Surfaces and embedded graphs

Francis Lazarus and Arnaud de Mesmay

October 5, 2017

Contents

1 Surfaces

1.1 Surfaces and cellularly embedded graphs

A **surface** is a Hausdorff and second countable topological space that is locally homeomorphic to the plane : that means that every point has a neighborhood homeomorphic to $\mathbb{R}^2$. Recall that a space is **Hausdorff** if every pair of distinct points have disjoint neighborhood and is **second countable** if it admits a countable base of open sets. In this course, we will only deal with compact surfaces, and will generally consider surfaces up to homeomorphism, which is why we say "the sphere" instead of "a sphere".

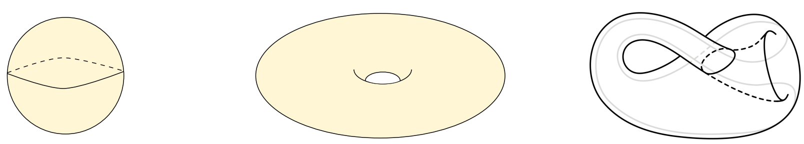

Figure 1: The sphere, the torus and the Klein bottle.

Examples of surfaces include the **sphere** $\mathbb{S}^2$, the **torus** $\mathbb{T}^2$, or the **Klein Bottle** $\mathbb{K}^2$, see Figure 1. Note that so far, we have not proved that they are different. We emphasize that surfaces are defined intrinsically, i.e., they do not have to be embedded in $\mathbb{R}^3$. For example, the Klein bottle cannot be embedded in $\mathbb{R}^3$: as in Figure 1, any representation of it in the usual space induces self-crossings. But this does not prevent it from being a surface : it is behaved locally like the plane which is all that matters here.

Exercise 1.1. Consider two copies of the sphere and identify all the corresponding points in the two copies, except for the North pole N. Formally, the resulting space is $\mathbb{S}^2 \times \{0,1\}/\sim$, where $(s,0) \sim (s,1)$ for all $s \in \mathbb{S}^2 \setminus \{N\}$. Show that this space is locally homeomorphic to the plane but that it is not Hausdorff.

Exercise 1.2. Show that the plane is second countable. Deduce that a compact space locally homeomorphic to the plane is second countable.

Following our approach outlined in the panorama, we will study surfaces by decomposing them into fundamental pieces, which can be seen as the faces of an embedded graph. Analogously to the planar case, an **embedding** of a graph G into a topological surface Σ is an image of G in Σ where the vertices correspond to distinct points and the edges correspond to simple arcs connecting the image of their endpoints, such that the interior of each arc avoids other vertices and arcs. We first remark that G can always be embedded in some surface. To see this, we can make a drawing of G in the plane and introduce a small handle at every edge intersection as on Figure 2 to obtain an embedding.

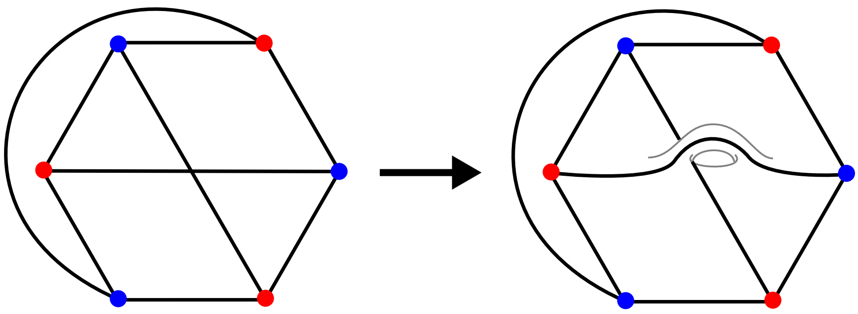

Figure 2: A plane drawing of $K_{3,3}$ with a crossing and an embedding in a genus 1 surface.

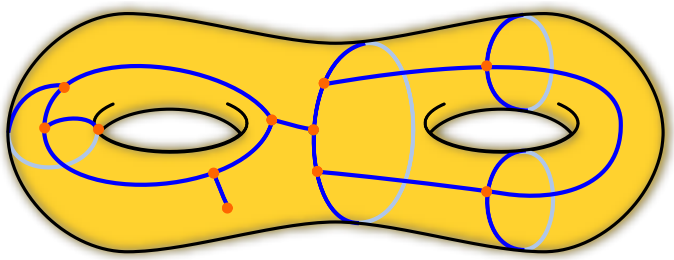

Figure 3: The complement of the graph in the surface is a disjoint union of open discs.

The **faces** of an embedding are the connected components of $S \setminus G$. A graph is **cellularly embedded** on S if it is embedded and all its faces are homeomorphic to a disk, see Figure 3. Thus, describing a cellularly embedded graph amounts to describing a combinatorial way to obtain a surface, by gluing disks together, and one can classify surfaces by studying the various possibilities. The following theorem shows that this approach loses no generality:

Theorem 1.3 (Kerékjártó-Radó)**.** *On any compact surface, there exists a cellularly embedded graph.*

Since disks can be triangulated, this is equivalent to saying that any compact surface can be triangulated, which is the way this theorem is generally stated in the literature.

PROOF. (Sketch) The result is obvious when the surface, call it S, is a sphere, so we assume this is not the case. Since S is compact and locally planar, it can be covered by a finite number of closed disks D_i, and up to the removal of the superfluous ones, we can assume that no disk lies in the union of any others. Then, if these disks intersect nicely enough (for example if two different boundaries ∂D and $\partial D'$ intersect in a finite number of points), one obtains a finite number of components in $S \setminus \cup \partial D_i$. One can easily show that each of these is a disk (because S is not a sphere!), and therefore one obtains a cellular graph by taking as vertices the intersection points, and as edges the arcs of circles.

So it suffices to show that one can assume that the disks intersect nicely. This can be done by repeated applications of the Jordan-Schoenflies Theorem, but it requires significant work. We refer to Thomassen [Tho92] or Doyle and Moran [DM68] for more details. □

Remark: The issue in this (non-)proof due to a possible infinite number of connected components might look like a mere technicality which one can *obviously* fix. However, we argue that there is a real difficulty lurking there, because the higher dimensional theorem is false: there exists a 4-manifold (A topological space locally homeomorphic to $\mathbb{R}^4$ that can not be triangulated[1], see for example Freedman [Fre82].

As in the planar case, a **triangulation** is a cellular embedding of a graph where all the faces have degree 3. A **subdivision** of a (triangular) face F is obtained by adding a vertex v inside the face, and adding edges between the new vertex v and all the vertices on F, or by adding a vertex w in the middle of an edge and adding edges between w and the non-adjacent vertices in the at most two incident faces. A triangulation is a **refinement** of another triangulation if it is obtained by repeated subdivisions. The same techniques can also be used to prove the following theorem:

Theorem 1.4 (Hauptvermutung in 2 dimensions)**.** *Any two triangulations on a given surface have a common refinement.*

We refer to Moise [Moi77] for a proof. As the name indicates ("main hypothesis" in German), this was widely believed to be true *in any dimension*, but once again counterexamples were found in dimensions 4 or higher (see for example [RCS+97]).

[1]On a first approximation, it means that it can not be built from a finite number of balls. More formally, it can not be realized as a simplicial complex, which we will introduce later on in the course.

1.2 Polygonal schemata

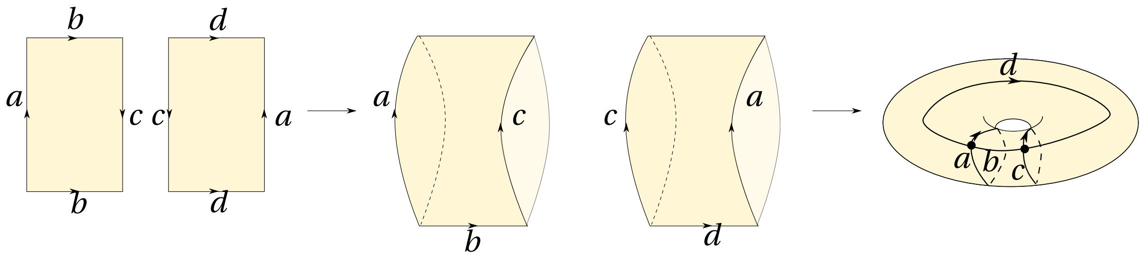

Figure 4: From the polygonal scheme $\{a\,b\,c\,\bar{b}, \bar{c}\,d\,\bar{a}\,\bar{d}\}$ to a cellular embedding.

In order to classify surfaces, we introduce **polygonal schemata**, which are a way of encoding the combinatorial data of a cellularly embedded graph : it describes a finite number of polygons with oriented sides identified in pairs. We will see later on, in Section 2, other avatars of this combinatorial description of a cellularly embedded graph.

