incident to $e$ and $e$ is incident to $f$, as one can see on a graph with a single vertex and a single loop edge.

The complexity of a plane graph described by flags is the total number of flags involved, which is linear in the number of edges, and thus, as we will see very shortly, in the number
of vertices and faces as well. Thus the complexity of a graph and of its embedding is linearly related.
Exercise 1.6. Think about naively how to compute a rotation system and a combinatorial map from a flag representation. Is it better?

Note that in this data structure, vertices, edges and faces are actually not stored. This can be unwieldy for some problems, and in this case, one can extend the data structure so that flags contain pointers to the underlying vertices, edges and faces. With this caveat, flags are versatile enough to handle any intuitive operation on a planar graph with the intuitive complexity. For example, navigating around a vertex $v$ has complexity $O($ degree $(v))$. Similarly, the dual of a graph is completely straightforward and just amounts to replacing $f i$ with $v i$ and vice versa. Thus, throughout the rest of these notes, and as is common in the literature, the precise translation to switch between an informal description of a navigation in the graph, and the precise procedures with the flags and their involutions, will be swept under the rug.

### 1.4 Euler's formula

Euler's formula shows that the number of vertices, edges and faces of a graph are related by a linear equation. It has a rich history, of which we will only say that it was likely not first discovered by Euler.
Lemma 1.7. Let $G$ be a 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 $p, p^{\prime}$ be two points in $f \backslash e$. There is a path in $f$ connecting $p$ and $p^{\prime}$. 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, $N_{e} \backslash e$ is connected). It follows that $p$ and $p^{\prime}$ belong to a same component of $f \backslash e$. We conclude that $G$ has only one face in $f$, so that $G$ and $G-v$ have the same number of faces.

Theorem 1.8 (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 1.7, $G$ 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 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 $f$ intersect 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 $f \backslash e$ has at most two components. It follows that $f$ contains exactly two faces of $G$, which concludes the claim.

Exercise 1.9. 1. A graph embedded on the plane is a triangulation if every face (including the outerface) is adjacent to exactly three edges. Show that if $G$ is a triangulation, its numbers of vertices and faces satisfy $2|V|=4+|F|$.
2. Let $G$ be a simple planar graph, and suppose that we arbitrarily color each edge of $G$ either blue or red. Prove that for any planar embedding of $G$, there exists a vertex around which the incident red edges are consecutive. Hint: Add edges to $G$ to obtain a triangulation $G^{\prime}$, then define a well-chosen bipartite graph out of the vertices and faces of $G^{\prime}$, look at the degree of its vertices and apply the previous question.

Solution of the exercise given very confusingly (sorry!) in class:

1. In a triangulation, every face is adjacent to exactly three edges, so by double-counting edges, we obtain that $3|F|=2|E|$. We plug that into the Euler formula $|V|-|E|+|F|=2$ to obtain that $|V|-|F| / 2=2$ and thus $2|V|=4+|F|$.
2. First, we add edges to $G$ until we obtain a triangulation $G^{\prime}$ and color its edges arbitrarily. If we prove the property for $G^{\prime}$, it also proves it for $G$ : indeed, if a vertex in $G^{\prime}$ has all its incident red edges consecutive, this is also the case for the same vertex in $G$, since it just has less edges. Now, let us define a graph $H$ as follows. It has two types of vertices, either vertices of $G^{\prime}$ (vertices of the first type) or faces of $G^{\prime}$ (vertices of the second type). There is an edge between a first-type vertex and a second-type vertex if the corresponding vertex and face are incident in $G^{\prime}$, and if the two edges of $G^{\prime}$ that are adjacent to that incidence have different colors. Then we remove the vertices of the second type with no adjacent edges. This graph $H$ has the following properties: (i) All the vertices of the second type have degree two and (ii) If $G$ has no vertex with all incident red edges consecutive, then the vertices of the first type all have degree at least four (since they see at least at least alternation red-blue-red-blue). But the number of vertices of the first type in $H$ is $\left|V\left(G^{\prime}\right)\right|$ and the number of vertices of the second type in $H$ is at most $\left|F\left(G^{\prime}\right)\right|$ (recall that we might have thrown some away). Therefore, by double counting the edges in this bipartite graph, we have that $4\left|F\left(G^{\prime}\right)\right| \geq 2\left|V\left(G^{\prime}\right)\right|$. But by the previous exercise, $2\left|V\left(G^{\prime}\right)\right|=4\left|F\left(G^{\prime}\right)\right|+4$. This is a contradiction, which concludes the exercise.

To illustrate the power of this formula, here is a standard application.


Figure 8: The graphs $K_{5}$ and $K_{3,3}$.
Theorem 1.10. The graphs $K_{5}$ and $K_{3,3}$ (see Figure 8) are not planar.
Proof. In any planar embedding of a graph, each edge is adjacent to two faces (perhaps twice the same). Since $K_{5}$ is simple, each face is adjacent to at least 3 edges. In any planar embedding of $K_{5}$, the Euler formula would give $|V|-|E|+|F|=2$, and by the previous two observations, we have $3|F| \leq 2|E|$, therefore $|V|-|E| / 3 \geq 2$. But plugging in the values gives $5 / 3 \geq 2$, a contradiction. For $K_{3,3}$, we observe that any facial walk in a planar embedding would have at least 4 edges because the graph is bipartite. This leads to $|V|-|E| / 2 \geq 2$, e.g., $3 / 2 \geq 2$, another contradiction.

The following two applications, very close to the proof we have just seen, are absolutely key.
Exercise 1.11. Every simple planar graph $G$ with $n \geq 3$ vertices has at most $3 n-6$ edges.
Exercise 1.12. Every simple planar graph has a vertex with degree less than 6.
The first one shows that simple planar graphs are sparse, i.e., verify $|E|=O(|V|)$. Therefore, when talking about the complexity of a simple connected planar graph, we can refer either to its vertices or edges.

The second exercise has the following nice application. A $k$-coloring of a graph is a map $c: V \rightarrow\{1, \ldots, k\}$ so that two adjacent vertices are mapped to different colors.

Proposition 1.13. Planar graphs are 6-colorable.
Proof. We prove the result by induction on the number of vertices. For low values, this is immediate. For the induction step, pick a vertex $z$ of degree less than 6 , and color inductively $G \backslash z$. Then the five neighbors have at most 5 different colors, and we can color $z$ with one of the remaining colors.

Exercise 1.14. Planar graphs are 5-colorable. Hint: look at paths connecting non-adjacent neighbors of a degree 5 vertex.

As is well known, planar graphs are actually 4 -colorable, and we will not prove this in the course.

To conclude, we prove a very strong variant of Exercise 1.11, which allows to quantify how non-planar dense graphs are. A drawing of a graph if 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 $\operatorname{cr}(G)$ of a graph is the minimal number of crossings over all the possible drawings of $G$. For instance, $\operatorname{cr}(G)=0$ if and only if $G$ is planar. The crossing number inequality provides the following lower bound on the crossing number.

Theorem 1.15. $\operatorname{cr}(G) \geq \frac{|E|^{3}}{64|V|^{2}}$ if $|E| \geq 4|V|$.
As an immediate corollary, we have that complete graphs $K_{n}$ have crossing number $\Omega\left(n^{4}\right)$. Figuring out the exact constant hidden in this $\Omega(\cdot)$ is a well-known open problem.

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 $G^{\prime}$ obtained by removing one edge for each crossing. This graph is planar since we removed all the crossings, and it has $|E|-c r(G)$ edges, so we obtain that $|E|-c r(G) \leq 3|V|$ (Note that we removed the -6 to obtain an inequality valid for any number of vertices). This gives in turn

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

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 $G^{\prime \prime}$ the resulting graph. Taking the previous inequality with expectations, we obtain $\mathbb{E}\left(\operatorname{cr}\left(G^{\prime \prime}\right)\right) \geq \mathbb{E}\left(\left|E^{\prime \prime}\right|\right)-3 \mathbb{E}\left(\left|V^{\prime \prime}\right|\right)$. Since vertices are removed with probability $1-p$, we have $\mathbb{E}\left(\left|V^{\prime \prime}\right|\right)=p|V|$. An edge survives if and only if both its endpoints survives, and a crossing survives if an 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 $\mathbb{E}\left(\left|E^{\prime \prime}\right|\right)=p^{2}|E|$ and $\mathbb{E}\left(\operatorname{cr}\left(G^{\prime \prime}\right)\right)=p^{4} \operatorname{cr}(G)$. So we obtain

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

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

### 1.5 Kuratowski's Theorem

We say that $H$ is a 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 nonplanar. It follows from Theorem 1.10 that a planar graph cannot have a subdivision of $K_{5}$ or $K_{3,3}$ as a subgraph. It turns out that the converse is also true.
Theorem 1.16 (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.

This is a striking instance of the interplay between the topology and the combinatorics: our topological constraint can be replaced by a purely combinatorial one, and vice versa. In the following $K_{5}$ and $K_{3,3}$ are called forbidden graphs.

Contracting an edge $e$ in a graph $G$ means removing it and identifying its two endpoints. When the graph is embedded, the contraction can be done while preserving the embedding. We denote this operation by $G / e$. A 3 -connected graph is a graph that stays connected when one removes an arbitrary pair of vertices.
Exercise 1.17. Let $G$ be a graph such that $G / e$ contains a subdivision of $K_{5}$ or $K_{3,3}$. Then $G$ also contains a such a subdivision. (Hint: this might not be the same forbidden graph!).

The proof of Kuratowski's theorem is very nice and clean in the 3 -connected case, and requires somewhat tedious machinery to reduce to that case. This machinery is summarized in the following two lemmas, whose proof is skipped in class.
Lemma 1.18. Any 3-connected graph $G$ with at least five vertices contains an edge e such that $G / e$ is 3-connected.
Lemma 1.19. Let $G$ be a graph with at least five vertices, containing no subdivision of $K_{5}$ nor of $K_{3,3}$ and such that the addition of any edge between non-adjacent vertices creates such a subdivision. Then $G$ is 3 -connected.

Proposition 1.20 (Kuratowski's theorem for 3-connected graphs). A 3-connected graph $G$ without any subgraph isomorphic to a subdivision of a forbidden graph is a planar graph.

Proof. We prove the contrapositive: any 3 -connected non-planar graph contains a subdivision of $K_{5}$ or $K_{3,3}$ as a subgraph. In order to do so, let us pick $G$ a contraction-minimal non-planar graph, i.e., a graph such that for any edge, contracting the edge $e$ yields a planar graph. We prove that $G$ contains a subdivision of $K_{5}$ or $K_{3,3}$ as a subgraph, which proves the theorem, since any non-planar graph contains a contraction-minimal non-planar graph as a subgraph, and the Exercise 1.17 guarantees that this property will still be satisfied after uncontracting.

Let us pick an edge $e=x y$ of $G$ so that $G / e$ is 3 -connected, as guaranteed by Lemma 1.18, Then $G / e$ is planar. In the graph $G / e, e$ corresponds to a single vertex $z$, and $G / e \backslash z$ is 2 -connected. Hence the face that contains $z$ is a cycle $C$ of the graph, by Lemma 1.4. Now we try to uncontract $z$ while preserving planarity: this will fail either if $x$ and $y$ have at least three adjacent vertices in common, see Figure 9 , right, or if $x$ and $y$ are adjacent to four vertices that interleave, that is, if there exist vertices $x_{1}, x_{2}$ adjacent to $x$, and $y_{1}$ and $y_{2}$ adjacent to $y$, that are cyclically ordered as $x_{1}, y_{1}, x_{2}, y_{2}$ on $C$, see Figure 9 , left. In the first case, $G$ contains a subdivision of $K_{5}$, and in the second case it contains a subdivision of $K_{3,3}$.


Figure 9: The two cases to find subdivisions of forbidden graphs.

Proof of Kuratowski's theorem. One direction has already been proved: a graph that contains a subdivision of $K_{5}$ or $K_{3,3}$ cannot be planar. For the reverse direction, let $G$ be a graph that does not contain a subdivision of $K_{5}$ or $K_{3,3}$. Add edges to $G$ until it is maximal with this property, i.e., adding any new edge yields a subdivision of $K_{5}$ or $K_{3,3}$. By Lemma 1.19, the new graph $G^{\prime}$ is 3 -connected, and thus by Proposition 1.20, it is planar. Since $G$ is obtained from $G^{\prime}$ by removing edges, it is also planar, concluding the proof.

