An FPT Algorithm for the Embeddability of Graphs into Two-Dimensional Simplicial Complexes*

Éric Colin de Verdière[†] Thomas Magnard[‡]

Abstract

We consider the embeddability problem of a graph G into a two-dimensional simplicial complex C: Given G and C, decide whether G admits a topological embedding into C. The problem is NP-hard, even in the restricted case where C is homeomorphic to a surface.

It is known that the problem admits an algorithm with running time $f(c)n^{O(c)}$, where n is the size of the graph G and c is the size of the two-dimensional complex C. In other words, that algorithm is polynomial when C is fixed, but the degree of the polynomial depends on C. We prove that the problem is fixed-parameter tractable in the size of the two-dimensional complex, by providing a deterministic $f(c)n^3$-time algorithm. We also provide a randomized algorithm with expected running time $2^{c^{O(1)}}n^{O(1)}$.

Our approach is to reduce to the case where G has bounded branchwidth via an irrelevant vertex method, and to apply dynamic programming. We do not rely on any component of the existing linear-time algorithms for embedding graphs on a fixed surface; the only elaborated tool that we use is an algorithm to compute grid minors.

1 Introduction

An embedding of a graph G into a host topological space X is a crossing-free topological drawing of G into X. The use and computation of graph embeddings is central in the communities of computational topology, topological graph theory, and graph drawing. A landmark result is the algorithm of Hopcroft and Tarjan [HT74], which allows to decide whether a given graph is planar (has an embedding into the plane) in linear time. Related results include more planarity testing algorithms [Pat06], algorithms for embedding graphs on surfaces [Moh99, KMR08] and for computing book embeddings [Mal94], Hanani-Tutte theorems [Sch13], and the theory of crossing numbers and planarization [BCG+06].

In this paper, we describe algorithms for deciding the embeddability of graphs into topological spaces that are, in a sense, as general as possible: two-dimensional simplicial complexes (or 2-complexes for brevity), which are made from vertices, edges, and triangles glued together. (We remark that every graph is embeddable in $\mathbb{R}^3$, so considering higher-dimensional simplicial complexes is irrelevant.) In a previous article, jointly written with Mohar [CdVMM18], we proved that, given a graph G and a 2-complex $\mathscr{C}$, one can decide whether G embeds into $\mathscr{C}$ in polynomial time for fixed $\mathscr{C}$; but the algorithm has running time $f(c) \cdot n^{O(c)}$, where n and c are the respective sizes of G and $\mathscr{C}$. Using a very different strategy, we describe algorithms for this problem, proving that it is fixed-parameter tractable in the complexity of the input complex:

Theorem 1.1. *One can solve the embeddability problem of graphs into 2-dimensional simplicial complexes in deterministic $f(c)n^3$ time or in expected time $2^{c^{O(1)}}n^{O(1)}$, where c is the number of simplices of the input 2-complex, n is the number of vertices and edges of the input graph, and f is some computable function of c.*

2-complexes are much more general than surfaces, and tools that are suitable for studying embeddability of graphs on surfaces do not generalize. For example, the set of graphs embeddable on a given

*Partially supported by the ANR projects Min-Max (ANR-19-CE40-0014) and SoS (ANR-17-CE40-0033).
[†]LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France. Email: eric.colindeverdiere@u-pem.fr
[‡]LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France. Email: thomas.magnard@univ-eiffel.fr

2-complex is not closed under minor, which makes many tools for dealing with graphs on surfaces unsuitable for 2-complexes. Moreover, the complexity of some topological problems increase drastically when we consider 2-complexes instead of surfaces, e.g., deciding homeomorphism [ÓDWW00], or deciding the contractibility of curves [DG99, LR12, EW13]. Some other topological problems, such as the existence of a drawing a graph with at most k crossings in the plane or in a surface [KR07], can be recast as deciding whether the graph embeds on a certain 2-complex. For more detailed motivations, see [CdVMM18, Introduction].

Comparison with previous works on surfaces. Since every surface is homeomorphic to a 2-complex, our problem has been largely considered in the special case where the input 2-complex is (homeomorphic to) a surface. That restricted problem is NP-hard [Tho89], but several algorithms that are fixed-parameter tractable in the genus have been given, which we review now.

Mohar [Moh99] has given an algorithm for this purpose that takes linear time in the input graph, for every fixed surface. This algorithm is very technical and relies on several other articles. The dependence on the genus is not made explicit, but seems to be doubly exponential [KMR08].

Kawarabayashi et al., in an extended abstract [KMR08], have given a simpler linear-time algorithm for this problem, but not all details are presented, which makes the approach hard to check [KP19, p. 3657, footnote].