Formally, let S be a finite set of **symbols**, and denote by $\bar{S} = \{\bar{s} \mid s \in S\}$. Then a polygonal scheme is a finite set R of **relations**, each relation being a non-empty word in the alphabet $S \cup \bar{S}$, so that for every $s \in S$, the total number of occurences of s or $\bar{s}$ in R is exactly two.

Starting from a cellularly embedded graph it induces a polygonal scheme in the following way: we first name the edges and orient them arbitrarily. Then for every face, we follow the cyclic list of edges around that face, with a bar if and only if an edge appears in the wrong direction. Every face gives us a relation of R and since every edge is adjacent to exactly two faces, possibly the same, we obtain a polygonal scheme. Conversely, starting from a polygonal scheme, for each relation of size n we build a polygon with n sides, and label its sides following the relation (with the bar indicating the orientation). Then, once all the polygons are built, we can identify the edges labelled with the same label taking the orientations into account. See Figure 4.

Exercise 1.5. The topological space obtained this way is a compact surface.

Thus, polygonal schemes and cellularly embedded graphs are two facets of the same object. Furthermore, by Theorem 1.3, every surface has a cellularly embedded graph, and thus can be obtained by some polygonal scheme. We leverage on this to classify surfaces.

1.3 Classification of surfaces

Theorem 1.6. *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. $a_1 b_1 \bar{a}_1 \bar{b}_1 \dots a_g b_g \bar{a}_g b_g$ *for some* $g \geq 1$,
3. $a_1 a_1 \dots a_g a_g$ *for some* $g \geq 1$.

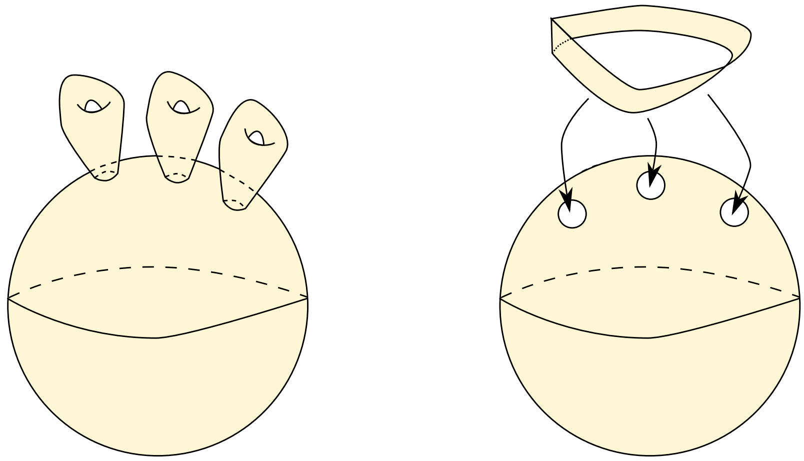

Figure 5: The orientable surface of genus 3 and the non-orientable surface of genus 3.

In the second case, the surface is said to be **orientable**, while in the third case it is **non-orientable**. The integer g is called the **genus** of the surface (by convention $g = 0$ for the sphere). In the orientable case, the genus quantifies the number of holes of a surface : an orientable surface of genus g can be built by adding g handles to a sphere. A non-orientable surface of genus g can be built by cutting out g disks of a sphere and gluing g **Möbius bands** along their boundaries. See Figure 5.

PROOF. We follow the exposition of Stillwell [Sti93]. Let S be a compact connected surface, and let G be a graph cellularly embedded on S, which exists by Theorem 1.3. Whenever an edge of G is adjacent to two different faces, we remove it. Whenever an edge of G is adjacent to two different vertices, we contract[2] it. When this is done, we obtain a cellularly embedded graph G' with a single face and a single vertex. If there are no more edges, then by uncontracting the single vertex into two vertices linked by an edge, we are in case 1 of the theorem and the surface is a sphere. Therefore we can now assume that there is at least one edge.

The graph G' induces a polygonal scheme consisting of a single relation. We will show that this relation can be transformed into either case 2 or case 3 of the theorem without changing the homeomorphism class of S.

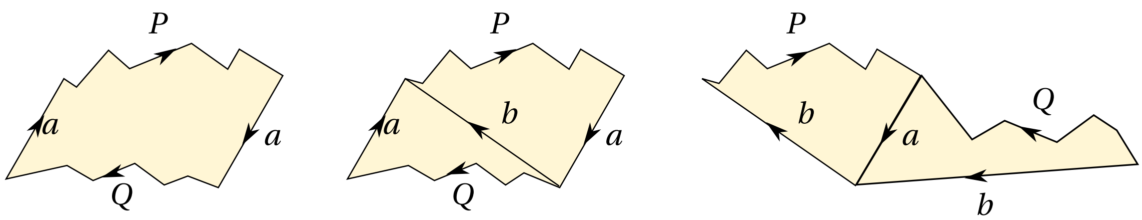

Figure 6: From $aPaQ$ to $bbP\bar{Q}$.

- If the polygonal scheme has the form $aPaQ$ where P and Q are possibly empty words, then we can transform it into $bbP\bar{Q}$ by adding a new edge and removing

[2]The contraction is not meant in the graph-theoretical sense introduced in the earlier chapter : it might result in loops and multiples edges, which we keep.

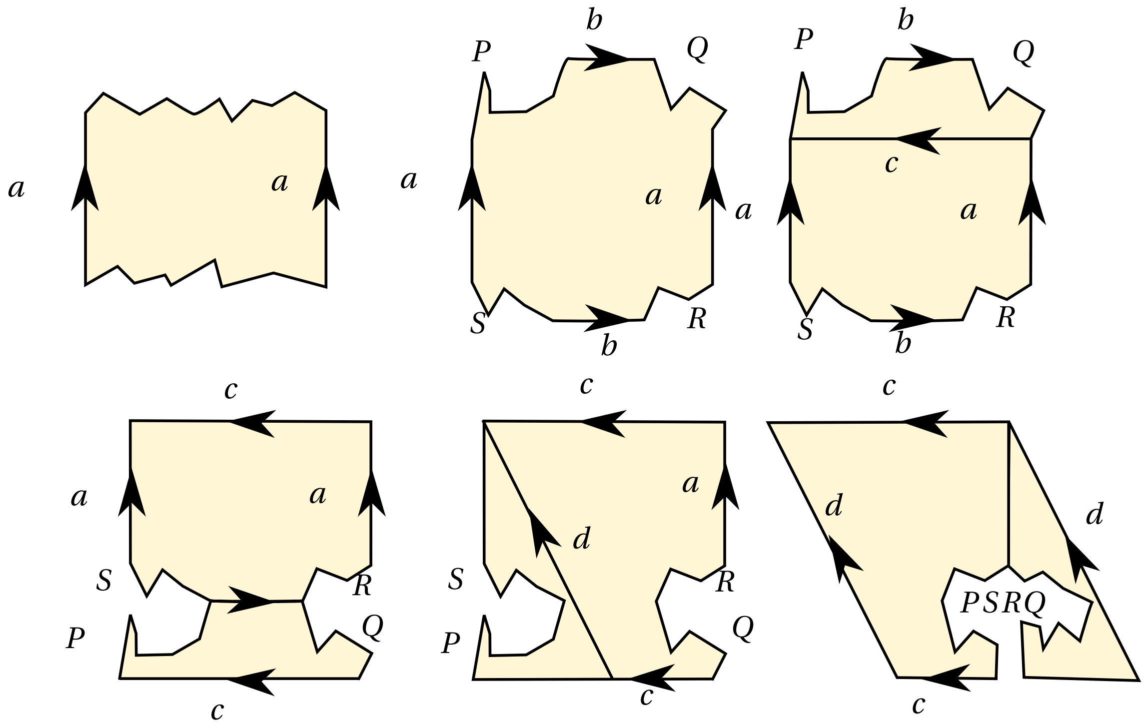

Figure 7: From $aPbQ\bar{a}R\bar{b}S$ to $cd\bar{c}\bar{d}PSRQ$.

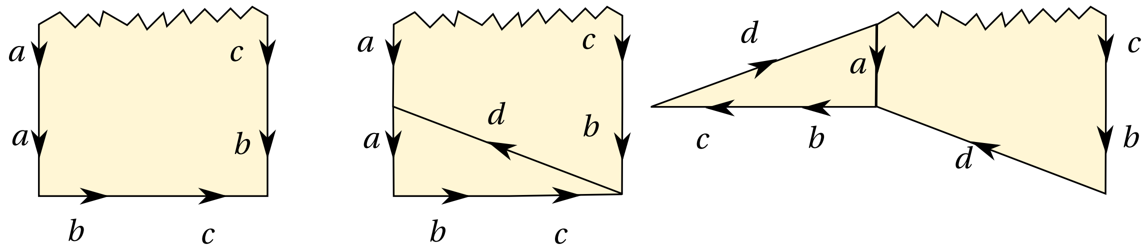

Figure 8: From $aabc\bar{b}\bar{c}$ to $\bar{d}\bar{c}\bar{b}\bar{d}\bar{b}\bar{c}$.

a, see Figure 6. 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 $aU\bar{a}V$, then U and V must share an edge b since otherwise G' 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 $aU\bar{a}V = aPbQ\bar{a}R\bar{b}S$. This can be transformed into $dc\bar{d}\bar{c}PSRQ$, as pictured in Figure 7. Inductively, at the end of this step the relation is a concatenation of blocks of the form aa or $ab\bar{a}\bar{b}$. 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 $aabc\bar{b}\bar{c}$. This can be transformed into $\bar{d}\bar{c}\bar{b}\bar{d}\bar{b}\bar{c}$, 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. □

Let G be a graph cellularly embedded on a compact surface. The **Euler characteristic** of this embedding equals $v-e+f$, where v is the number of vertices, e is the number of edges and f is the number of faces of the embedded graph.