Now for the missing proofs, in a quite hasty style:
Proof of Lemma 1.18. Let $e=x y$ be an edge of $G$. If $G / e$ is not 3 -connected, then $G$ contains a vertex $z$ such that $G \backslash\{x, y, z\}$ is disconnected. Let $G_{1}$ be the smallest component of $G \backslash\{x, y, z\}$. Then $z$ is joined to $G_{1}$ by an edge $e_{1}=z u$. If $G / e_{1}$ is not 3 -connected, then there is a vertex $v$ such that $G \backslash\{z, u, v\}$ is disconnected, and since $G \backslash z$ is 2-connected, $v$ is in $G_{1} \cup\{x, y\}$ and the smallest component is a subgraph of $G_{1}$. So we continue until we get the edge that we are looking for.

Proof of Lemma 1.19. We pick a minimal counter-example: a smallest graph $G$ with at least 4 vertices that is not 3 -connected, containing no subdivision of a forbidden graph and such that the addition of any edge between non-adjacent vertices creates such a subdivision.

We claim that $G$ is 2 -connected. Otherwise, there is a cut vertex $x$, adjacent to vertices $y_{1}$ and $y_{2}$, each on one side. Adding an edge between $y_{1}$ and $y_{2}$ creates a subdivision $K$ of a forbidden graph. Since $K_{5}$ and $K_{3,3}$ are both 3-connected, all the vertices of degree at least 3 of $K$ are on the same side. Then the only way for $K$ to use the other side is if a path uses the edge $y_{1} y_{2}$ and comes back via $x$, but it might as well use the edge $y_{i} x$, and so $K$ was already present in $G$.

Now, if $G$ is not 3 -connected, there are two vertices $x$ and $y$ disconnecting it. We claim that the edge $x y$ must be in $G$. Otherwise, adding it creates a subdivision $K$ of a forbidden graph, and by the same previous argument, all of its vertices of degree at least 3 must be on the same side. Then we can reroute a path going through the other side with the edge $x y$, and so $K$ was already present in $G$.

So now we are in the case where $\{x, y\}$ disconnects $G$ into $G_{1}$ and $G_{2}$ and the edge $x y$ is in $G$. Adding some edge in $G_{i}$ creates a subdivision of a forbidden graph, which, by the same arguments as before, is contained in $G_{i}$. By minimality of $G, G_{i}$ is 3 -connected or has at most three vertices. By Proposition 1.20 , they are both planar, pick one of their embeddings. Let $z_{i} \neq x, y$ be a vertex of a face $F_{i}$ of $G_{i}$ bounded by $x y$. Note that $F_{i}$ must be a triangle, as otherwise one could add a diagonal within the face, preserving planarity and thus also the fact that it contains no subdivision of a forbidden graph. Then adding $z_{1} z_{2}$ creates a subdivision of a forbidden graph $K$.

At this stage, in the graph $G+z_{1} z_{2}$, there are now three vertices connecting $G_{1}$ and $G_{2}$, so it might be that the vertices of degree at least 3 of $K$ are not all contained in the same side. Actually, this has to be the case, otherwise, some path uses $z_{1} z_{2}$ and either $x$ and $y$,
and thus can be rerouted to $z_{i} x$ or $z_{i} y$. So $K$ has at least one vertex on each side. If $K$ is a subdivision of $K_{5}$, then there are four independent paths connecting those two vertices, which is not possible here. So $K$ is a subdivision of $K_{3,3}$, and by a similar argument, five of its vertices are on one side, say $G_{1}$ and one on the other side. But then we could as well put that last vertex at some place $p$ inside $F_{1}$ and connect it to $x, y$ and $z_{1}$. This graph $G_{1}+\{p x, p y, p z\}$ is planar by construction, but contains a subdivision of $K_{3,3}$. This last contradiction concludes the proof.

A minor of a graph $G$ is any graph obtained from a subgraph of $G$ by contracting a subset of its edges. If $G$ contains a subdivision of $H$ then $H$ is a minor of $G$, but the converse is not true (can you find a counter-example?). Nevertheless, a similar theorem to Kuratowski's hold for minors:

Theorem 1.21 (Wagner, 1937). A graph $G$ is planar if and only if none of $K_{5}$ and $K_{3,3}$ is a minor of $G$.

Proof. If $G$ is not planar, then by Kuratowski's theorem it contains a subdivision of $K_{5}$ or $K_{3,3}$ as a subgraph, and thus a $K_{5}$ or a $K_{3,3}$ minor. For the converse direction, a minor of a planar graph is planar, so this follows from Theorem 1.10 .

Wagner's theorem admits generalizations to graphs embedded on other surfaces, and in what constitutes a jewel of structural graph theory, to any minor-closed family of graphs.

The following theorem is a strong application of Kuratowski's theorem. Two edges are independent is they do not share a vertex.

Theorem 1.22 (Strong Hanani-Tutte theorem, 1934). Any drawing of a non-planar graph contains two independent edges that cross an odd number of times.

Conversely, if one can draw a graph so that all independent edges cross an even number of times, then the graph is planar.

Proof. We first claim that it suffices to prove the theorem for $K_{5}$ and $K_{3,3}$. Indeed, let $G$ be a non-planar graph. By Kuratowski's theorem, it contains a subdivision of $K_{5}$ or $K_{3,3}$. If the theorem is proved for those, then we can find a pair of disjoint paths in our graph $G$ that cross an odd number of times. This means that there exists a pair of independent edges that cross an odd number of times.

In order to prove the theorem for $K_{5}$ and $K_{3,3}$, we prove the stronger result that in any drawing of these two graphs, the sum of the number of crossings over all independent pairs of edges is odd. The reason is that this quantity mod 2 does not actually depend on the drawing. Indeed, given two drawings of $K_{5}$ or $K_{3,3}$, one can continuously morph one drawing into the other by moving the vertices so that they match, and then matching the edges. This can be done in a way that all the crossings stay transverse except at a finite number of singularities: when two edges start crossing, when the pattern of intersection between edges changes or when an edges passes on the other side of a vertex ${ }^{2}$. Note that whether a pair of edges crosses an odd number of times only changes when an edges passes to the other side of a vertex. But in the case of $K_{5}$ and $K_{3,3}$, the evenness of the sum of these crossings actually does not change: for $K_{5}$ it changes by 4 , and for $K_{3,3}$ it changes by 2 . We conclude the proof by exhbiting drawings of $K_{5}$ and $K_{3,3}$ where this sum is odd, and we have already done so in Figure 8 .

[^1]

Figure 10: A finger move for a pair $p=(e, v)$.

This theorem and its proof provide a polynomial-time algorithm to test whether a graph is planar. Let $G$ be a graph of which we want to test planarity. Let us consider the vector space $S$ over $\mathbb{F}_{2}$ whose basis consists of pairs of independent edges $(e, f)$ in $G$.

1. Start with any drawing of the graph $G$. Write down the vector $x(D) \in S$ of the number of crossings mod 2 of independent edges.
2. For every pair $p=(v, e)$ in the graph $G$, define a vector $u(p)=\sum_{f \text { incident to } v}(f, e) \in S$ : it contains a 1 for each pair of edges $(f, e)$ such that $f$ is incident to $e$, and 0 otherwise.
3. Denoting by $U$ the subspace spanned by the family $u(p)$. Test whether $x(D)$ belongs to $U$. If yes, output that the graph is planar, otherwise, output that it is not planar.

Lemma 1.23. This algorithm is correct and has polynomial complexity.

Proof. The reason why this algorithm is correct follows from Theorem 1.22. Indeed, starting from any drawing $D$ of $G$, one can change its vector $x(D)$ by exactly $u(p)$ by applying a finger move (see Figure 10) from the edge $e$ around the vertex $p$. If $x(D)$ belongs to $U$, there exists a sequence of finger moves that brings the vector of crossings $x(D)$ to zero mod 2 . Then by Theorem 1.22, the graph is planar. Conversely, if the graph is planar, one can morph the drawing similarly to the proof of Theorem 1.22 until it is planar, and this morphing gives a sequence of finger moves, and thus a family of vectors $u(p)$ showing that $x(D)$ belongs to $U$.

The third step can be solved in Gaussian elimination in cubic time (or there are faster algorithms, even more so since the underlying field is $\mathbb{F}_{2}$ ). Since $S$ has size $O\left(|E|^{2}\right)=O\left(|V|^{2}\right)$ (if the graph is not sparse, we can reject it straight away), this yields a complexity $O\left(|V|^{6}\right)$.

For the algebraic-minded reader, this algorithm actually amounts to computing an obstruction class in the equivariant cohomology of the deleted product (ask me in class if you are interested about what these words mean).

Famously:
Theorem 1.24 (Hopcroft-Tarjan). Given a graph $G$, one can, in linear time, determine whether $G$ is planar, and, if so, compute a combinatorial map of $G$ in the plane.

As far as I know, there is no known simple algorithm that runs in linear time, but see the references listed at the start of these lecture notes for planarity-testing algorithms closer in spirit to Theorem 1.24 .
Exercise 1.25. Let $G$ be a graph, and $C$ be a cycle in that graph. An edge in $E(G) \backslash E(C)$ connecting two vertices of $C$, or a subgraph of $G$ induced by a connected component of $G \backslash C$ and the edges attaching this connected component to $C$, are called bridges. If all of the endpoints of a bridge are contained between two endpoints of an other bridge, they are nested. Otherwise, they overlap. The overlap graph for $G$ and $C$ is the graph whose vertices are bridges and two bridges are connected if they overlap. Prove that a graph $G$ is planar if and only if for each cycle $C$, the bridge-overlap graph of $G$ with respect to $C$ is bipartite.


Figure 11: The Tutte magic
Solution: If $G$ is planar, $C$ is embedded as a Jordan curve, which defines an inside and an outside. Partition the bridges according to the side they are on. Bridges on the same side cannot be adjacent, otherwise by the definition of overlap this would induce crossings. So this is a bipartition. For the reverse direction, if a graph is non-planar, by Kuratowski it has a subdivision of $K_{5}$ or $K_{3,3}$ as a subgraph. Following the defintiions, one sees that it suffices to prove the result for $K_{5}$ and $K_{3,3}$, for which it is immediate.

### 1.6 The Tutte barycentric embedding

The previous algorithm allows us to test whether a graph is planar. In this section, we explain a very clean conceptual way to find a nice embedding of a planar graph, with just the knowledge of a facial cycle of the graph. We will assume that the graph is simple and triangulated, that is, that each facial cycle, except possibly the unbounded on\& ${ }^{3}$, is adjacent to three edges. Note that if one is given a planar embedding of a graph, one can add edges inside the faces of degree more than 3, subdividing those until the graph is triangulated. By the second exercise in the homework, a triangulated graph with at least 4 vertices is 3 -connected.

Then the Tutte embedding for a triangulated ${ }^{4}$ graph proceeds as follows

1. Pick a facial cycle of $G$, which will bound the outer face.
2. Nail all its vertices in some strictly convex position on a wall.
3. Let the other vertices move freely, and put springs in the place of the edges, each with some arbitrary positive stiffness.
4. Let the system relax.
5. Enjoy the nice planar embedding of the graph at equilibrium.

This is illustrated in Figure 11.
More formally, denoting the vertices not on the outer face by $V_{I}$, we have
Theorem 1.26 (Tutte, 1963). Let $G$ be a triangulated graph and $C$ be the facial cycle bounding the outer face. 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 $\lambda_{v w}$, for $w \in N(v)$ :

$$
\begin{equation*}
\forall v \in V_{I}, \quad \sum_{w \in N(v)} \lambda_{v w}(\tau(v)-\tau(w))=0 . \tag{1.1}
\end{equation*}
$$

Moreover, $\tau$ induces an embedding of $G$ by connecting the images of every pair of neighbor vertices with line segments.

[^2]When one takes symmetric weights $\lambda_{v w}=\lambda_{w v}$, the weights can be interpreted as the stiffness of springs and the system of equations of this theorem is the equilibrium condition of the physical system (obtained for example by differentiating the energy). Even more simply, taking all the weights equal to one puts each vertex that is not on the outer face at the barycenter of its neighbors. The theorem then stipulates that this very intuitive placement yields a planar embedding. In particular, one obtains:

Theorem 1.27 (Fàry-Stein-Wagner). Every planar graph can be drawn in the plane with straight line edges.

There are easier proofs of this theorem, but I like this one because of how conceptual it is. The following notes will be alas pretty devoid of pictures because it is very hard to picture things that do not happen.