Lemma 1.7. *The Euler characteristic of a graph G cellularly embedded on a surface S only depends on the surface S.*

PROOF. Let G and G' be two cellular embeddings on the same surface S. Since triangulating faces does not change the Euler characteristic, one can suppose that they are triangulated. By Theorem 1.4, they have a common refinement. Since subdividing faces does not change the Euler characteristic, this proves the Lemma. □

The Euler characteristic of the surfaces in Theorem 1.6 are readily computed from their polygonal schemes: for the sphere we obtain two, for the orientable surfaces $2-2g$ and for the non-orientables ones $2-g$. Therefore the orientable surfaces are all pairwise non-homeomorphic, as are the non-orientable ones. Can orientable surfaces be homeomorphic to non-orientable ones?

Lemma 1.8. *A surface S is orientable if and only if it has a cellularly embedded graph G such that the boundary of its faces can be oriented so that each edge gets two opposite orientations by its incident faces.*

PROOF. If the surface S is orientable, then it can be obtained by a polygonal scheme of type 2, for which the boundaries of the faces can be oriented as the lemma requires. If the surface is non-orientable, then any cellularly embedded graph G has a common refinement with one having a polygonal scheme of type 3. Observing that such a graph can not be oriented as the lemma requires, and that this property is maintained when refining, this proves the lemma. □

Corollary 1.9. *Orientable surfaces are not homeomorphic to non-orientable ones.*

Therefore, we have established that all the surfaces in Theorem 1.6 are pairwise non-homeomorphic. Conversely, any pair of connected surfaces with the same orientability (as defined by Lemma 1.8) and Euler characteristic are homeomorphic.

Remark: This classification of surfaces can be extended to the setting of surfaces with boundary: a **surface with boundary** is a topological space where every point is locally homeomorphic to either the plane or the closed half-plane. The **boundary** of such a surface is the set of points that have no neighborhood homeomorphic to the plane. One can show that up to homoemorphism, in line with the above classification, surfaces with boundaries are classified by their genus, their orientability and the number of boundaries (i.e., connected components of the boundary). One way to obtain this is on the one hand to observe that the number of boundaries is a topological invariant, and on the other hand that by gluing disks on the boundaries of a surface with boundary, one obtains a surface without boundary, for which the usual classification applies. The Euler characteristic of the orientable, respectively non-orientable surface of genus g with b boundaries is $2-2g-b$, respectively $2-g-b$.

2 Maps

To make things simpler we shall restrict ourselves from now on to orientable surfaces. Up to homeomorphism, a cellular embedding of a graph can be described by the graph itself together with the circular ordering of the edges incident to each vertex. These are purely combinatorial data referred to as a **rotation system**, a **cellular embedding** (of a graph), a **combinatorial surface**, a **combinatorial map**, or just a **map**. The theory of combinatorial maps was developed from the early 1970's, but can be traced back to works of Heffter [Hef91, Hef98] and Edmonds [Edm60] for the combinatorial description of a graph embedded on a surface. The notion of combinatorial map relies on oriented edges rather than just edges. An oriented edge is also called an **arc** or a **half-edge**. Formally, a combinatorial map is a triple (A,ρ,ι) where

- A is a set of arcs,
- $\rho : A \to A$ is a permutation of A,
- $\iota : A \to A$ is a fixed point free involution.

This data allows to recover the embedded graph easily: its vertices correspond to the orbits, or cycles (of the cyclic decomposition), of ρ and its edges correspond to the orbits of ι (so that a and $\iota(a)$ correspond to the two orientations of a same edge). The source vertex of an arc is its ρ-orbit. We shall often write $\bar{a}$ for $\iota(a)$. There are two basic ways of visualizing the corresponding cellular embedding. One way consists in placing disjoint disks in the xy-plane of $\mathbb{R}^3$, one for each vertex, then attaching rectangular strips to the disks, with one strip per edge. The strips should expand in $\mathbb{R}^3$ so that they do not intersect. The counterclockwise ordering of the strips attached to a disc should coincide with the cycle of ρ defining the corresponding vertex. See Figure 9 for an illustration. The resulting *ribbon graph* is topologically equivalent to a surface

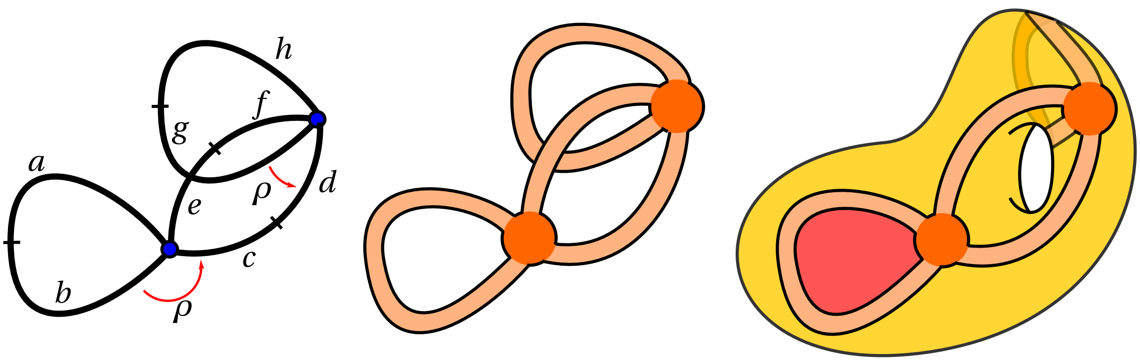

Figure 9: A cellular embedding associated to the map (A,ρ,ι) with $A = \{a,b,c,d,e,f,g,h\}$, $\rho = (a,b,c,e)(g,d,h,f)$ and $\iota = (a,b)(c,d)(e,f)(g,h)$. The corresponding graph has a loop edge and a multiple edge.

with boundary. Finally glue a disk along each boundary component to a obtain a closed surface where the graph is cellularly embedded. Note that the boundary of each face, traversed with the face to the right, visits the arcs according to the permutation $\varphi := \rho \circ \iota$. The φ-orbits are called **facial walks**. A facial walk need not be simple as can be seen on Figure 3. Note that this construction is dual to the concept of polygonal scheme that we saw earlier : another way of visualizing the cellular embedding is to draw one polygon per facial walk, marking its sides with the arcs of the orbit. Then glue the sides of the polygons that correspond to oppositely oriented arcs (related by the involution ι). Figure 10 illustrates this second construction. The numbers $|V|, |E|, |F|$

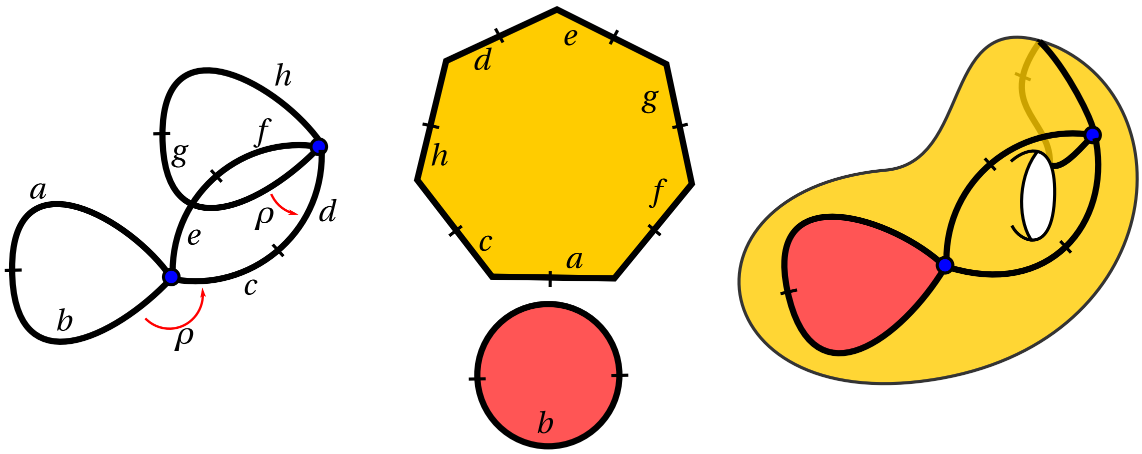

Figure 10: Left, the same map as above. Middle, the facial walks of $\varphi = (a,c,h,d,e,g,f)(b)$. Right, the resulting graph embedding.