Throughout the proof, we number the internal vertices from 1 to $k$ and the remaining ones from $k+1$ to $n$, and to homogenize notation, we consider the weight $\lambda_{i j}$ for two non-adjacent vertices to be zero (this corresponds to a spring that does not pull). We first look at the system of equations 1.1:

Lemma 1.28. The system of linear equations 1.1 has a unique solution.
Proof. The unknowns are the positions of the vertices $\tau(i)$ for $i$ from 1 to $k$.
(1.1) can be written

$$
\Lambda\left[\begin{array}{c}
\tau_{1} \\
\vdots \\
\tau_{k}
\end{array}\right]=\left[\begin{array}{c}
\sum_{j>k} \lambda_{1 j} \tau_{j} \\
\vdots \\
\sum_{j>k} \lambda_{k j} \tau_{j}
\end{array}\right]
$$

where $\tau_{i}$ stands for $\tau(i)$ and

$$
\Lambda=\left[\begin{array}{cccc}
\sum_{j=1}^{n} \lambda_{1 j} & -\lambda_{12} & \cdots & -\lambda_{1 k} \\
-\lambda_{21} & \sum_{j=1}^{n} \lambda_{2 j} & \cdots & -\lambda_{2 k} \\
\vdots & \vdots & \ddots & \\
-\lambda_{k 1} & -\lambda_{k 2} & \cdots & \sum_{j=1}^{n} \lambda_{k j}
\end{array}\right]
$$

Beware that the notations used are a bit misleading, as each $\tau_{i}$ is two-dimensional. To conclude, it suffices to prove that $\Lambda$ is invertible. It follows from the fact that it is close to being diagonally dominant, i.e., each diagonal entry is bigger than the sum of the other entries on the same row. A proof goes as follows: if there exists $x$ such that $\Lambda x=0$, take $x_{i}$ to be one of its components of maximal absolute value. Then the only way for $(\Lambda x)_{i}=\sum_{j \in N(i)} \lambda_{i j}\left(x_{i}-x_{j}\right)$ to be zero is that for all $j$ such that $\lambda_{i j}>0$, we have $x_{j}=x_{i}$. By connectivity of $G$, this extends to all the $x_{j}$, in particular those for vertices adjacent to a boundary vertex. But for those, the row is diagonally dominant, and thus $x_{j}=0$. We conclude that $\Lambda$ is invertible.

Note that this theorem is readily algorithmic, as solving this system directly gives the coordinates of the vertices.

For an interior vertex $v$ placed at $\tau(v)$, let $h: \mathbb{R}^{2} \rightarrow \mathbb{R}$ denote an affine function vanishing on $h(\tau(v))$. We say that $v$ is $h$-inactive if all its neighbors have the same $h$-value. Otherwise, it is $\boldsymbol{h}$-active. If a vertex is active, since it is a barycenter of its neighbors, it can be connected to a vertex with strictly bigger $h$-value. In turn, that vertex has to be active, and thus we can find a rising path from $v$ to an exterior vertex, i.e., a path where the $h$ values strictly increase. Similarly, one can find a falling path.


Figure 12: An inactive vertex leads to a $K_{3,3}$ subdivision.

Lemma 1.29. For any inner vertex $v, \tau(v)$ is contained in the strict convex hull of the outer vertices.

Proof. We first prove that $\tau(v)$ is contained in the convex hull of the outer vertices. Otherwise there exists a vertex $v$ and an affine form $h$ so that $h(\tau(v))>0$ but for all the outer vertices $w$, $h(\tau(w)) \leq 0$. Pick such a $v$ of maximal $h$-value. Then all of its neighbors must have the same value, i.e., $v$ is inactive. We continue inductively: all their neighbors must have the same value, until by connectivity we reach an exterior vertex, and a contradiction.

So let us now assume that for some affine form $h$, we have $h(\tau(v))=0$, while $h$ is negative for all the outer vertices except for two of them (since the starting embedding of the outer face is strictly convex). Denote by $H$ the subgraph formed by inner vertices in $h^{-1}(0) . H$ is nonempty since it contains $v$, and all the vertices in $H$ are $h$-inactive. By 3 -connectedness of $G, H$ is connected to the rest of $G$ in at least 3 ways, and is thus connected to at least one vertex such that $h$ is negative. This means that at least one vertex in $H$ is $h$-active, a contradiction.

Lemma 1.30. For any affine form $h$, all the interior vertices are $h$-active.
Proof. Assume otherwise that some interior vertex $v$ is inactive. Pick a vertex $w$ on the outer face so that $h(\tau(w))>h(\tau(v))$. Then, by 3 -connectedness, there exists three independent paths $P_{i}$ between $v$ and $w$. Denote by $u_{1}, u_{2}$ and $u_{3}$ the first active vertices on $P_{i}$. Note that the three of them satisfy $h\left(\tau\left(u_{i}\right)\right)=0$. Then we take three rising paths from $u_{i}$ to a vertex $y$ of maximal $h$-value. While the three rising paths may intersect, it is easy to see (and asked as an exercise just below this proof) that from them, we can always find a vertex $y^{\prime}$ such that there are three independent paths from $u_{i}$ to $y^{\prime}$. Similarly, we take three falling paths from $u_{i}$ to a vertex $z$ of minimal $h$-value, and out of them we can find three independent paths to a vertex $z^{\prime}$. Now we look at the paths between $v, u_{1}, u_{2}$ and $u_{3}$, as well as the six other paths we have constructed, and recognize a subdivision of $K_{3,3}$, see Figure 12. This contradicts the planarity of the graph, and proves the lemma.

Exercise 1.31. Let $u_{1}, u_{2}, u_{3}$ and $y$ be distinct vertices of a graph. Assume that for each $i$, there exists a path $P_{i}$ connecting $u_{i}$ to $y$ which avoids all the other $u_{j}$ for $j \neq i$. Then there exist three independent paths $P_{i}^{\prime}$ connecting $w_{i}$ to a common vertex $y^{\prime}$.

Now that every vertex has been proved to be active, we can prove that $\tau$ is an embedding locally:

Lemma 1.32. Let uvx and uvy be two facial triangles incident to an edge uv of $G$ not in $C$, then $\tau(x)$ and $\tau(y)$ are each on one side of any line through $\tau(u)$ and $\tau(v)$.

Note that the lemma does not assume that $\tau(u) \neq \tau(v)$.
Proof. As usual, pick $h$ an affine form whose kernel goes through $\tau(u)$ and $\tau(v)$. Pick two rising paths from $\tau(u)$ and $\tau(v)$ to an exterior vertex. (A subset of) These two rising paths and the edge $u v$, plus possibly some exterior edges, give rise to a cycle $K$ for which all the $h$-values are positive. As in the proof of Lemma 1.29, all the vertices interior to $K$ have positive $h$-value. In particular, one of $x$ or $y$ has positive $h$-value. Doing the same with falling paths, we get that one of $x$ or $y$ has negative $h$-value, which proves the lemma.

Say that a facial triangle is non-degenerate if its embedding has positive area. From this lemma, we conclude that all the facial triangles are non-degenerate. Indeed, this is clear for the triangles adjacent to outer edges, and this lemma shows that a triangle adjacent to a nondegenerate triangle is also non-degenerate. Therefore the claim follows by connectivity.

We can now conclude the proof of Theorem 1.26.
Proof. We have now proved that all the facial triangles are non-degenerate. Let us prove that the embeddings of two different triangles have distinct interiors. If not, let $p$ be a point contained in two triangles $t$ and $t^{\prime}$. Throw a ray from $p$ that avoids all the vertices $\tau(v)$. This ray exits $t$ at some edge, where it enters a different triangle $t_{2}$ (by Lemma 1.32). Continuing like that until we hit the outer face, we have a sequence $t, t_{2}, \ldots, t_{n}$ where $t_{n}$ is adjacent to the outer face. Likewise, the ray crosses $t^{\prime}$ and induces a sequence of triangles $t^{\prime}, t_{2}^{\prime}, \ldots, t_{m}^{\prime}$ where $t_{m}^{\prime}$ is adjacent to the outer face. But the ray crosses the outer face at a given point, where there is a unique triangle incident to the outer edge. So $t_{n}=t_{m}^{\prime}$, and then, following the ray back, we see that all of the triangles in the sequence are equal, including $t$ and $t^{\prime}$.

A nice corollary of this theorem is that we can always embed a planar 3-connected graph so that its faces are convex.

## 2 Some planar graph algorithms

This section takes a more algorithmic turn than the previous one, the overarching idea is to investigate how algorithmic problems become easier when the input graph is planar.

### 2.1 Efficient algorithms for planar graphs

Many algorithmic problems can be solved faster when the input graph is planar. This includes some NP-hard problems in general which can be solved in polynomial time in the planar case: for example Max-Cut, (uniform) Sparsest Cut, Feedback arc set, computing the Branchwidth if a graph. We will not study any of those (choices have to be made). Instead, we will look at a few problems where planarity allows for faster and conceptually simpler algorithms than in the general case, but not that dramatically so. At the risk of overly simplifying things, there are in my opinion two main reasons as to why planar graphs tend to be easier to handle algorithmically. The first reason is duality: sometimes a problem that is not easy to solve in the primal graph becomes much easier to graph in the dual graph. Sometimes juggling between both the primal and the dual setting allows to make good progress. The second reason is the existence of small separators, which allows for efficient divide and conquer approaches for a


Figure 13: A spanning tree and the spanning co-tree. These are sometimes called interdigitating tree.
lot of problems. Many proofs of the existence of these small separators (as ours does) rely on duality, so arguably there is only one reason here.

### 2.1.1 Minimum spanning trees

Let $G$ be a graph with a weight function $w: E \rightarrow \mathbb{R}^{+}$. A minimum spanning tree of a graph $G=(V, E)$ is a tree $T=\left(V, E^{\prime}\right)$ with $E^{\prime} \subseteq E$ that spans all the vertices of $G$ (this was actually already implied by the notations of the vertices) and such that it has minimal total weight $\sum_{e \in E^{\prime}} w(e)$ among all the spanning trees. Computing a minimum spanning tree is a fundamental primitive in algorithm design, and also an important practical problem in its own right: think about an electric company wanting to wire all the houses in a neighborhood at a minimal cost. For general graphs with $n$ vertices and $m$ edges, classical algorithms (e.g., Prim's or Kruskal's) run in time $O(m \log n)$ (note that in the sparse case, $m=O(n)$, but this is still not linear). Using some very fancy data structures, one can do better, but there is no known deterministic algorithm running in linear time.

In contrast, for planar graphs, one can compute a minimum spanning tree very easily in linear time. The key is the use of duality:

Theorem 2.1. If $G$ is a planar graph with $n$ vertices, one can compute a minimum spanning tree in time $O(n)$.

We start with exploring how spanning trees interact with duality, as illustrated in Figure 13 .
Lemma 2.2. Let $G$ be a planar graph and $G^{*}$ be its dual graph. Then if $T=\left(V, E^{\prime}\right)$ and $E^{\prime} \subseteq E$ is a spanning tree, then $T^{*}=\left(F^{*},\left(E \backslash E^{\prime}\right)^{*}\right)$ is also a spanning tree, called the co-tree. If one is minimal, the other is maximal.

Proof. If $T$ is a spanning tree, it does not contain any cycle, which happens if and only if its complement $\mathbb{R}^{2} \backslash E^{\prime}$ is connected, by the Jordan curve theorem. But $\mathbb{R}^{2} \backslash E^{\prime}$ is connected if and
only if $T^{*}$ is connected. Indeed, any two points in $\mathbb{R}^{2} \backslash E^{\prime}$ are in faces of $G$, and thus can be connected without crossing edges of $E^{\prime}$ if and only if those faces can be connected in the dual graph where one removed the edges dual to those of $E^{\prime}$. Since $T$ and $T^{\prime}$ are both connected, they are both acyclic and thus are both trees. If one of them, say $T$, is minimal, then $T^{*}$ is maximal, since otherwise one could increase the weight of $T^{*}$ which would mechanically decrease the weight of $T$.

Exercise 2.3. Prove that any tree with $v$ vertices and $e$ edges satisfies $v-e=1$. Deduce from Lemma 2.2 a new proof of Euler's formula.

The following exercise is a good way to get familiar with these interdigitating trees.
Exercise 2.4. Let $G$ be a graph embedded into the plane. Assume that $G$ is cellular (each face is homeomorphic to an open disk) and specified by its combinatorial map. Suppose that every face has a real weight. Let $T$ be a spanning tree of $G$ (a subgraph of $G$ with the same vertex set as $G$ that is a tree). Give an algorithm (as efficient as possible) that, for all edges $e$ of $G$ not lying in $T$, computes the sum of the weights of all the faces inside the cycle formed by $e$ and the unique path in $T$ joining its endpoints. Give the complexity of your algorithm.

Sketch of solution: We consider the graph of $G^{*}$ obtained by dualizing the edges not in $T$. It is a tree because of the previous lemma. We root it at the outer face. Then the exercise amounts to computing for each edge the sum of the vertex weights below it. This is done in linear time by going up from the leaves of the tree (dynamic programming).

Observe that when we contract an edge in a graph, the corresponding edge in the dual graph gets removed, and vice-versa.

The following is an easy consequence of Euler's formula.
Lemma 2.5. In a planar graph $G$, there is either a vertex of degree at most 3, or a face of degree at most 3 .

Proof. Assume otherwise. Then we can double count edges in two different ways and get $4 v \leq 2 e$, and $4 f \leq 2 e$. Then $v-e+f \leq 0$ contradicting Euler's formula.

We now have all the tools to describe our algorithm.
Proof of Theorem 2.1. The algorithm is based on alternating actions between the primal and the dual graph. We initialize the set of edges $E^{\prime}$ that we pick at the empty set, and:

- Let $v$ be a vertex of the graph $G$. If $v$ is only surrounded by loops, then the solution is the trivial empty tree. Otherwise, at least one non-loop edge is adjacent to $v$. Among all these non-loop edges, one of minimal weight $e$ necessarily belongs to the minimal spanning tree: otherwise, one could add it and remove another edge. So we can take it, add it to $E^{\prime}$ and recurse on $G / e$.
- Let $f$ be a face of the graph $G$, and thus a vertex of the dual graph $G^{*}$. If all the edges adjacent to $f$ in $G^{*}$ are loops, then $G^{*}$ has a single face, thus $G$ is a tree, and the spanning tree is $G$ itself. Otherwise, the dual of an edge $e$ of maximal weight incident to $f$ necessarily belongs to the maximal spanning co-tree, and thus $e$ does not belong to the minimum spanning tree. So we can recurse on $G \backslash e$.

Each of these two actions removes an edge in one way or another from the graph that we consider, so the number of recursions is $O(n)$. For each of the two actions, the cost of finding
which edge to contract or remove is of the order of the degree of the vertex or the face that we consider. So if we always pick a vertex or an edge of degree at most 3 (provided by Lemma 2.5), this will take constant time. Note that contracting (respectively removing) an edge adjacent to a vertex (respectively a face) of constant degree means updating $O(1)$ flags in the representation, so the recursive call can be made in constant time.

So there remains to explain how to find the vertex or face of low degree provided by Lemma 2.5 in constant time. We will amortize this search, i.e., we will prove that the total time spent looking for these is $O(n)$, and thus it is $O(1)$ in average per round. In order to do that, we first compute a list $L$ containing all the vertices and the faces of degree 3 of the initial graph. This takes linear time by traversing both the primal and the dual. The list $L$ will be updated throughout the algorithm so that it always contains the list of vertices or faces that may have degree at most 3 . When contracting or removing an edge, the degree of the four adjacent vertices and faces can change, so we add them all to the list $L$. Since there are $O(n)$ iterations, this process adds $O(n)$ elements to the list $L$ throughout the algorithm.

Now, whenever we want to take an action, we look at the first element of the list $L$. If it does not exist anymore, or if it has degree more than 3, we remove it from the list, and continue. By Lemma 2.5, we always end up finding a vertex or a face. Since $O(n)$ vertices and faces were added to the list, we remove $O(n)$ of those throughout the algorithm. So in total, the search procedure takes $O(n)$ time, which proves our amortized complexity and finishes the proof of the algorithm.

Note that this algorithm crucially relies on the fact that we work with non-simple graphs (even if the initial graph is simple, it becomes non-simple after contractions). This is why we rely on Lemma 2.5 and duality and not merely Exercise 1.12 ,

### 2.1.2 Planar separators

A (balanced) separator of a graph is a subset of the vertices $S \subseteq V$ so that each component of $V \backslash S$ contains at most two thirds of the vertices of the graph. The size of the separator is the number of vertices of $S$. Separators of small size are key to algorithm design, as they allow for very intuitive divide and conquer algorithms: cut the graph into two parts of roughly equal size, solve your problem recursively on both sides and glue back the solutions. Hopefully the gluing back is facilitated by the small size of the separators. We shall see a specific example of that approach.

Here we prove that planar graphs have (somewhat) small separators.
Theorem 2.6. Let $G$ be a planar graph with $n$ vertices. Then one can compute in $O(n)$ time a separator for $G$ of size $O(\sqrt{n})$.

This is tight, as one can see by looking at a $\sqrt{n} \times \sqrt{n}$ grid.

Proof of Theorem 2.6. In the proof, we will be intentionally vague on the constants. We can safely assume that the graph $G$ is triangulated: one can clearly add edges until the graph is triangulated, which takes linear time, and then a separator for the new graph is also a separator for the old graph.

We start from an arbitrary vertex $r$ and compute a breadth-first search tree $T$ rooted at $r$. The level of a vertex is the distance from that vertex to the root in $T$. Denote by $L_{i}$ the set of vertices of level $i$. For any edge $e$ that is not in $T$, denote by $C_{e}$ the cycle induced by the edge $e$ and the tree. We first prove that one of the levels $L_{i}$ is a separator, and that one of the cycles $C_{e}$ is a separator.


Figure 14: An edge separating the dual co-tree into balanced parts induces a balanced separator in the primal graph.


Figure 15: The separator is made of $L_{1}, L_{2}$, and the subpaths of the fundamental cycle between them. It cuts the graph into at most four small components.

- At least one the $L_{i}$ is a balanced separator. Indeed, if we sort all the vertices based on their level and look at the level $L_{m}$ containing the $n / 2$ th vertex, it is necessarily a balanced separator.
- At least one of the $C_{e}$ is a balanced separator. Indeed, look at the co-tree $C$ corresponding to $T$. Root $C$ at the leaf of your choice. Below that leaf, the tree $C$ is a binary tree because $G$ is triangulated. We walk down from the root, always picking the edge going towards the larger subtree. We stop when we reach the first vertex $v$ for which the subtree has less than $2 n / 3$ dual vertices. Then the edge $v w$ (where $w$ is the parent of $v$ ) cuts the tree into two subtrees, each of which contains at most $2 n / 3$ dual vertices. Therefore, the cycle $C_{(v w)^{*}}$ separates $G$ into two disks, each containing at most two thirds of the faces. From Euler's formula and the fact that the graph is triangulated it follows that this cycle is a balanced separator. See Figure 14 for an illustration.

Both of these balanced separators can be computed in time $O(n)$. If one of them is short, i.e., has size $O(\sqrt{n})$, we win. Otherwise, let us assume that both $L_{m}$ and $C_{(v w)^{*}}$ are long. Then in particular the depth of the breadth-first search tree is bigger than $\sqrt{n}$. We look at at the levels $L_{i}$ for $i$ between $m$ and $m+\sqrt{n}$. Since there are $\sqrt{n}$ levels, one of the levels has less than $\sqrt{n}$ elements, and we call it $L_{2}$. Likewise, looking at the levels $L_{i}$ for $i$ between $m$ and $m-\sqrt{n}$, we find a level $L_{1}$ with at most $\sqrt{n}$ elements. The cycle $C_{(v w)^{*}}$ connects $L_{1}$ to $L_{2}$ with two paths $\alpha$ and $\beta$, which both have length $O(\sqrt{n})$. We claim that $L_{1} \cup L_{2} \cup \alpha \cup \beta$ is the balanced separator that we have been looking for. First, it is short. Second, it cuts $G$ into at least four connected components (see Figure 15):

- The component between the root and $L_{1}$ has at most $n / 2$ vertices since it is contained in one of the regions cut by $L_{m}$.
- The component after $L_{2}$ has at most $n / 2$ vertices since it is contained in one of the regions cut by $L_{m}$.
- There are two components between $L_{1}$ and $L_{2}$, separated by $\alpha$ and $\beta$. Each of them has at most $2 n / 3$ vertices since they are contained in one of the components cut by $C_{(v w)^{*}}$.

All the steps in this proof are readily algorithmic and can be implemented in $O(n)$ time.

Here is a simple algorithmic application of this theorem, using a divide and conquer framework.

Theorem 2.7. In a planar directed graph, one can find the shortest cycle in time $O\left(n^{3 / 2}\right)$.

The naive algorithm in general graphs consists in computing a shortest path tree at each vertex, which has at least quadratic complexity.

Proof. When one cuts the graph using planar separators, the shortest cycle is either on one side, or on the other side, or it crosses the separator. The first two cases can be handled recursively, and the last case can be handled with $O(\sqrt{n})$ computations of a shortest path tree. More precisely, the algorithm proceeds as follows:

1. Find a planar separator $S$ cutting $G$ into two balanced subsets $A$ and $B$,
2. Recursively search $A$ and $B$,
3. For each $s$ in $S$, compute the shortest cycle through $s$ using a breadth-first search tree,
4. Return the shortest of the cycles output in steps 2 . and 3 .

Step 3 takes $O\left(n^{3 / 2}\right)$ time, and this dominates the complexity of the algorithm, even in the recursive calls. Indeed, the complexity of this algorithm obeys a recurrence of the form $f(n)=O\left(n^{3 / 2}\right)+2 f(2 n / 3)$. Applying blindly the master theorem to this does not yield the result (hence why I got confused in the classroom). Instead, one can leverage the fact that $|V(A)|+|V(B)| \leq n$. Therefore we get the stronger recurrence $f(n)=O\left(n^{3 / 2}\right)+f\left(n_{1}\right)+f\left(n_{2}\right)$, where $n_{1}+n_{2} \leq n, n_{1} \leq 2 / 3 n$ and $n_{2} \leq 2 / 3 n$. We can conclude using a stronger version of the master theorem. Or we can do it by hand: by convexity of the function $x \mapsto x^{3 / 2}$, given these assumptions we have that $n_{1}^{3 / 2}+n_{2}^{3 / 2} \leq(2 n / 3)^{3 / 2}+(n / 3)^{3 / 2}$. Then, if $c_{1}$ is the constant such that Step 3 takes at most $c_{1} n^{3 / 2}$ time, we want to show that there exists a constant $c_{2}$ such that $c_{1} n^{3 / 2}+c_{2}\left((2 n / 3)^{3 / 2}+(n / 3)^{3 / 2}\right) \leq c_{2} n^{3 / 2}$. This is the case because $(2 / 3)^{3 / 2}+(1 / 3)^{3 / 2}<1$.

Similarly, but with much more work, any problem involving shortest paths can be sped up on planar graphs using planar separators. In particular, a theorem that we will not prove is the following:

Theorem 2.8 (Henzinger, Klein, Rao, Subramanian 1997). On an edge-weighted planar graph, a shortest path tree from any given vertex can be computed in linear time.

[^3]More generally, many problems for which the best algorithm runs in time $2^{O(n)}$ or $n^{O(n)}$ can be sped up on planar graphs to obtain algorithms running in time $2^{O(\sqrt{n})}$ or $n^{O(\sqrt{n})}$. This is sometimes called the Square root phenomenon. A good framework to prove such upper bounds is to introduce the notion of treewidth and use planar separators to prove that planar graphs have treewidth $O(\sqrt{n})$.

### 2.2 Minimum cut

Our last foray into algorithms for planar graphs concerns the computation of a minimum cut: given two vertices $s$ and $t$, called terminals, on a planar graph $G=(V, E)$, we want to compute the minimum set of edges $X$ so that removing $X$ from $E$ separates $s$ and $t$.

Theorem 2.9. Let $G=(V, E)$ be a cellularly embedded edge-weighted planar graph, and $s$ and $t$ be two distinct vertices of $G$. Then the problem of computing a minimum $s-t$-cut of $G$ can be solved in $O(n \log n)$ time.

The basic idea of an efficient algorithm for min-cut on planar graphs is to look at it through the lens of duality: a cut on the primal graph separating $s$ and $t$ dualizes to a cycle in the dual graph separating the faces dual to $s$ and $t$.

Proposition 2.10. $X \subseteq E$ is an $(s, t)$-cut in $G$ if and only if $X^{*}$ contains the edge set of some cycle of $G^{*}$ separating $s$ and $t$.

This proposition is considered obvious pretty much anywhere, but we will prove it, if only to emphasize that it does not hold on other surfaces (as usual, we will use the Jordan curve theorem).

Proof. The reverse direction is straightforward: if $X^{*}$ contains a cycle of $G^{*}$ separating $s$ and $t$, then any path in $G$ between $s$ and $t$ must cross this cycle, and thus $X$ is an $(s, t)$-cut.