of vertices, edges and faces of the resulting surface are thus given by the number of cycles of the permutations ρ, ι and φ respectively. Obviously, the number of cycles of the involution ι is just $|E| = |A|/2$. The Euler characteristic of this surface can then be computed by the formula

$$\chi = |V| - |E| + |F|.$$

Basic operations on maps. The contraction or deletion of an edge in a graph extend naturally to embedded graphs. Given a map $M = (A, \rho, \iota)$ with graph G, the **contraction** of a non-loop edge $e = \{a, \bar{a}\}$ in G leads to a new map M/e obtained by merging the circular orderings at the two endpoints of e. See Figure 11. Formally,

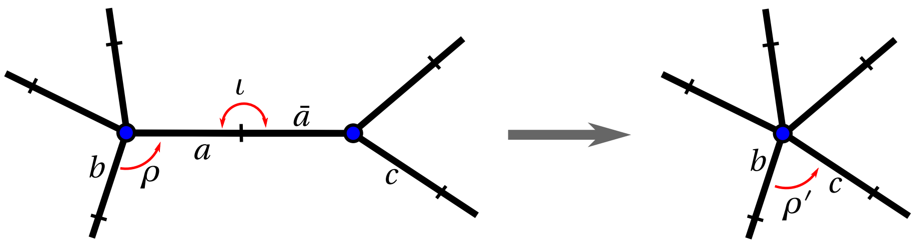

Figure 11: The contraction of a non-loop edge. $\rho(b) = a \implies \rho'(b) = \rho \circ \iota(\rho(b)) = c$.

$M/e = (A \setminus e, \rho', \iota')$ where ι' is the restriction of ι to $A \setminus e$ and ρ' is obtained by merging the cycles of a and $\bar{a}$, i.e.,

$$\forall b \in A \setminus e, \rho'(b) = \begin{cases} \rho(b) & \text{if } \rho(b) \notin e, \\ \rho \circ \iota(\rho(b)) & \text{if } \rho(b) \in e \text{ and } \rho \circ \iota(\rho(b)) \notin e, \\ (\rho \circ \iota)^2(\rho(b)) & \text{otherwise.} \end{cases}$$

Likewise, if e has no degree one vertex, the **deletion** of e in G leads to new map $M-e=(A\setminus e,\rho',\iota')$ where ι' is the restriction of ι to $A\setminus e$ and ρ' is obtained by deleting a and $\bar{a}$ in the cycles of ρ, *i.e.*,

$$\forall b\in A\setminus e, \rho'(b)=\begin{cases}\rho(b) & \text{if } \rho(b)\notin e,\\ \rho^2(b) & \text{if } \rho(b)\in e \text{ and } \rho^2(b)\notin e,\\ \rho^3(b) & \text{otherwise.}\end{cases} \quad (1)$$

Figure 12 illustrates the deletion of a loop edge. Let us look at the effect of an edge

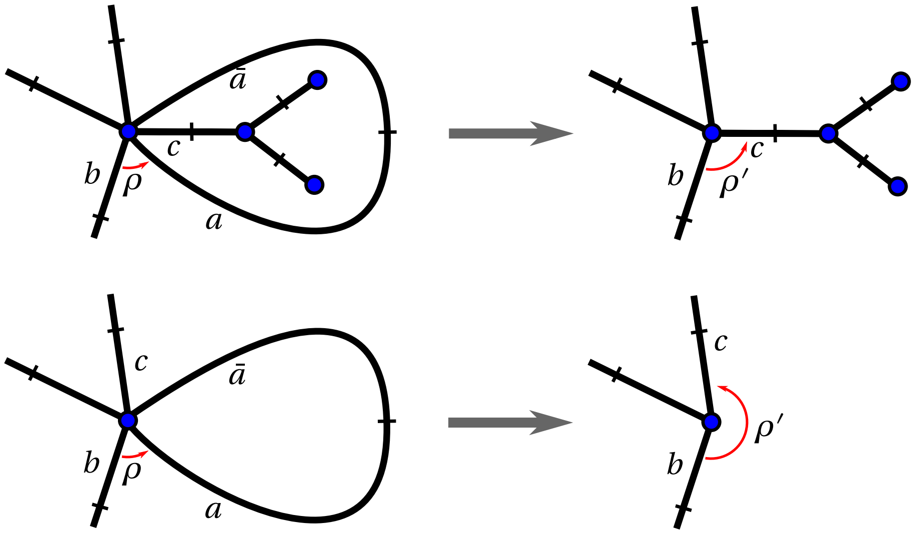

Figure 12: The deletion of a loop edge. Above, We have $\rho^2(b)\notin\{a,\bar{a}\}$ implying $\rho'(b)=\rho^2(b)=c$. Below, $\rho^2(b)\in\{a,\bar{a}\}$ so that $\rho'(b)=\rho^3(b)=c$.

contraction or deletion on the topology of a cellular embedding.

Lemma 2.1. *If M is a connected map with at least two edges and $e=\{a,\bar{a}\}$ is a non-loop edge of M then M/e is connected and has the same Euler characteristic as M.*

PROOF. The lemma is quite clear if one remarks that M/e has the same number of faces as M but has one edge less and one vertex less than M. Hence,

$$\chi(M\setminus e)=(|V(M)|-1)-(|E(M)|-1)+|F(M)|=\chi(M)$$

□

An edge of an embedded graph is said **regular** if is it incident to two distinct faces and **singular** otherwise.

Lemma 2.2. *Let e be an edge of a map M with at least two edges. If e has no vertex of degree one, then*

$$\chi(M-e)=\begin{cases}\chi(M) & \textit{if } e \textit{ is regular}\\ \chi(M)+2 & \textit{otherwise.}\end{cases}$$

Note that the deletion of e may disconnect the map.

PROOF. Clearly, M' has the same number of vertices has M and one edge less. Let $\varphi = \rho \circ \iota$ and $\varphi' = \rho' \circ \iota'$ be the facial permutation of M and M' respectively. Writing $e = \{a, \bar{a}\}$, we have that e is regular if and only if the φ-cycles of a and $\bar{a}$ are distinct. Using formula (1), we see that the cycles of φ' are the same as for φ except for those containing a and $\bar{a}$, which are merged if they are distinct for φ and which is split otherwise. We infer that M' has one face less in the former case and one more in the latter. We conclude that

$$\chi(M-e) = |V(M)| - (|E(M)|-1) + (|F(M)|-1) = \chi(M)$$

if e is regular and

$$\chi(M-e) = |V(M)| - (|E(M)|-1) + (|F(M)|+1) = \chi(M)+2$$

otherwise. □

We also define an **edge subdivision** in a map by introducing a vertex in the middle of one of its edges. Likewise, a **face subdivision** consists in the splitting of a face by the insertion of an edge between two vertices of its facial walk. Remark that by contracting one of the two new edges in an edge subdivision one recovers the original map. Similarly, the new edge in a face subdivision is regular and its deletion leads to the original map. It follows from Lemmas 2.1 and 2.2 that any subdivision of a map preserves the characteristic. Define the **genus of a graph** as the minimum genus of any orientable surface where the graph embeds.

Corollary 2.3. *The genus of (a subdivision of) a minor of a graph G is at most the genus of G.*

PROOF. Let M be a cellular embedding of G with minimal genus g. Any subdivision H of a minor of G can be obtained by a succession of edge contractions, deletions and subdivisions. We can perform the same operations on M. By Lemmas 2.1 and 2.2 and the preceding discussion, the characteristic may only increase during these operations. It follows that the resulting embedding of H has genus at most g, implying that the genus of H is at most g. □

3 The Genus of a Map

Thanks to Euler's formula it is quite easy to recover the genus of a map given by a triple (A, ρ, ι). We have:

$$g = 1 - \chi/2 = 1 - (|V| - |E| + |F|)/2 \tag{2}$$

where V, E, F are the set of vertices, edges and faces of the map. For a graph G, a certificate that it can be embedded in a surface of genus at most g may be given in the form of a rotation system for G, checking that the genus of the resulting map is at most g. In particular, a graph is planar if and only if it admits a rotation system of genus zero. It follows from the above certificate that computing the genus of a graph is an NP problem. A greedy approach to compute the genus of G is to compute the minimum genus of every possible rotation system for G. For a vertex v the number of possible circular orderings of the incident arcs is $(d_v - 1)!$ where d_v is the degree of v in G. It ensues that the greedy approach needs to consider as much as $\prod_{v \in V(G)}(d_v - 1)!$ rotation systems. It appears that the problem is hard to solve. Indeed,

Theorem 3.1 (Thomassen, 1993)**.** *The graph genus problem is NP-complete.*

We first remark that we can restrict the problem to connected simple graphs with at least three vertices. Moreover, given such a graph G, the existence of a rotation system on G that triangulates a surface, *i.e.* such that every facial walk has length three, reduces to the genus problem. Indeed, the number of vertices and edges being fixed by G, only the number of faces may vary among rotation systems. Since the faces of a map correspond to its φ-cycles and since every face has length at least 3, we have $3|F| \leq |A| = |E|/2$ with equality if and only if the map is a triangulation. Formula (2) shows that g is minimal in this case. In other words, we can directly deduce from its genus whether G triangulates a surface or not. It is thus enough to show the NP hardness of the triangulation problem. The proof relies on a reduction of the following problem to the triangulation problem.