For the forward direction, we take $X$ an $(s, t)$-cut in $G$, and choose $C$ to be an inclusionwise minimal subset of $X$ that is also an $(s, t)$-cut in $G$. We show that $C^{*}$ is a cycle separating $s$ and $t$. By minimality of $C$, each vertex of $G$ can be connected to either $s$ or $t$ without taking edges of $C$, and the two cases are exclusive. We label vertices with " S " or " T " depending on which one they are connected to. Moreover, for any edge in $C$, its two endpoints cannot have the same label, and for any other edge, its endpoints have the same label. So if we look at a face $f$ of $G$ adjacent to an edge of $C$, the labels on the facial walk on the face alternate only when the edge is in $C$, and thus there is an even number of edges of $C$ adjacent to $f$. So $C^{*}$ is an Eulerian subgraph of $G^{*}$, that is, a subgraph where each vertex has even degree. Pick any cycle in that subgraph. By the Jordan curve theorem, it separates $s$ and $t$. By minimality, $C^{*}$ is that cycle.

This proposition transforms a combinatorial problem into a topological one, namely, finding in the dual graph the shortest cycle enclosing a given face and not another given one. Even more topologically, if we remove the faces $s^{*}$ and $t^{*}$ from the dual graph, we obtain a surface homeomorphic to an annulus, and we look for the shortest cycle that goes around this annulus. However, this runs into an interesting technical issue: is $s^{*}$ and $t^{*}$ are adjacent, then removing $s^{*}$ and $t^{*}$ does not actually yield an annulus. Morally, this should not pose a problem: we just want to add an infinitesimally small buffer between $s^{*}$ and $t^{*}$ and we will be fine. One way to do this is to enlarge a tiny bit the set of curves that we look at. When working on the primal graph, we generally work with walks on the primal graph. When working on the dual graph, we work with walks on the dual graph, which correspond by duality to closed curves that are


Figure 16: The annulus obtained after removing a small disk around $s$ and $t$.
in general position with respect to $G$, i.e., they do not meet the vertices of $G$ and cross the edges of $G$ transversely. So our solution is to directly work in this setting of curves in general position with respect to $G$. The length of such a curve is defined to be the number of edges of $G$ that it crosses. Note that this is a bit more general than just looking at walks on the dual graph: now we can define a pair of small curves in general position around $s$ and $t$ that are disjoint, even if $s$ and $t$ are adjacent in $G$. Yet from an algorithmic perspective, all the curves in general position can be pushed on the dual graph in a way that does not change the length, so any computation, for example shortest paths, can be made in the dual graph. In this new setting, Proposition 2.10 becomes:

Proposition 2.11. Let $\gamma$ be a simple closed curve in general position with respect to $G$, that separates $s$ from $t$ and that has minimal length among all such curves. Then the set of edges crossed by $\gamma$ is a minimum ( $s, t$ )-cut in $G$.

Recall that a simple closed curve on a sphere is an injective map $\gamma: \mathbb{S}^{1} \rightarrow \mathbb{S}^{2}$.
Proof. Any path connecting $s$ to $t$ in $G$ crosses $\gamma$, thus the set of edges crossed by $\gamma$ is an $(s, t)$-cut. Conversely, a minimum ( $s, t)$-cut dualizes to a cycle in $G^{*}$ separating $s$ and $t$, which corresponds to a simple closed curve $\gamma$ in general position with respect to $G$, that separates $s$ from $t$.

Now, we remove a small disk around $s$ and $t$, getting an annulus $A$, and a portion of the graph $G$ embedded on this annulus as in Figure 16. How do we compute a shortest closed curve that goes around the annulus $A$ ? We can fix one point on each boundary, and compute a shortest path between these two points, and call it $p$. Then we show that some shortest closed curve does not cross $p$ more than once, and thus can be found by a shortest path computation.
Lemma 2.12. Some shortest closed curve separating the two boundaries of $A$ is simple and crosses $p$ exactly once.


Figure 17: The two shortcutting arguments of Lemma 2.12 and Lemma 2.13 .

Proof. This is illustrated in the left of Figure 17. Let $\gamma$ be such a curve. If we cut $A$ along $p$, we obtain a topological disk $D$, where the path $p$ got cut into two paths $p^{\prime}$ and $p^{\prime \prime}$ on the boundary. The curve $\gamma$, once cut on $D$, contains a simple path $q$ connecting $p^{\prime}$ to $p^{\prime \prime}$, otherwise $\gamma$ would not separate $s$ from $t$. Now we reglue $D$ along $p$, and connect the two endpoints $q_{1}$ and $q_{2}$ of $q$ by running parallel to $p$. We obtain a closed curve that is simple, crosses $p$ exactly once and is no longer than $\gamma$ since $\gamma$ was connecting $q_{1}$ to $q_{2}$ as well.

From this we get a naive quadratic algorithm. We compute the shortest path $p$, which has some length $k$. We pick points $v_{0}, \ldots, v_{k}$ on this path such that the subpath $\left[v_{i}, v_{i+1}\right]$ has length one. Then, cutting along $p$, each vertex $v_{i}$ gets duplicated into $v_{i}^{\prime}$ and $v_{i}^{\prime \prime}$, and we compute all shortest paths between each $v_{i}^{\prime}$ and $v_{i}^{\prime \prime}$. Following all the lemmas, one of them is the dual of a minimal cut. Since shortest paths in the dual graph can be computed in linear time (cf. the unproved Theorem 2.8), this takes $O\left(n^{2}\right)$ time.

We can speed this up doing some divide-and-conquering.
Lemma 2.13. Let $(x, y, z)$ be points on $p$, appearing in this order. When cutting $A$ along $p$, these get duplicated into $\left(x^{\prime}, y^{\prime}, z^{\prime}\right)$ and ( $\left.x^{\prime \prime}, y^{\prime \prime}, z^{\prime \prime}\right)$. If $\gamma_{x}$ and $\gamma_{z}$ are disjoint shortest paths between $x^{\prime}$ and $x^{\prime \prime}$, respectively $z^{\prime}$ and $z^{\prime \prime}$, then some shortest path between $y$ and $y^{\prime}$ does not cross $\gamma_{x}$ nor $\gamma_{z}$.

Proof. This is illustrated in the right of Figure 17. Say that $\gamma_{y}$ crosses $\gamma_{x}$, then it crosses it in at least two points. Let $a$ and $b$ be the first and the last crossing points when going from $y^{\prime}$ to $y^{\prime \prime}$. Then we can replace whatever $\gamma_{y}$ was doing between $a$ and $b$ with the subpath of $\gamma_{x}$ between $a$ and $b$ (or more precisely, some shortest paths infinitesimally close to it). Since $\gamma_{x}$ is a shortest path, the new curve is at most as along as $\gamma_{y}$. Doing the same for the crossings with $\gamma_{z}$ concludes the proof.

This suggest the following recursive approach. Having computed our shortest path $p$ of length $k$ between the two boundaries of $A$, if $k>2$,

1. we pick a vertex $v:=v_{\lfloor k / 2\rfloor}$, cut along $p$ and compute on $D$ a shortest path $p$ between the two vertices $v^{\prime}$ and $v^{\prime \prime}$ corresponding to $v$,
2. we reglue $D$ into $A$, and the shortest path $p$ is a simple closed curve $\gamma$. We cut $A$ along $\gamma$ and get two annuli $A_{1}$ and $A_{2}$.
3. we recurse on $A_{1}$ and $A_{2}$ and output the shortest of the two solutions.

We stop the recursion when we reach one of the following two base cases for the recursion: (1) when $k \leq 2$, we can compute the shortest closed curve by brute forcing the problem in $O(n)$ time as in the quadratic algorithm, and (2) if there exists a face adjacent to both boundaries, we can compute a shortest cycle going through that face directly in time $O(n)$.

Here again, this algorithm would be quite a bit more annoying to describe purely in the dual graph, as the shortest path $p$ might be following the boundary of an annulus, and thus when cutting along $\gamma$ in step 2 we do not obtain annuli. This can be dealt with by appropriately subdividing in the correct places, which, when thinking about it, is exactly what this algorithm does - but I believe that the description using curves in general position is more transparent (this perspective is directly taken from Éric Colin de Verdière's notes).

To conclude the proof of the theorem, we establish the correctness and the complexity analysis:

Proof of Theorem 2.9. The algorithm terminates in $O(\log n)$ recursion levels since the length of the path $p$ in the recursive calls shrinks by a half at each recursive call. The correctness follows from Proposition 2.11, Lemmas 2.12 and 2.13 , since they prove that some minimal cut will be dual to the shortest cycle that our recursive calls will find.

The proof of correctness is not as immediate as one could expect, as recursive calls share a lot of structure with their parent: for example it looks like the same non-boundary edge of $G$ might be cut several times by the shortest paths in the recursion, and thus appear in several of the annuli. But note that if an edge $e$ is cut into subedges $e_{1}, \ldots, e_{\ell}$, then in all the recursive calls involving the annulus between $e_{i}$ and $e_{i+1}$, the recursion stops directly, since there is a face adjacent to both sides of the annulus. Therefore, only $e_{1}$ and $e_{\ell}$ actually lead to recursive subcalls, and thus at each level of the recursion, throughout all branches of the recursion tree, each non-boundary edge only appears a constant number of times.

Similarly, the number of boundary edges on an annulus is at most twice the number of non-boundary edges, and thus the total number of boundary edges at a level of a recursion is $O(n)$.

Each computation in steps 1 and 2 of the algorithm takes time linear in the complexity of the annulus at this stage. There are $O(\log n)$ levels, and by the previous observations, each them costs $O(n)$ time in total, so the total complexity is $O(n \log n)$.

## 3 Surfaces

We now turn our attention to other topological spaces of dimension 2, and the graphs embedded thereon.

### 3.1 Definition and classification

A surface is a topological space locally homeomorphic to the plane, i.e., every point has an open neighborhood homeomorphic to $\mathbb{R}^{2}$. In this course, we restrict our attention to compact and connected surfaces. Behind this very wide-looking definition, there are actually only a few (yet infinitely many) different surfaces, and their classification is the topic of this subsection. Examples of surfaces are the sphere, the torus, and the projective plane, see Figure 18. As before, we consider spaces up to homeomorphism. It turns out that these are exhaustive in the sense that all the other surfaces can be built by gluing those together.

In the first chapter, we saw how a connected graph embedded on the plane naturally cuts the plane into disks. Similarly, we will investigate and classify surfaces by the graphs embedded on them. The definition is the same as in the planar case: a graph $G=(V, E)$ is embeddable on a surface $S$ if there exists an injective map $f: G \rightarrow S$, i.e., $G$ can be drawn without crossings


Figure 18: A sphere, a torus, a Klein bottle, and a projective plane (in the last one, the right boundary is identified to the left boundary in the direction indicated by the arrow).
on $S$. A graph is cellularly embedded if every connected component of $S \backslash G$ is homeomorphic to an open disk.

Theorem 3.1 (Kerékjártó-Radó). On any compact connected surface, there exists a cellularly embedded graph.

Start of a proof. By definition of a surface, every point has an open neighborhood homeomorphic to the plane, or an open two-dimensional disk. By compactness, we can extract a finite covering out of this collection of open sets. Their boundaries form a family of simple closed curves embedded on the surface. If they cross finitely many times, the interiors of their intersections can be seen to be bounded by simple closed curves, and thus, by the Jordan-Schoenflies theorem (which applies in the small neighborhoods since they are homeomorphic to the plane), they bound disks. Hence we have cut the surface into disks, and we win. If they cross infinitely many times, we can tinker with them to reduce to the case of finite crossings.

This is more subtle than it appears: this sketch can be made correct in two and three dimensions, but not in higher dimensions, as there are example of higher dimensional manifolds that can not be triangulated (a higher dimensional analogue of our cutting into disks).

While we are not proving theorems, let us add the following one. A triangulation is a cellularly embedded graph where every face has degree 3. A refinement of a triangulation is obtained by either subdividing a face (adding a vertex in that face and edges adjacent to each vertex incident to that face), or an edge (adding a vertex on an edge and edges adjacent to each non-adjacent vertex in the two incident faces).

Theorem 3.2. Any two triangulations of a surface have a common refinement.
The same not-a-proof (does not) work(s): overlay the two triangulations, if they have a finite number of intersections, we are done. Otherwise, we tinker things. Once again, this is quickly not true in higher dimensions.

For the reader disappointed in these two omissions (and I understand them), one alternate way to view this is that we are only looking at surfaces that are defined as disks glued together in a finite way, with two surfaces being isomorphic if they have a common refinement (this can be taken as a definition). Then the two omitted proofs show that we obtain the same surfaces as with the more topological definition, but if we are content within the purely combinatorial world, we do not need this equivalence.

We describe a graph cellularly embedded on a surface via polygonal schemata, which encodes the way that the disks are glued together. Starting from a cellularly embedded graph, we first name the edges and orient them in an arbitrary way. Each facial walk induces a word (considered up to cyclic permutation), where an edge $e$ taken in the reverse direction is denoted by $\bar{e}$ or $e^{-1}$. Each such facial walk is called a relation. The polygonal scheme is the data of


Figure 19: Attaching handles or Möbius bands to a sphere.
all these facial walks. Reciprocally, if one is given a collection of words where each letter of the alphabet appears exactly twice, one can interpret those as disks glued to each other, which together form a surface.

With Theorem 3.1 in hand, we are now ready to classify surfaces:
Theorem 3.3. Every compact connected surface is homeomorphic to a surface given by one of the following polygonal schemata, each made of a single relation:

1. $a \bar{a}$ (the sphere),
2. $a_{1} b_{1} \overline{a_{1}} \overline{b_{1}} \ldots a_{g} b_{g} \overline{a_{g}} b_{g}$ for some $g \geq 1$,
3. $a_{1} a_{1} \ldots a_{g} a_{g}$ for some $g \geq 1$.

The surfaces of the first and second category are called orientable, those of the third category are non-orientable. The integer $g$ is the genus of the surfaces. The second case corresponds to $g$ tori glued together (this is called a connected sum), or equivalently, a sphere to which we have glued $g$ handles. The third case corresponds to $g$ projective planes glued together, or equivalently, to a sphere with $g$ disks removed on which we have glued $g$ Möbius bands. See Figure 19, and Figure 20 for the polygonal schemata depicted on the surfaces. Many non-trivial homeomorphisms are hidden behind this apparently simple classification, and it almost equally simple proof: for example, the theorem implies that the connected sum of two projective planes is homeomorphic to a Klein bottle, and that the connected sum of a torus and a projective plane is homeomorphic to the connected sum of three projective planes.

Proof. Let $S$ be a compact connected surface, and let $G$ be a graph embedded on $S$, which exists by Theorem 3.1. We iteratively remove each edge adjacent to two different face, until there is just a single face. For each edge adjacent to two different vertices, we contract it, and keep the multiple edges or loops that might result, until there is just a single vertex. We now have a graph with a single vertex and a single face embedded on $S$. If there are no more edges, by uncontracting once we obtain a sphere as in case 1 . of the theorem. Thus there is at least one edge.

The single face induces a polygonal scheme with a single relation. The rest of the proof aims at transforming this single relation into one of the two forms of Theorem 3.3 via cut-and-paste operations:


Figure 20: Polygonal schemata of the orientable and non-orientable surfaces. In the nonorientable cases, $X$ denotes a disk on which a Möbius band has been glued.




Figure 21: From $a P a Q$ to $b b P \bar{Q}$.

- If the polygonal scheme has the form $a P a Q$ where $P$ and $Q$ are possibly empty words, then we can transform it into $b b P \bar{Q}$ by adding a new edge and removing $a$, see Figure 21 . Inductively, we conclude that each pair of symbols with the same orientation appears consecutively in the polygonal scheme.
- If the polygonal scheme has the form $a U \bar{a} V$, then $U$ and $V$ must share an edge $b$ since otherwise $G^{\prime}$ would have more than one vertex. By the preceding step, $b$ must appear in opposite orientations in $U$ and $V$, so we have the form $a U \bar{a} V=a P b Q \bar{a} R \bar{b} S$. This can be transformed into $d c \bar{d} \bar{c} P S R Q$, as pictured in Figure 22. Inductively, at the end of this step the relation is a concatenation of blocks of the form $a a$ or $a b \bar{a} \bar{b}$. If all the blocks are of one of these types, we are in case 2 or 3 and we are done.
- Otherwise, the relation has a subword of the form $a a b c \bar{b} \bar{c}$. This can be transformed into $\bar{c} \bar{c} \bar{b} \bar{d} \bar{b} \bar{c}$, as in Figure 23, and then using the first step again this can be transformed into eeffgg. Inductively, we obtain a relation of the form 3.

This concludes the proof.
The Euler characteristic of a surface $S$ is defined to be $v-e+f$, where $v, e$ and $f$ are the numbers of vertices, edges and faces of a cellularly embedded graph on $S$. The fact that this actually does not depend on the graph is the object of the following proposition:
Proposition 3.4. For any graph $G$ cellularly embedded on $S$ with $v$ vertices, e edges and $f$ faces, the value $v-e+f$ is the same.

It is tempting to prove this using Theorem 3.3: any cellularly embedded graph can be transformed to one of the three graphs stipulated by the theorem, in a way that does not





Figure 22: From $a P b Q \bar{a} R \bar{b} S$ to $c d \bar{c} \bar{d} P S R Q$.


Figure 23: From $a a b c \bar{b} \bar{c}$ to $\bar{d} \bar{c} \bar{b} \bar{d} \bar{b} \bar{c}$.
change the Euler characteristic. But to conclude, we need to prove that two different outputs of the theorem are not homeomorphic, which we have not done yet.

Proof. We pick two graphs cellularly embedded on $S$. We add edges until they are both triangulations, which can easily be seen not to change the Euler characteristic. By Theorem 3.2, they have a common refinement. The proof follows from the fact that subdividing edges or faces does not changes the Euler characteristic.

A look at the graphs corresponding to the cases 1,2 and 3 of Theorem 3.3 shows that the sphere has Euler characteristic 2 (which should not come as a surprise after the first chapter), the orientable surface of genus $g$ has Euler characteristic $2-2 g$, and the non-orientable surface of genus $g$ has Euler characteristic $2-g$.

Proposition 3.5. A surface is orientable if and only if for any cellularly embedded graph $G$, the boundaries of the faces can be oriented so that each edge appears in opposite directions in the two adjacent faces.

Proof. First note that the property of the proposition is invariant under cut and pasting, as well as edge or face subdivision. Since the polygonal scheme of orientable surfaces satisfies the proposition, this proves the first implication. For the reverse implication, it suffices to observe that for non-orientable surfaces, such orientations of the faces are not possible.

As a corollary, a non-orientable surface is never homeomorphic to an orientable surface. Since all the orientable surfaces (respectively non-orientable surfaces) have a different Euler characteristic and the Euler characteristic is a topological invariant, this shows that all the
surfaces given by Theorem 3.3 are pairwise non-homeomorphic. Therefore we have found them all, and they are all different.
Exercise 3.6. 1. Provide an example of two cellular embeddings of the same graph $G$ on two different surfaces.
2. What is the surface described by the polygonal schemes $a b c \overline{a b c}$ ? $a b c d \overline{a b c d}$ ?

We conclude with a few remarks:
Surfaces with boundary: It is often very convenient to also consider surfaces where points can be homeomorphic either to $\mathbb{R}^{2}$ or to an open half-space $\mathbb{R}^{2} \cap\{x \geq 0\}$. These are surfaces with boundary, where the boundary is the set of points homeomorphic to an open half-space. Examples are the disk, the annulus, the torus with a disk removed, etc. A similar classification theorem holds for those: surfaces are still classified by their genus, their orientability, but also by the number of connected components of their boundary. It can be proved by gluing disks on the boundary to reduce to the usual case of surfaces. For $G$ a graph embedded on a surface $S$ (in particular for $G$ a simple closed curve), one can cut $S$ along $G$ and obtain a surface with boundary which we denote by $S \ll G$.

Maps: As in the planar case, the actual description of a cellularly embedded graph can be made purely combinatorial. This is already the case with the polygonal schemes, and can be made more convenient algorithmically by using flags, in exactly the same way as in the planar case.

Duality: A graph cellularly embedded on a surface naturally induces a dual graph, as in the planar case.

Subdivisions and minors: There is no analogue of Kuratowski's theorem for surfaces, but there is an analogue for Wagner's theorem. However, the proofs are non-constructive, and the precise list of minors is unknown except for the projective plane.

Separators: An $n$-vertex graph embedded on a surface of genus $g$ admits an $O(\sqrt{g n})$-sized balanced separator. A good reference for that is this paper of David Eppstein, Algorithms for planar graphs that are exclusively leveraging the existence of small separators can therefore directly be generalized to surface-embedded graphs.

Testing Embeddability: The genus of a graph is the smallest genus of a surface it embeds on. It is NP-hard to compute the genus of a graph, but for a fixed surface, there exists a linear-time algorithm to test embeddability on that surface. This will not be covered here (the original algorithm of Mohar and technical and intricate, a more recent one by Kawarabayashi, Mohar and Reed is simpler but requires strong familiarity with graph minor theory techniques).

### 3.2 Some topological algorithms

The main difference between the plane and more complicated surfaces is that there is no Jordan curve theorem outside of the plane, and thus that some simple closed curves can be nonseparating. Even among separating curves, there are differences, as is outlined in the following exercise:

Exercise 3.7. - Show that on an orientable surface of genus $g>0$, there are exactly $\lfloor g / 2\rfloor+2$ non-homeomorphid ${ }^{6}$ simple closed curves.

- (Harder) How many homeomorphism classes of curves are there on a non-orientable surface?

[^4]A standard criterion to differentiate closed curves on surfaces is homotopy. Two closed curves $\gamma_{1}$ and $\gamma_{2}$ are homotopic if there is a homotopy between them ${ }^{7}$, a continuous map $h:[0,1] \times \mathbb{S}^{1} \rightarrow S$ so that $h(0, \cdot)=\gamma$ and $h(1, \cdot)=\gamma_{2}$. The analogues of the nice curves in the planar case are the contractible curves: A simple closed curve $\gamma$ on a surface $S$ is contractible if it is homotopic to a trivial map at a point $p$. The following lemma, which could be taken as a definition, explains why:

Lemma 3.8. A simple closed curve is contractible if and only if it bounds a disk.

Half-proof: The reverse direction is immediate. For the forward direction, the subtlety comes from the fact the homotopy could introduce self-crossings. It turns out to be the case that such self-crossings are actually not needed, but we will not prove this.

Conversely, non-contractible curves are the main object of interest on surfaces. Note that non-contractible curves can be separating. More generally,

We will address two basic problems related to non-trivial curves:

1. How to test whether a closed curve is contractible?
2. How to compute the shortest non-contractible, or non-separating curve?

These algorithmic questions will also serve as an opening for the end of the course, as we use them to illustrate how classical tools from algebraic topology (universal covers, homology) impact the design of algorithms for topological problems on surface-embedded graphs.

### 3.3 Computing shortest interesting cycles

In this subsection, we provide algorithms to compute shortest non-contractible and non-separating curves. This is a natural and important topological primitive, and here are two applications for such algorithms:

- The first thing one often wants to do when given a surface is to cut it into something more planar. This is relevant for practical purposes (e.g., topological noise removal, texture mapping) and algorithm design, since then one can use the good old planar algorithms.
- In the planar case, as we saw, the min-cut problem reduces to computing a shortest cycle that "goes around" the annulus in the dual, i.e., the shortest non-contractible cycle. We are considering the direct generalization of this. This turns out to be relevant to solve min-cut on higher genus surfaces, but we will not have enough time to cover this in this course.

We will prove the following theorem:
Theorem 3.9. For a graph $G$ with $n$ edges embedded on an orientable surface of genus $g$, we can compute a shortest non-contractible cycle, respectively a shortest non-separating cycle, in time $O\left(n^{2} \log n\right)$.

[^5]

Figure 24: A graph and its dual. The walk $a b \bar{a} c$ is definitely not a cycle in the primal graph, but when we look at it in the cross-metric perspective and push it a bit, it is a nice simple closed curve (in red).

As we saw during the min-cut algorithm, when cutting surfaces, cutting along cycles in the primal or the dual graph can lead to annoying issues. For example when one cuts once along a cycle, and then along a second cycle sharing edges with the first one, this yields some degeneracies. In order to deal with these issues in a clean way, we formalize the approach that we used for min-cut, based on considering curves in general position with respect to a graph.