Proposition 3.2 ([Tho93])**.** *Deciding whether a cubic bipartite graph contains two Hamiltonian cycles intersecting in a perfect matching is an NP complete problem.*

Recall that a graph is **cubic** if all its vertices have degree three and it is **bipartite** if its vertices can be split into two sets such that no edge joins two vertices in a same set. A cycle in the graph is **Hamiltonian** if it goes through all the vertices. Finally, a **perfect matching** is a subset of edges such that every vertex is incident to exactly one edge in the subset.

PROOF OF THEOREM 3.1. We shall reduce the problem of Proposition 3.2 to the triangulation problem. By Proposition 3.2 and the discussion after the theorem, the claim implies that the genus problem is NP hard, hence NP complete as we already know it is in NP. Let G be a cubic bipartite graph with a bipartition $A \cup B$ of its vertex set. We construct another graph by first taking a copy G' of G, adding one edge between each vertex v of G and its copy v' in G'. We further add four vertices v_1, v_2, v_1', v_2' and join v_1 and v_2 to every vertex in G and similarly join v_1' and v_2' to every vertex in G'. Let H be the resulting graph. We next construct a graph Q by contracting all the edges of H of the form vv' with $v \in A$ and v' its copy in G'. We claim that
Q triangulates a surface if and only if G admits two Hamiltonian cycles intersecting in a perfect matching.
We first prove the direct implication in the claim, assuming that Q triangulates a surface. In other words there is map with graph Q all of whose faces are triangles. The local rotation of this map around v_1 directly provides a Hamiltonian cycle C_1 in G, which is the boundary of the union of the triangles incident to v_1. Note that every vertex of C_1 is incident to three edges in this union: two edges along C_1 and one edge toward v_1. Similarly, the local rotation around v_2 provides a Hamiltonian cycle C_2. If C_1 and C_2 had two consecutive edges in common, then their shared endpoint would

be incident to exactly two more edges: one toward v_1 and one toward v_2. However, by construction v is also incident to its copy in G' or to v_1' and v_2' if $v \in A$, leading to a contradiction. Moreover, since G is cubic C_2 cannot miss two consecutive edges of C_1. It follows that C_1 and C_2 share one of every two edges, hence a perfect matching.

To prove the reverse implication of the claim, suppose now that G has two Hamiltonian cycles C_1 and C_2 intersecting in a perfect matching. We consider the following rotation system on H. If $v \in A$, we let e_i, for $\{i, j\} = \{1,2\}$, be the arc of $C_i \setminus C_j$ with source vertex v and we let e_3 be the third incident arc, hence in $C_1 \cap C_2$. The local rotation around v is then given by the cycle

$$(e_1, v v_1, e_3, v v_2, e_2, v v')$$

The local rotation around a vertex in $V(G')\setminus A'$ is defined analogously and so is the local rotation around a vertex in A' or in $V(G)\setminus A$, except that we exchange the indices 1 and 2 in the above cycle. See Figure 13. For $\{i, j\} = \{1,2\}$, we define the local rotation

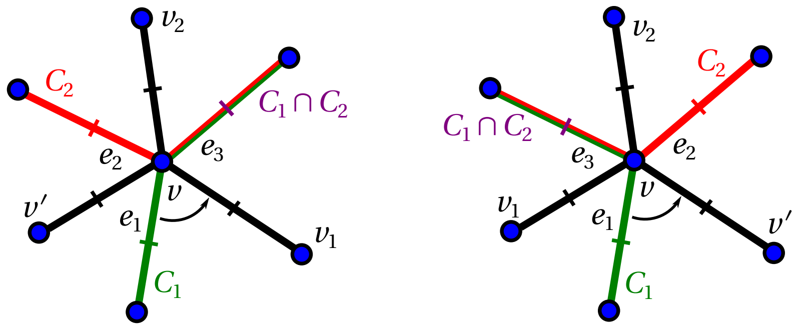

Figure 13: Left, the local rotation at a vertex $v \in A \cup V(G')\setminus A'$. Right, the local rotation at a vertex $v \in A' \cup V(G)\setminus A$. The vertices v, v' are copies of the same vertex in G and G'.

around v_i as the cycle C_i oriented so that the edges in common with C_j (previously denoted by e_3) are directed from $V(G)\setminus A$ to A for $i = 1$ and from A to $V(G)\setminus A$ for $i = 2$. Similarly, we define the local rotation around v_i' as the cycle C_i' oriented oppositely to C_i, *i.e.* so that the edges in common with C_j' are oriented from A to $V(G)\setminus A$ when $i = 1$ and from $V(G')\setminus A'$ to A' when $i = 2$. It is an exercise to check that the facial walks of the resulting rotation system have the following form, where $\{i, j\} = \{1,2\}$:

- A triangle defined by v_i and some edge of C_i, or
- A triangle defined by v_i' and some edge of C_i', or
- a quadrilateral defined by an edge of $C_i \setminus C_j$, its copy in G' and the two edges between their endpoints from G to G'.

Figure 14 shows some of those facial walks. Hence, by contracting the edges $v v'$ of H with $v \in A$ (and its copy $v' \in A'$), the quadrilaterals are transformed into triangles and one obtain a triangular embedding for Q. □

It can be shown that the graph genus problem is fixed parameter tractable (FPT) with respect to the genus. More precisely, the question whether a graph of size n

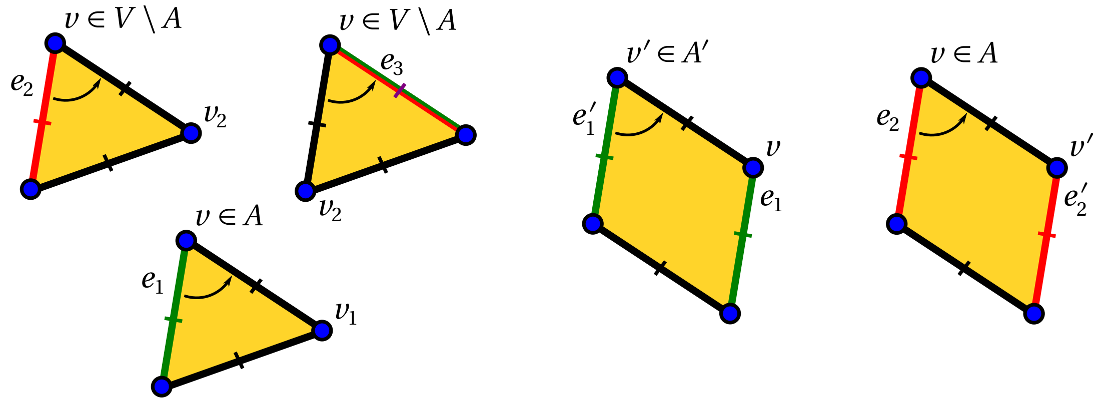

Figure 14: Some faces of the rotation system for H.

has genus at most g can be answered in $O(f(g)n)$ time where $f(g)$ is some singly exponential function of g [Moh99, KMR08]. Interestingly, the genus of the complete graphs are known. It was conjectured by Heawood in 1890 that the genus of the complete graph K_n over $n \geq 3$ vertices is

$$g(K_n) = \left\lceil \frac{(n-3)(n-4)}{12} \right\rceil$$

This conjecture was eventually established in 1968 by Ringel and Youngs. The long proof [Rin74] provides explicit minimal genus embeddings of K_n with a different construction for each residue of n modulo 12.

4 Homotopy

In a nutshell, two curves drawn on a surface are homotopic if there exists a continuous deformation between them. This intuitive notion, dating back to Poincaré, naturally leads to a very rich theory drawing a bridge between topology on one side and group theory on the other side. We start by introducing the relevant background in group theory.

4.1 Groups, generators and relations

Although most groups we will be dealing with in this course are infinite, they can often be very succinctly encoded in terms of **generators** and **relations**. The attentive reader will probably notice some similarities between the formalism established here and the definitions on polygonal schemata : as we will see later on, this is no coincidence.

Let us consider a set G of **generators**, and denote by G^{-1} their **inverses** $G^{-1} = \{g^{-1} \mid g \in G\}$. A word is a string over the alphabet $G \cup G^{-1}$, and we denote by ε the empty word. We consider that two words are equivalent if one can switch from one to the other by adding or removing words of the form gg^{-1} or $g^{-1}g$. The set of finite words quotiented by this equivalence relation can naturally be endowed with the structure of a group: the law is the concatenation, and the neutral element is ε. Indeed:

- The concatenation is well-defined with respect to the equivalence relation and is associative.
- For any word w, $w\varepsilon = \varepsilon w = w$.
- Each element has an inverse obtained by reversing the order of the letters and inverting them, e.g., $c^{-1}b^{-1}a^{-1}$ is the inverse of abc.