Recall that a curve is in general position with respect to an embedded graph if they cross transversely, away from the vertices, and a finite amount of times. A cross-metric surface $\left(S, G^{*}\right)$ is a topological surface $S$ with a (possibly edge-weighted) graph $G^{*}$ embedded on it. A cross-metric surface assigns lengths to any curve $\gamma$ in general position by simply counting the (possibly weighted) number of intersections with $G^{*}$. For a (possibly edge-weighted) graph $G$ cellularly embedded on a surface $S$, measuring the length of a walk in $G$ is the same as measuring its length in the cross-metric surface $\left(S, G^{*}\right)$, where $G^{*}$ is the dual (hence the notation). Similarly, shortest paths in the cross-metric surface ( $S, G^{*}$ ) can be computed by the usual graph algorithms on $G$. The added value with the cross-metric surface, compared to just using duality, is that we are considering more curves than in the pure graph-theoretical world: for example there are often walks on a graph that are not really self-crossing, in the sense that one could clearly push them infinitesimally to make them simple. In the cross-metric setting, we can directly pick them to be simple, which makes proofs more streamlined. See Figure 24.

Algorithmically speaking, while we will be considering arbitrary topological curves in transverse position with the graph $G^{*}$, we do not need to encode the precise location of a curve $\gamma$, merely its crossing points with $G^{*}$, its self-crossing points if there are any, and what it does in between. Equivalently, we can encode the superposition of $\gamma$ with $G^{*}$, which is also a cellularly embedded graph. Likewise, when there is more than one curve, we simply encode their superposition with the graph $G^{*}$.

Now that the setup is set up, we move forward. We first compute shortest loops: a loop $\ell$ is a closed curve $\ell: \mathbb{S}^{1} \rightarrow S$ going through a fixed point $b$, called the basepoint of the loop.

For a point $b$ in a face of $\left(S, G^{*}\right)$, we denote by $T$ a shortest path tree rooted at $b$, which can be computed by using Dijkstra's algorithm in the primal graph. The cut locus $C$ of ( $S, G^{*}$ ) with respect to $b$ is the set of edges not crossed by $T$. Informally, we are blowing a balloon based at $b$, and the cut locus is the set of points where it self-intersects. This should be very reminiscent of the tree-cotree duality that we explored on planar graphs, and thus the following lemma should not come as a surprise. See Figure 25 for an illustration.

Lemma 3.10. The cut locus cuts $S$ into a disk.

Proof. While growing the shortest path tree, the set of open faces visited by the tree union all the edges that it crosses is an open disk, and this is maintained until the end. At the end, the complement of this disk is exactly the set of of edges in the cut locus, which proves the lemma.


Figure 25: A shortest path tree, in black, defining a cut locus (in blue). Cutting the surface along the cut locus yields a disk. On the right, an edge of the cut locus corresponds to loop $\sigma(e)$ going through the basepoint $b$.

For an edge $e$ in the cut locus, we denote by $\sigma(e)$ the loop obtained by starting from $b$, taking a shortest path to one of the two faces of $G^{*}$ adjacent to $e$, crossing $e$, and coming back to $e$ via the shortest path on the other side of $e$. The weight of $e$ is defined to be the length of $\sigma(e)$. The following key lemma shows that the shortest non-contractible loop can be found among the $\sigma(e)$ :

Lemma 3.11. Some shortest non-contractible loop has the form $\sigma(e)$.
Proof. Let $L$ be a shortest non-contractible loop crossing the cut locus $C$ as few times as possible. If $L$ crosses $C$ at least twice, there is a point $p$ between the two crossings, cutting $L$ into $L_{1}$ and $L_{2}$. This point $p$ is connected to the root via a path $\rho$ on the shortest path tree $T$. Then either $L_{1}$ concatenated with $\rho$ or $L_{2}$ concatenated with $\rho$ is non-contractible, since otherwise the contraction of both would yield a contraction of $L$. This argument is sometimes called the 3-path condition, after Thomassen.

So $L$ crosses $C$ at most once. It has to cross it at least once, since otherwise it bounds a disk by the Jordan-Schoenflies theorem on the surface cut along the cut locus, and is thus contractible. So $L$ crosses the cut locus at some edge $e$, and since $L$ goes through the root, it has to have at least the length of $\sigma(e)$, which concludes the proof.

This immediately suggests a brute-force algorithm to find the shortest non-contractible loop: try all the $\sigma(e)$, test their contractibility (which is easy since they are simple) and output the shortest one.

But one can be smarter, and figure out a combinatorial criterion to decide whether a $\sigma(e)$ is contractible:

Lemma 3.12. Let e be an edge of $C$. Then $\sigma(e)$ is contractible if and only if some component of $C \backslash e$ is a tree.

Proof. For the reverse direction, if some component of $C \backslash e$ is a tree, one can homotope $\sigma(e)$ into a trivial loop by following this tree. For the forward direction, if $\sigma(e)$ is contractible, it bounds a disk by Lemma 3.8. If no component of $C \backslash e$ is a tree, then in particular the component that is a disk is not a tree, and thus it contains a cycle. But this contradicts Lemma 3.10.

We now have all the tools to prove the first half of Theorem 3.9:
Finding a shortest non-contractible cycle in $O\left(n^{2} \log n\right)$ time. For each face of $G$, we fix a root $r$ at $G$. Then we compute in $O(n \log n)$ time the cut locus based at $r$ as well as the weight
of each of its edges, i.e., the length of $\sigma(e)$. There remains to prune the cut locus, that is, to remove all of its useless arborescent parts. Since each tree has a degree-one vertex (its leaves), this can be done by removing all the degree one vertices, and then the new degree one vertices, etc. Our shortest non-contractible loop through $r$ is then the smallest of the remaining $\sigma(e)$ s. Looping through all the possible $r$ s increases the complexity to $O\left(n^{2} \log n\right)$, and we output the shortest of the resulting cycles.

## The following algorithm was not covered in class and is not part of the syllabus, I include it just for completeness.

In order to find the shortest non-separating loop, we first want to establish that one of them is of the form $\sigma(e)$. For non-contractible loops, this relied on the three-path condition, and we need something similar here, something like "Let $a$ and $b$ be two points on $S$ and $p, q$ and $r$ be three paths from $a$ to $b$, oriented from $a$ to $b$. If $p \bar{q}$ and $r \bar{q}$ are both separating, then so is $p \bar{p}$."

This poses an annoying issue, as it is not clear what it means for a curve that is not simple to be separating. So we first address this, and the convenient language for that is the language of $(\bmod 2)$ homology.

Let $G$ be a graph embedded on a surface $S$, with its set of vertices, edges and faces. We think of those as being 0 -dimensional, 1-dimensional and 2-dimensional objects. A $k$-chain, for $k=0,1$ or 2 is a subset respectively the set of vertices, edges or faces. We think of a chain as being an element in a vector space over $\mathbb{Z}_{2}$, the set of integers $\bmod 2$. Therefore, chains can be added, using the rule $1+1=0$. These vector spaces are denoted by $C_{0}, C_{1}$ and $C_{2}$. We define the boundary of an edge to be the sum of its endpoints, and the boundary of a face to be the sum of its boundary edges. These two boundary maps extend by linearity on the whole spaces $C_{1}$ and $C_{2}$, defining linear maps $\partial_{1}: C_{1} \rightarrow C_{0}$ and $\partial_{2}: C_{2} \rightarrow C_{1}$. A 1-chain is a cycle if its boundary is trivial, and it is a boundary if it is the boundary of some 2-chain. Cycles are generally denoted by $Z_{i}$, and boundaries by $B_{i}$. Convince yourself that the boundary of a boundary is empty.

Now, a closed walk $\gamma$ on $G$ or a closed curve in general position with $G$ can naturally be considered as a 1-chain (either for the graph $G$ or in the graph that is the overlay of $G$ and $\gamma$ ). The following lemma shows that being a homology boundary naturally generalizes being separating.

Lemma 3.13. A non-trivial simple closed curve $\gamma$ is a homology boundary if and only if it is separating.

Proof. If $\gamma$ is separating, then it is equal to the boundary of the sum of the faces of (either) one of the two connected components. If $\gamma$ is a homology boundary, then it is the boundary of a sum of faces $\mathcal{F}$. Note that this set of faces cannot be all the faces, as the boundary of the sum would be empty. Then the faces in $\mathcal{F}$ are separated from the faces not in $\mathcal{F}$, since any path connecting them would cross the boundary $\gamma$.

The homology group $H_{i}$ is defined as the quotient of the space $Z_{i}$ by the space $B_{i}$ : it is the space of cycles which are not boundaries. So it directly generalizes the separating curve.

With this language, we have the two tools needed to prove our missing lemma: a notion of sum of cycles, and a notion of separating for non-simple curve:

Lemma 3.14. Some shortest non-separating loop going through the root $r$ has the form $\sigma(e)$.
Proof. Since a non-separating loop is a non-trivial homology cycle, we can equivalently look for a shortest loop that is non-trivial in homology if we can prove (which we will) that one of them
is simple. As in the non-contractible case, any shortest non-trivial homology loop $L$ must cross $C$ at least once. We pick one that crosses $C$ a minimal number of times. If it crosses it more than once, we take $p$ to be a point between two crossings, which is connected to the root via a path $\rho$. The point $p$ cuts $L$ into $L_{1}$ and $L_{2}$, and we look at the three cycles $L, L_{1}$ concatenated with $\rho$ and $L_{2}$ concatenated with $\rho$. Note that the sum (as chains) of any two of these cycles forms the third cycle, and the set of homology boundaries is a vector space, and thus is closed under addition. So if both $L_{1}+\rho$ or $L_{2}+\rho$ were homology boundaries, then so would be $L$, which is a contradiction. So the shortest homology loop crosses $C$ exactly once, and since it contains $r$ it must be at least as long as $\sigma(e)$. Hence some shortest homology loop has the form $\sigma(e)$. Thus it is simple, and is thus non-separating.

Now we can bruteforce and try all the $\sigma(e)$ s to find a shortest non-separating one. Or we can try to be smarter, and prove that:

Lemma 3.15. A loop $\sigma(e)$ is separating if and only if e separates $C$.

Proof. If $\sigma(e)$ is separating, it separates $C$ into at least two components. In the other direction, if $e$ separates $C$, then any path on the surface connecting these two components and not crossing $\sigma(e)$ can be pushed back to $C$ since $S \backslash C$ is a disk.

The set of edges $e$ separating $C$ are called bridge edges. Note that edges corresponding to a contractible $\sigma(e)$, as characterized by Lemma 3.12 are bridge edges, which makes sense since contractible curves are separating. One can determine in linear time all the bridge edges of $C$ using depth-first search (exercise). Plugging everything together:

Finding a shortest non-separating cycle in $O\left(n^{2} \log n\right)$ time. For each face of $G$, we fix a root $r$ at $G$. Then we compute in $O(n \log n)$ time the cut locus based at $r$ as well as the weight of each of its edges, i.e., the length of $\sigma(e)$. We compute all the bridge edges and compare the remaining ones to find the shortest non-separating loop based at $r$. Looping through all the possible $r$ s increases the complexity to $O\left(n^{2} \log n\right)$, and we output the shortest of the resulting cycles.

Zooming out a bit: These two algorithms work because of the algebraic structure behind the curves: the concatenation of two contractible curves is contractible (i.e., the set of homotopy classes forms a group under concatenation), and the concatenation of two "separating" curves is "separating" (the homology classes form a group under addition). There are some other algebraic structures of interest, in particular, we can leverage relative homology to compute shortest systems of loops (see the lecture notes of Éric Colin de Verdière). Yet when there is no such known algebraic structure behind the problem that we consider, any optimization problem becomes much harder: how to compute the shortest polygonal scheme of the form $a_{1} b_{1} \overline{a_{1}} \overline{b_{1}} \ldots a_{g} b_{g} \overline{a_{g}} b_{g}$ for an orientable surface? How to compute the shortest collection of closed curves cutting a surface into a collection of sphere with three holes (a pants decomposition)? No polynomial algorithm nor hardness proof is known for these two problems.

### 3.4 Homotopy testing

To test whether a closed curve on a surface $S$ is contractible, if the curve is simple, we can simply cut along it and see whether one of the components is a disk. This works because of Lemma 3.8 , But if the curve is not simple, the problem is significantly less obvious to tackle. This is a good excuse to do some more topology. We will focus on getting polynomial algorithms. With quite some work, everything can actually be made linear.


Figure 26: The universal covers of the circle, the annulus and the torus.

The main strategy that we use is to use an auxiliary space from $S$, called the universal cover (or at least a portion of it), which differentiates in a natural way curves that are not homotopic. A covering space $(\hat{S}, p)$ for a surface $S$ is a topological space with a continuous map $p: \hat{S} \rightarrow S$ that is a local homeomorphism: any point $x$ in $S$ has an open neighborhood $U$ so that $p^{-1}(U)$ is a disjoint union of open sets $U_{i}$ which are all homeomorphic to $U$. The universal cover is a covering space where every simple closed curve is contractible - such a space is called simply connected. Such a universal cover is unique in some (natural) sense, though we will not prove this nor use it.