This group is called the **free group** on the set G, denoted by $F(G)$.

Now, let us also fix a set R of words called the **relations**, and consider the free group $F(G)$, where one also identifies a word with the word obtained by inserting at any place a word taken from R or their inverses. This defines another group, which is formally the quotient of $F(G)$ by the normal subgroup generated by R[3]. This group is said to admit the **presentation** $< G \mid R >$. So the free group admits the presentation $< G \mid \emptyset >$, generally abbreviated by $< G >$.

Here are a few examples:

- The group $\mathbb{Z}$ is the free group on one letter $F(\{a\})$.
- The group $< a \mid aa >$ is the group $\mathbb{Z}_2$ (or $\mathbb{Z}/2\mathbb{Z}$).
- The group $< a, b \mid aba^{-1}b^{-1} >$ is the two-dimensional lattice $\mathbb{Z}^2$: indeed, the relation $aba^{-1}b^{-1}$ implies that $ab = ba$, thus the group is abelian, and the isomorphism with $\mathbb{Z}^2$ is the map $a \mapsto (1,0)$, $b \mapsto (0,1)$.

Exercise 4.1. Recognize the groups $< a, b \mid aab, ab^{-1}b^{-1} >$ and $< a, b \mid a^m, b^n, aba^{-1}b^{-1} >$.

Exercise 4.2. Show that any group admits a presentation (with possibly an infinite number of generators and relations), and that any finite group admits a finite presentation.

4.2 Fundamental groups, the combinatorial way

We start by introducing homotopy in a combinatorial setting, which makes computations very convenient. The baby case is the case of graphs, which corresponds directly to free groups.

4.2.1 Fundamental groups of graphs

Let G denote a graph (not necessarily embedded) where the edges are oriented. An **arc** is an oriented edge or its inverse, it has an **origin** $o(a)$ and a **target** $t(a) = o(a^{-1})$. A **path** in G is a sequence of arcs $(e_1, \ldots, e_n)$ such that the target of e_i coincides with the origin of e_{i+1}. A **loop** is a path such that the target of e_n coincides with the origin of e_1, this point is called the **basepoint** of the loop. The **trivial loop** is the empty loop. Two loops with a common basepoint x are **homotopic** if they can be related to each other by adding or removing subpaths of the type (e, e^{-1}), and a path is **reduced** if it does not contain such a subpath.

[3] Recall that a subgroup $H \subseteq G$ is normal if $gH = Hg$ for any $g \in G$.

Let x be a vertex of G. The set of homotopy classes of loops in G forms a group, where the law is the concatenation and the neutral element the trivial loop. Indeed:

- The concatenation is well-defined with respect ot the equivalence relation and is associative.
- Concatenating with the trivial loop does not change a loop.
- Every loop $(e_1,\dots,e_n)$ has an inverse $(e_n^{-1},\dots,e_1^{-1})$.

This group is called the **fundamental group** of G, denoted by $\pi_1(G,x)$.

Theorem 4.3. *Let T be a spanning tree of G containing the vertex x. Then the fundamental group $\pi_1(G,x)$ is isomorphic to the free group generated by the edges of G that are not in T.*

PROOF. Let C denote the set of edges not in T. For every arc a, one can associate a loop based at x denoted by $\gamma_a^T=\gamma^T_{x\to o(a)}\cdot a\cdot\gamma^T_{t(a)\to x}$, where $\gamma^T_{x\to y}$ denotes the unique reduced path in T between x and y. Then, every loop $(e_1,\dots,e_n)$ based at x in G is homotopic to the loop $\gamma^T_{e_1},\dots,\gamma^T_{e_n}$, and for every arc a in T, γ_a^T is homotopic to the constant loop. Therefore, $\pi_1(G,x)$ is generated by the loops γ_a^T for a and arc not in T, and since $\gamma^T_{a^{-1}}=(\gamma^T_a)^{-1}$, it is enough to pick one arc for every edge of C. Finally, since the loop γ_a^T for a not in T is the only one containing a, there is no non-trivial relation between the loops γ_a^T. This proves the theorem. □

An alternative way of seeing this proof is to observe that the fundamental group of G is the same as the fundamental group of G obtained after contracting a spanning tree of G. The resulting graph is a bouquet of circles, and there is one generator for each circle.

4.2.2 Fundamental groups of surfaces

Now, let S denote a connected surface and let G be a graph cellularly embedded on S. Similarly as before, a **loop** and a **path** in (S,G) is a path, respectively a loop in G. An **elementary homotopy** between two loops is either a reduction (deletion/addition of ee^{-1}) or the deletion/addition of a subpath bounding a face of G. This corresponds to the idea that in a continuous deformation between two curves, one can flip a face of a cellularly embedded graph, see Figure 15. Elementary homotopies induce an equivalence relation, called **homotopy** between loops based at a common point, denoted by $\equiv$.

As before, the set of homotopy classes of loops based at a vertex x of G forms a group, where the law is the concatenation. Indeed,

- If $\gamma_1\equiv\gamma_1'$ and $\gamma_2\equiv\gamma_2'$ are two pairs of homotopic loops, then their concatenations are homotopic : $\gamma_1\gamma_2\equiv\gamma_1'\gamma_2'$.
- The concatenation law is associative.

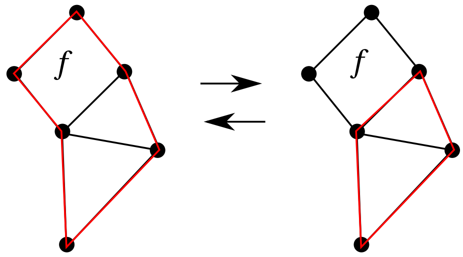

Figure 15: The two red paths are homotopic since they are related by going over the face f.

- The trivial loop, denoted by 1_x or simply 1 is a neutral element for the concatenation law.
- Every loop $(e_1,\ldots,e_n)$ admits $(e_n^{-1},\ldots,e_1^{-1})$ as an inverse.

This group is called the **fundamental group** of S, denoted by $\pi_1(S,x)$[4].

Exercise 4.4. Let x and y be two vertices of G, then show that $\pi_1(S,x)$ and $\pi_1(S,y)$ are isomorphic. This justifies the common abuse of notation to just write $\pi_1(S)$ without specifying a base point.

Let T be a spanning tree of G containing a vertex x, and let C denote the set of edges not in T. The fundamental group of G is the free group on C, and to obtain the fundamental group of S from it, one just needs to add the relations corresponding to the faces of G. Formally, for every face $f=(e_1,\ldots,e_n)$ of G, denote by r_f the **facial relation** induced by f on C, that is, the word obtained by only keeping the e_i that are in C. Then we have the following theorem:

Theorem 4.5. *Let S be a connected surface, G be a cellularly embedded graph on S with a vertex x and a set of faces F, and T be a spanning tree of G containing x. Denote by C the set of edges not in T and by r_f the facial relation induced by a face f of G on C. Then $\pi_1(S)$ is isomorphic to the group π presented by*

$$< C \mid \{r_f\}_{f\in F} > .$$

PROOF. As before, every arc a in G corresponds to a loop γ_a^T obtained by adjoining the reduced paths in T between x and the endpoints of a. Let us consider the map $\gamma: C \rightarrow \pi_1(S)$ mapping every arc a in C to γ_a^T. This map induces a morphism of groups $\gamma': F(C) \rightarrow \pi_1(S)$. As for homotopy in graphs, this map is surjective since any loop of S is the image by γ of its arcs in C. We will show that its kernel equals the normal subgroup N generated by the elements r_f for $f \in F$, which proves the theorem.

Let $w = e_1,\ldots e_n$ be an element of $F(C)$ such that $\gamma'(w) \equiv 1$. Then it means that $\gamma'(w)$ can be reduced to the trivial loop by a sequence of elementary homotopies. Then for every reduction over edges of C, the same reduction can be applied to w, and for every face flip over a face f, the corresponding facial relation r_f can be used to modify

[4]To be accurate, we should write $\pi_1(S,G,x)$ but as we will shortly see, this actually does not depend on G.

the word w. Thus w can be reduced to the trivial word using the set of relations r_f and the kernel of γ' is included in N. Reciprocally, an element of N is mapped to the trivial loop by γ' since the facial relations r_f dictate the face flips to do to simplify the corresponding relations. This concludes the proof. □

Now, let us observe that the group $\pi(S)$ stays the same when one

1. contracts an edge of G between two different endpoints,

2. or one removes an edge of G between two different faces.

For 1, let e be the edge we are contracting, and T be a spanning tree of G containing e. This contraction yields a new tree T' with one less edge, but the set C of non-tree edges stays the same, as well as the set of facial relations. So the group stays the same.

For 2, when one removes an edge e of G between two different faces, one merges the two adjacent faces f_1 and f_2 into a single face f. One can pick the spanning tree so that e is **not** in it, and thus $\pi_1(S)$ lost one generator g, and the two relations r_1 and r_2 containing g have been merged into one. Observing that this amounts to deleting every appearance of g in $\pi_1(S)$ using r_1 (or r_2), we see that this operation does not change the group.

Therefore, by the classification of surfaces (or rather its proof), we see that the graph G can be transformed into one of the polygonal schemata of Theorem 1.6 without changing the fundamental group. In particular, $\pi_1(S)$ only depends on the surface S and not the graph G, and it is isomorphic to

- the trivial group if S is a sphere,
- the group $< a_1, b_1, \dots a_g, b_g \mid a_1 b_1 \bar{a}_1 \bar{b}_1 \dots a_g b_g \bar{a}_g b_g >$ if S is the orientable surface of genus g,
- or $< a_1, \dots, a_g \mid a_1 a_1 \dots a_g a_g >$ if S is the non-orientable surface of genus g.

Remark: The operations of contraction and deletion of edges used above can be interpreted in the light of **dual graphs**: a graph G embedded on a surface S has a dual graph G^*, defined by placing one vertex in each face of G and edges between adjacent faces. Then, contracting an edge in the primal graph amounts to removing an edge in the dual graph, and vice versa. Contracting every edge between different endpoints and removing every edge between different faces amounts to contracting a spanning tree T of G and a spanning tree T^* of G^* which are **interdigitating**, that is, such that the edges of T are not duals of edges of T^*, see Figure 16. The use of such interdigitating trees, also called a **tree-cotree decomposition**, is an important tool in the study of embedded graphs, especially from an algorithmic point of view, but we will not rely on it further in this course.

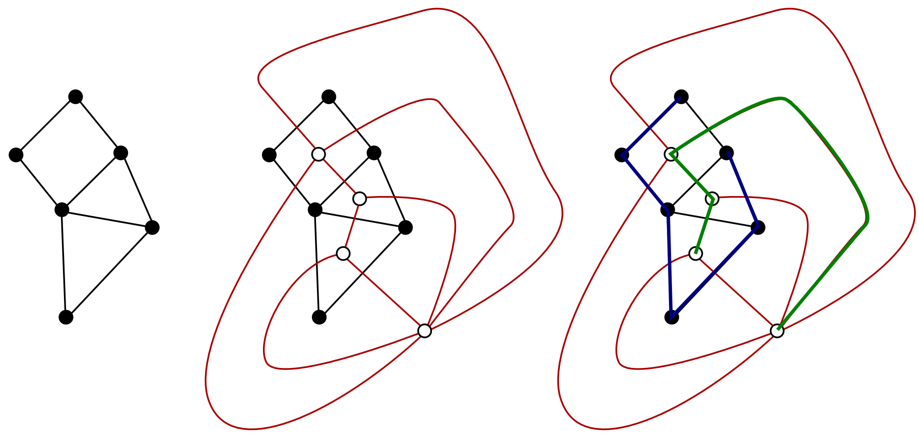

Figure 16: A planar graph, its dual graph and a pair of interdigitating spanning trees.

4.3 Fundamental groups, the topological way

The homotopies we defined in the previous section are very combinatorial, and do not match our a priori intuition of a continuous deformation. In this section, we define homotopies in a topological way, and show that the corresponding notion of fundamental group matches the one obtained before.

In a purely topological setting, we are now only considering a surface S, without any mention of a cellularly embedded graph. A **path** on S is a continuous map $p:[0,1]\rightarrow S$, and a **loop** based at x is a path $p:[0,1]\rightarrow S$ where $p(0)=p(1)=x$. A **homotopy** with basepoint x between two loops ℓ_1 and ℓ_2 is a continuous map $h:[0,1]\rightarrow[0,1]\rightarrow S$ such that $h(0,\cdot)=\ell_1$, $h(1,\cdot)=\ell_2$ and $h(\cdot,0)=h(\cdot,1)=x$. The **constant loop** at x is the loop $p:[0,1]\mapsto x$. The **inverse** of a loop ℓ^{-1} is defined by $\ell^{-1}(t)=\ell(1-t)$. The **homotopy class** of a loop is the set of loops homotopic to it.

The **concatenation** of two loops ℓ_1 and ℓ_2 is the loop defined by $\ell_1(2t)$ for $t\in[0,1/2]$ and $\ell_2(2t-1)$ for $t\in[1/2,1]$. The set of homotopy classes of loops based at x forms a group for the concatenation law, where the neutral element is the constant loop. Indeed:

- If $\gamma_1\equiv\gamma_1'$ and $\gamma_2\equiv\gamma_2'$ are two pairs of homotopic loops, then their concatenations are homotopic : $\gamma_1\gamma_2\equiv\gamma_1'\gamma_2'$.
- The concatenation law is associative.
- The constant loop, denoted by 1_x or simply 1 is a neutral element for the concatenation law.
- The inverse of a loop is its inverse for the concatenation law.

This group is called the **fundamental group** $\pi_1(S,x)$.

Remark: We are always working with loops with basepoints, and the homotopies preserve this basepoint. This is needed to obtain a nice algebraic structure: otherwise there is no natural way to concatenate loops. But the study of homotopies without basepoints, called **free homotopies**, is arguably more natural. We will see later on how to include it in this framework.

The following exercise mirrors Exercise 4.4:

Exercise 4.6. If S is a connected surface, then for every $(x, y) \in S$, the groups $\pi_1(S, x)$ and $\pi_1(S, y)$ are isomorphic. This justifies the common abuse of notation to just write $\pi_1(S)$ without specifying a base point.

As the notations suggest, the topological fundamental group and the combinatorial group turn out to be isomorphic, this is the point of the following theorem.

Theorem 4.7. *The topological fundamental group is the same as the combinatorial fundamental group.*

PROOF (SKETCH). Just for the time of this proof, let us denote respectively by $\pi_1^{comb}(S)$ and $\pi_1^{top}(S)$ the combinatorial and topological fundamental groups. We pick a cellularly embedded graph G on S, which will be used to study $\pi_1^{comb}(S)$ (but as we saw, the group itself does not depend on G). We will study the map $\varphi : \pi_1^{comb}(s) \to \pi_1^{top}(S)$ mapping a homotopy class of loops to the corresponding topological homotopy class of loops.

Claim 1: The map φ is well-defined and is a morphism of groups.

Indeed, if two combinatorial loops γ_1 and γ_2 are homotopic, then they are related by a sequence of reductions and face flips. Such reductions and face flips can be realized using topological homotopies, so their images $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ are homotopic. This map behaves nicely under composition laws, so it is a morphism.

Claim 2: The map φ is surjective.

Let γ be a topological loop on S. It suffices to prove that it is homotopic to a combinatorial loop of G. By perturbing by a very local homotopy if needed, one can assume that γ crosses G a finite number of times. Then between each pair of crossings, on can push γ on one side (for example the left one), so that one obtains a homotopic loop lying entirely in G. Now, it may happen that γ backtracks in the middle of an edge of G, but using a homotopy, one can reduce it so that it does not happen, and thus we obtain a combinatorial loop of G.

Claim 3: The map φ is injective.

Let γ be a combinatorial loop such that $\varphi(\gamma) = 1$. This means that some topological homotopy contracts $\varphi(\gamma)$ to the empty loop. We want to discretize this homotopy so that it becomes a concatenation of face flips and reductions. To do that, push every loop in the topological homotopy into a combinatorial loop of G using the previous technique. By construction, the difference between two consecutive such loops will be a series of reductions or face flips, and thus we obtain a combinatorial homotopy.

□

Remark: The "map" π_1 associates a group for any surface, and it can furthermore be seen as a **functor** : continuous maps between surfaces also induce morphisms between their fundamental groups: indeed, such a continuous applications maps loops to loops, and by taking their homotopy classes, one obtains a morphism. Such functors are the playground of **category theory**, which has deep connections with algebraic topology, but we will not delve at all into these aspects in this course.

All in all, we have just seen that properties regarding continuous deformations of loops can be rephrased in a purely group-theoretical point of view. This is very fruitful from a conceptual perspective, as it provides a strong algebraic structure to work with. But from an algorithmic perspective, the benefits are not that immediate : the issue is that working with presentations of groups is very unwieldy, as struggling with the following exercise showcases:

Exercise 4.8. Show that the fundamental groups of non-homeomorphic surfaces are not isomorphic.

In fact, most computational problems for group presentations are actually **undecidable**. This is the case for example for the following problems:

- Deciding whether two groups provided by a finite presentation are isomorphic.
- Deciding whether a given group provided by a finite presentation is trivial.
- Deciding whether an element in a group provided by a finite presentation is trivial.