This abstract definition makes much more sense when looking at specific examples, see Figure 26 for some illustrations:

- The universal cover of a circle $\mathbb{S}^{1}$ is an infinite spiral over this circle.
- The universal cover of an annulus $\mathbb{A}$ is an infinite strip over this annulus.
- The universal cover of a sphere is the sphere itself.
- The universal cover of a torus is the plane $\mathbb{R}^{2}$ with a tiling into squares.

Exercise 3.16. What is the universal cover of the projective plane?
Since the map $p$ of a universal cover is a local homeomorphism, every path, or closed curve $\gamma$ on $S$, once a specific preimage of one of its points has been chosen, can be traced on $\hat{S}$, yielding a lift $\hat{\gamma}$. This is pictured in Figure 26 in the annulus example. But this lift might not be a closed curve. The following lemma shows that it is a closed curve exactly when it is contractible:

Lemma 3.17. A closed curve $\gamma$ on a surface lifts in the universal cover to a closed curve if and only if it is contractible.

Proof. The universal cover is simply connected, and thus each closed curve there can be contracted to a point. This contraction projects via $p$ to a contraction on $S$, and thus a closed curve on $S$ that lifts to a closed curve must be contractible. Conversely, the lift of a trivial curve is clearly a closed curve, and the contraction on $S$ lifts to a contraction in $\hat{S}$.


Figure 27: Lifting a closed curve on a disk with with two holes. Note that in the cover, the path starts on the first "floor" and ends on the third floor. Hence the curve is not contractible

Surfaces with boundary: For a surface with boundary, one can compute a portion of the universal cover in the following way. Starting from a graph $G$ embedded on a surface with boundary, we first contract and remove edges until there is a single face and all the vertices are on the boundary. Denote by $F$ the remaining edges. We cut along them, thus getting a disk $D$. The edges of $F$ will act as fences: whenever we cross such an edge, and nothing has been put yet on the other side, we glue a new copy of $D$. This process is naturally infinite, but is already useful for testing contractibility of a curve $\gamma$. First, observe that when an edge of $G$ gets removed, one can reroute $\gamma$ without changing its homotopy type, which might increase its length by an $O(g)$ factor. Then one can build the portion of the universal cover on which the lift of $\gamma$ will live, trace the lift of $\gamma$ on this portion and check whether the end point is the same as the starting point, see Figure 27. Thinking a bit more about this algorithm, it really only boils down to words: it amounts to writing the word formed by the edges of $F$ that we intersect, with a $\cdot^{-1}$ if the orientation is reversed, and checking whether the word we obtain can be reduced to a trivial word by reducing subwords of the kind $a a^{-1}$. We clearly get a polynomial algorithm.

Orientable surfaces without boundary: We first look at the easy example of the torus. In this case, the universal cover is $\mathbb{R}^{2}$ tiled by squares, and the boundaries of the square can be taken to be any pair of edges $(a, b)$ in a single vertex, single face graph embedded on the torus. One can test the contractibility of a curve by walking on this tiling of $\mathbb{R}^{2}$ and checking whether we come back to the same vertex at the end of the walk, see Figure 28 . We emphasize here the main difference compared to the case with a boundary: after following the path for some time, we will sometimes come back to the same tile in $\mathbb{R}^{2}$ despite having never backtracked. Such a case happens in the example. Here again, in practice, this is a problem on words: this amounts to counting how many times the walk crosses $a$ and $b$ (counting the crossings positively or negatively depending on how we cross it) and checking whether at the end of the walk, those algebraic counts are zero.

$\qquad$


Figure 28: Lifting a closed curve on a torus. In the cover, it lifts to a path, hence the curve is not contractible

For the more general case, as in the start of the classification of surfaces, up to removing and contracting edges, we can obtain from any graph a graph with a single vertex and a single face, which is often called a system of loops. By Euler's formula, this system of loops has $4 g$ loops. Cutting along it yields a $4 g$-gon. By analogy with the toroidal case, we want to look for the universal cover at a tiling of $\mathbb{R}^{2}$ where each tile is a $4 g$-gon, and every vertex has degree $4 g$. This is not possible in a Euclidean way, but there are such hyperbolic tilings, which are tilings in a space of negative curvature, see Figure 29 for an example with the genus 2 surface. (One model of) Hyperbolic geometry is a geometry where the ambient space is an open disk (thus homeomorphic to $\mathbb{R}^{2}$ ) and the straight lines are arcs of circle which are orthogonal to the boundary of the disk. The reader can check on the picture that we indeed get a tiling of the disk with octagons, eight of which meet at each vertex. Note that the definition of universal cover is purely topological, so the metric structure is just here at this stage to help us with intuition. One could try to trace on this hyperbolic tiling the walk and check whether we come back to a point, but now it is much less trivial to decide when we come back to the same tile. Instead, the following argument, dating back to Dehn, allows us to make progress locally until the curve is trivial:

A closed curve on a graph $G$ which is a system of loops can be seen as a word on $A \cup A^{-1}$, where the alphabet $A$ is the set of loops. A spur is a subword of the form $a a^{-1}$.

Lemma 3.18. Let $\gamma$ be a closed curve on a graph $G$ which is a cellularly embedded system of loops on an orientable surface of genus at least 2 . If $\gamma$ has no spurs and is contractible, then it has a subpath consisting of more than half of a facial walk of $G$.

Note that this is not true on a torus: indeed, this is a property that results from the hyperbolicity of the tiling.

The last lecture ended here (and I skipped the surfaces with boundary to gain some time). I include the proof of Dehn's lemma below using a discrete form of the Gauss-Bonnet theorem for completeness.

One can prove this using Euler's formula and a bit of sweat, but since we have introduced a bit of geometry, it is a good occasion to explore a geometric version of the Euler formula argument. It relies on a discrete notion of curvature, which quantifies how non-Euclidean a space is, which we first introduce. Let $D$ be a disk on which a graph is embedded. Around a vertex, between each pair of consecutive edges, there is a corner. For each corner $c$, let $\theta(c)$ be a positive number which we think of as the angle (in fractions of full turns) at this corner. In a Euclidean polygon, the sum of angles on a polygon with $d$ sides is always $(d-2) \pi$, so with our normalizations, this becomes $d / 2-1$. Therefore, it makes sense to define the curvature of


Figure 29: Covering a genus-two surface by a hyperbolic tiling of the open disk.
a face $\kappa(f)$ to be

$$
\kappa(f)=\sum_{c \in f} \theta(c)-\operatorname{deg}(f) / 2+1 .
$$

Thus, a triangle where the sum of angles is less than $\pi$ has negative curvature. Similarly, in usual space, the sum of angles around a vertex is $2 \pi$, or 1 when normalizing by full turns. Thus we define the curvature of an interior vertex $\kappa(v)$ is defined to be

$$
\kappa(v)=1-\sum_{c \in v} \theta(c) .
$$

In these discrete terms, a vertex with negative curvature is a vertex with too much stuff around it. Finally, the curvature of a boundary vertex $\tau(v)$ is

$$
\tau(v)=1 / 2-\sum_{c \in v} \theta(c)
$$

A magical formula called the Gauss-Bonnet formula stipulates that on a surface, the sum of the curvature equals the Euler characteristic (up to some constant normalizing factor). In our setting, the discrete version is the following:

Lemma 3.19. For a graph $G$ embedded on a disk $D$, and any choice of angles $\chi(c)$ on its corners, if we denote by $V_{i}$ and $V_{\partial}$ its interior and boundary vertices, we have:

$$
\sum_{v \in V_{i}} \kappa(v)+\sum_{v \in V_{\partial}} \tau(v)+\sum_{f \in F} \kappa(f)=\chi(D)=1 .
$$

Proof. Observe that in the three summands, the terms involving $\theta(c)$ cancel each other. What remains is $\left|V_{i}\right|+1 / 2\left|V_{\partial}\right|+|F|-\sum_{f} \operatorname{deg}(f) / 2$. The last sum counts the inner edges and half of
the edges on the boundary, or equivalently, all the edges minus half the edges on the boundary. Since on the boundary of the disk, there are as many edges as vertices, this sums to $\left|V_{i}\right|+$ $1 / 2\left|V_{\partial}\right|+|F|-\left(|E|-\left|V_{\partial}\right| / 2\right)=|V|-|E|+|F|$ which is exactly the Euler characteristic.

Of course, this generalizes to higher genus surfaces, with an identical proof. With this tool in hand, we proceed to the proof of Dehn's lemma.

Proof of Dehn's Lemma. If $\gamma$ is contractible, it lifts to a closed curve in the universal cover. This closed curve might self-intersect, and thus partitions our tiling of $\mathbb{R}^{2}$ into disks (and the outer face). Therefore it suffices to prove that any disk $D$ bounded by a closed curve on a $(4 g, 4 g)$-tiling of $\mathbb{R}^{2}$ contains a subpath of length at least $2 g+1$ on the boundary of one tile. The idea of the proof is to choose the angles so that a lot of curvature has to happen on the boundary vertices, which will force the subpath that we are looking for.

Therefore, we set all the corners to have angle $1 / 4$. Then all the faces and all the interior vertices all have negative curvature. The vertices of the boundary have positive curvature if and only if they are adjacent to a single face, in which case they have curvature $1 / 4$. We call such a vertex convex. So by the combinatorial Gauss-Bonnet formula, there will be a lot of convex vertices:

$$
\begin{aligned}
\sum_{v \in V_{i}} \kappa(v)+\sum_{v \in V_{\partial}} \tau(v)+\sum_{f \in F} \kappa(f) & =1 \\
|F|(1-g)+\left|V_{i}\right|(1-g)+\left|V_{\text {convex }}\right| / 4 & \geq 1 \\
\left|V_{\text {convex }}\right| & \geq(4 g-4)|F|+4
\end{aligned}
$$

So some face is adjacent to $4 g-3$ convex vertices. Since edges adjacent to these convex vertices are on the boudary of the disk, $4 g-3$ of these convex vertices must be consecutive. So some face has $4 g-2$ consecutive edges on $\partial D$, which is strictly bigger than $2 g$ for $g \geq 2$.

With Lemma 3.18 in hand, we can describe a combinatorial algorithm to test contractibility: after having reduced to a system of loops, we look at the word formed by the walk. We first inductively remove all the spurs. Then we scan it for a subword consisting of more than half of the facial walk of the system of loops. Whenever there is one, we replace it by the complementary part of the facial walk: this is a homotopy, and thus does not change the contractibility of the walk. Each of these changes reduces the complexity of the word, and we induct. By Lemma 3.18, if the closed curve is contractible, we will reach the trivial word. The complexity is clearly polynomial. With quite a lot of care, it can be made linear.

Zooming out a bit: In both variants, we have seen that the problem boils down to a problem on words: deciding whether a given word reduces to a trivial word under spur reduction (case with boundary), or with spur reductions and a more complicated relation. The underlying reason behind this is that the problem can be phrased in terms of combinatorial group theory. Indeed, the homotopy classes of loops on a surface, with the concatenation law, forms a group called the fundamental group of the surface, and a presentation for this group can be readily computed from a graph embedded on the surface: in the case with boundary, it is a free group with the fences as generators, and in the case without boundary, it is the single-relator group obtained with the system of loops as generators and the facial walk as a relation. Then the contractibility test amounts to testing the triviality of an element of this group. This perspective is easily misleading: in general testing triviality in a group defined by generators and relations
is undecidable. The fact that this can be solved for the fundamental group of surface is therefore remarkable.

More generally, we can try to test whether two given curves are homotopic. This is more subtle, but can also be made in linear time, using similar tools.


[^0]:    ${ }^{1}$ In 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.

[^1]:    ${ }^{2}$ It is not that obvious how to do so for continuous edges which might be pathological, but up to homeomorphism one can assume that they are polygonal, or even linear thanks to Theorem 1.26

[^2]:    ${ }^{3}$ There is no consensus on this definition, sometimes the outer one is also forced to be triangular, sometimes not
    ${ }^{4}$ It actually suffices for the graph to be 3-connected for this to work, we do the triangulated case for simplicity.

[^3]:    ${ }^{5}$ The same problem for undirected graphs is made quite a bit easier because there is always a cycle of length at most 5 .

[^4]:    ${ }^{6}$ Technical remark: any two simple closed curves are homeomorphic (since they are all homeomorphic to $\mathbb{S}^{1}$ ). Here, we mean "homeomorphism of the surface that sends one curve to the other one".

[^5]:    ${ }^{7}$ This definition requires curves to be oriented, we will often disregard this by considering two curves to be homotopic if some orientation makes them homotopic.