We will establish such undecidability results in a later chapter and see that they translate into undecidable topological problems in higher dimensions. Fortunately, fundamental groups of surfaces are simpler than general groups, and thus we will be able to devise algorithms for homotopy on surfaces, but these algorithms will have a very strong geometric or topological appeal, instead of a group theoretical one. More generally, studying groups by realizing them as fundamental groups of some topological space is one of the drives of **combinatorial group theory**.

4.4 Covering spaces

A **covering space** of S is a space $\widehat{S}$ together with a continuous surjective map $\pi : \widehat{S} \to S$ such that for every $x \in S$, there exists an open neighborhood U of x such that $\pi^{-1}(U)$ is a disjoint union of homeomorphic copies of U. The reader scared by this definition should look at the example of the annulus on the left of Figure 17.

We will only deal with covering spaces in an informal way, and refer to a standard textbook in algebraic topology like Hatcher [Hat02] for more precise statements.

The reason we are interested in covering spaces is that they are deeply connected with homotopy and fundamental groups. Indeed, a covering space allows to **lift** a path: if p is a path on S such that $p(0) = x = \pi(\widehat{x})$ for some $\widehat{x} \in \widehat{S}$, there is a unique path $\widehat{p}$ on $\widehat{S}$ starting at $\widehat{x}$ such that $p = \pi \circ \widehat{p}$. This is pictured in Figure 17. Note that loops do not lift necessarily to loops! This "unique lifting property" derives from the definition of covering spaces: when one sits at a point $\widehat{x}$ of the covering space, the local homeomorphism π specifies how to move on $\widehat{S}$ so that one follows the path p.

Every surface has a unique covering space $\widetilde{S}$ that is simply connected, that is, where every loop in $\widetilde{S}$ is homotopic to a trivial loop, it is called the **universal cover** of S. If S is a sphere, its universal cover is itself, so let us assume that it is not the case. One way to build this cover is as follows: pick a graph G cellularly embedded on S with a single vertex and a single face (for example one of the graphs used in the classification theorem). Cutting S along G gives a polygon and one can **tile** the plane with this

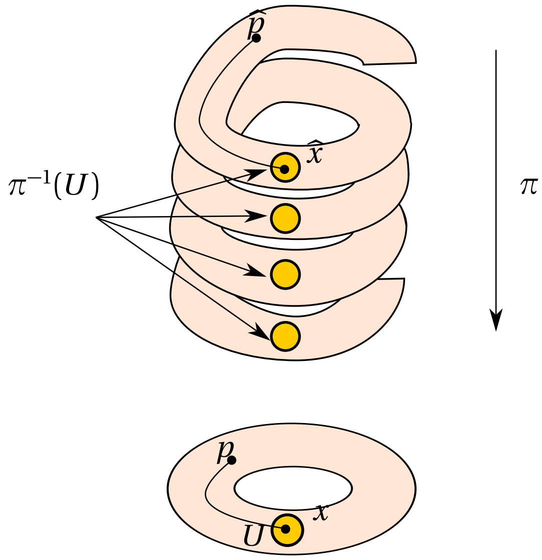

Figure 17: A covering of the annulus, and one lift of a path p.

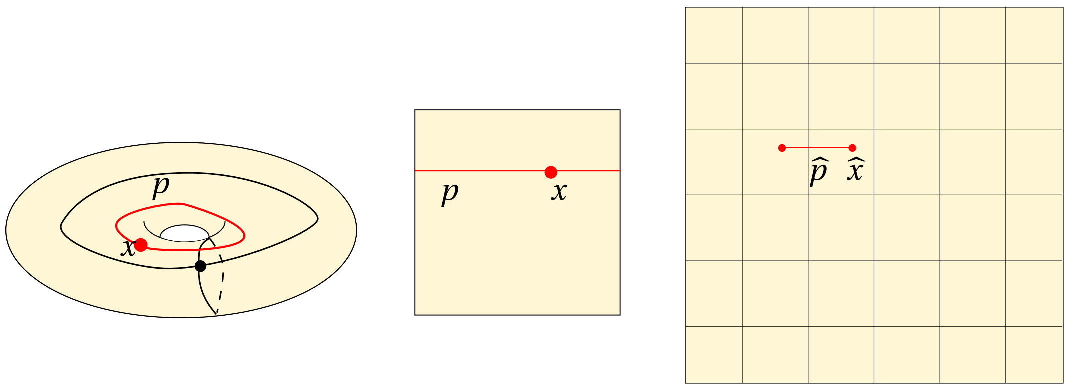

Figure 18: Lifting a loop p on the torus to a path on its universal cover.

polygon by putting adjacent copies of this polygon next to each other, so that every vertex of the tiling is adjacent to the correct number of polygons. This construction is pictured in Figure 18 for the torus.

For surfaces of higher genus, the same construction works, but the tiling will not look as symmetric : indeed it is for example impossible to tile the plane with regular octagons. This is not an issue for our construction, since any tiling with octagons will do, even if they do not have the same shape. However, an insightful way to deal with this issue is to use **hyperbolic geometry**: it is a non-Euclidean geometry on the open disk that allows for regular tilings of polygons with an arbitrary number of faces. Figure 19 pictures the universal cover of a genus 2 surface as a hyperbolic tiling.

One can readily check that the spaces we obtain are universal covering spaces of their respective surfaces, since they are simply connected and the map π can naturally be inferred by the tiling. One key property of universal covers is that a loop on S is contractible if and only if all its lifts in the universal cover are also loops, as can be tested on the above examples. This will be leveraged in the next section to design an algorithm to test contractibility of loops on surfaces.

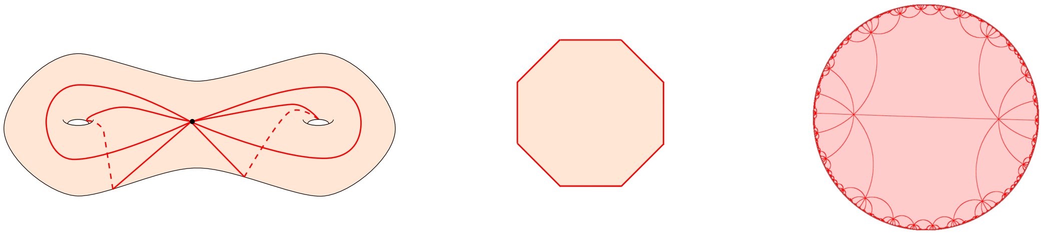

Figure 19: The universal covering space of the genus 2 surface.

Remark: The irruption of hyperbolic geometry here is not random at all: one can show that surfaces of genus at least 2 do not admit Euclidean metrics, but do admit hyperbolic ones. It is the lift of such a metric that one uses to obtain a hyperbolic tiling. Hyperbolic geometry plays a primordial role in the study of the geometric properties of surfaces, and has been used increasingly as well in the design of algorithms for computational topology.

References

[DM68] P. H. Doyle and D. A. Moran. A short proof that compact 2-manifolds can be triangulated. *Inventiones Mathematicae*, 5(2):160–162, 1968.

[Edm60] Jack R. Edmonds. A combinatorial representation of polyhedral surfaces. *Notices of the American Society*, 7:646, 1960.

[Fre82] Michael Hartley Freedman. The topology of four-dimensional manifolds. *J. Differential Geom.*, 17(3):357–453, 1982.

[Hat02] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. Avalaible at http://www.math.cornell.edu/ hatcher/AT/ATpage.html.

[Hef91] Lothar Heffter. Über das Problem der Nachbargebiete. *Math. Ann*, 38:477–508, 1891.

[Hef98] Lothar Heffter. Über metacyklische Gruppen und Nachbarconfigurationen. *Math. Ann*, 50:261–268, 1898.

[KMR08] Ken-ichi Kawarabayashi, Bojan Mohar, and Bruce Reed. A simpler linear time algorithm for embedding graphs into an arbitrary surface and the genus of graphs of bounded tree-width. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 771–780. IEEE, 2008.

[Moh99] Bojan Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM J. Discrete Math.*, 12:6–26, 1999.

[Moi77] Edwin E. Moise. *Geometric topology in dimensions 2 and 3*, volume 47 of *Graduate Texts in Math.* Springer-Verlag, 1977.

[RCS+97] AA Ranicki, A Casson, D Sullivan, M Armstrong, C Rourke, and G Cooke. The hauptvermutung book. *Collection of papers by Casson, Sullivan, Armstrong, Cooke, Rourke and Ranicki, K-Monographs in Mathematics*, 1, 1997.

[Rin74] Gerhard Ringel. *Map Color Theorem*. Springer-Verlag, 1974.

[Sti93] John Stillwell. *Classical topology and combinatorial group theory*. Springer-Verlag, New York, 1993.

[Tho92] Carsten Thomassen. The Jordan-Schönflies Theorem and the Classification of Surfaces. *American Mathematical Monthly*, 99(2):116–130, feb 1992.

[Tho93] Carsten Thomassen. Triangulating a surface with a prescribed graph. *Journal of Combinatorial Theory, Series B*, 57(2):196–206, 1993.