General graph minor theory provides an algorithm for the same purpose. The graph minor theorem by Robertson and Seymour [RS04] implies that, for every fixed surface $\mathscr{S}$, there is a finite list of graphs $\mathcal{O}_{\mathscr{S}}$ such that a graph G can be embedded on $\mathscr{S}$ if and only if G does not contain any graph in $\mathcal{O}_{\mathscr{S}}$ as a minor. Moreover, there is an algorithm that given any surface $\mathscr{S}$ (specified by its genus and orientability) outputs the list $\mathcal{O}_{\mathscr{S}}$ [AGK08], and there is an algorithm to decide whether a graph M is a minor of another graph G running, for fixed M, in time cubic in the size of G [RS95] [CFK$^+$15, Theorem 6.12]. These considerations thus lead to an algorithm to decide embeddability of a graph on a surface that runs, if the input surface is fixed, in cubic time in the size of the input graph.

Finally, in the same vein, Kociumaka and Pilipczuk [KP19] have studied the following more general problem than the embeddability problem of graphs on surfaces: Given a surface $\mathscr{S}$, a graph G, and an integer $k \geq 0$, is it possible to remove a set U of at most k vertices from G so that $G-U$ is embeddable on $\mathscr{S}$? They provide an algorithm that is fixed-parameter tractable in k and the genus of $\mathscr{S}$, where the dependence on the genus is unspecified. In particular, as a special case, they decide the embeddability of a graph on a surface; however, they use one of the previous algorithms [Moh99, KMR08] as a subroutine. The problem that we study, the embeddability of graphs on 2-complexes, is independent from the problem studied by Kociumaka and Pilipczuk, in the sense that there is, a priori, no obvious reduction from one problem to the other. However, we will reuse some ingredients from that paper.

Our algorithms, restricted to the case where we want to embed graphs on surfaces, are not as efficient as the existing algorithms mentioned above. Indeed, the deterministic one runs in cubic time in the size of the input graph (for a fixed complex); the dependence on the size of the complex is not made explicit, because the algorithm uses, as a subroutine, an algorithm to compute grid minors [RS95]. The second algorithm is randomized, because it uses an algorithmic version of the excluded grid theorem [CC16] that uses randomness; for every fixed surface, it runs in expected time that is a polynomial of fixed (but large) degree in the size of the input graph. However, our algorithms are independent from the existing algorithms for embedding graphs on surfaces; the only elaborated tool that we use is an algorithm to compute grid minors [RS95, CC16].

Overview and structure of the paper. We use a standard strategy in graph algorithms and parameterized complexity (see, e.g., the book by Cygan et al. [CFK$^+$15, Chapter 7]): we show by dynamic programming that the problem can be solved efficiently for graphs of bounded branchwidth, and then, using an irrelevant vertex method, we prove that one can assume without loss of generality that the input graph G has branchwidth bounded by a polynomial in the size of the input 2-complex. In the context of surface-embedded graphs, this paradigm has been used in the extended abstract by Kawarabayashi et al. [KMR08] and in the article by Kociumaka and Pilipczuk [KP19]; our algorithm takes inspiration from the former, for the idea of the dynamic programming algorithm, and from the latter, for some arguments in the irrelevant vertex method. However, handling 2-complexes requires significantly more effort. More precisely, Theorem 1.1 follows immediately from the following two theorems.

Theorem 1.2 (algorithm for bounded branchwidth)**.** *One can solve the embeddability problem of graphs into two-dimensional simplicial complexes in time* $(c+w)^{O(c+w)}n$*, where* c *is the number of simplices of the input 2-complex, and where* n *and* w *are the number of vertices and edges and the branchwidth of the input graph, respectively.*

Theorem 1.3 (algorithm to reduce branchwidth)**.** *Given a 2-complex* $\mathscr{C}$ *with* c *simplices, and a graph* G *with* n *vertices and edges, we can correctly report that* G *is embeddable on* $\mathscr{C}$*, or correctly report that* G *is not embeddable on* $\mathscr{C}$*, or compute a subgraph* H *of* G*, of branchwidth polynomial in the number of simplices of* $\mathscr{C}$*, such that* G *embeds on* $\mathscr{C}$ *if and only if* H *does:*

- *in deterministic time* $f(c) \cdot n^3$ *for some computable function* f*,*
- *or in expected polynomial time.*

We now present the structure of the paper, indicating which techniques are used. We also emphasize which components would be simpler if we were just aiming for an algorithm for embedding graphs on surfaces.

We introduce some standard notions in Section 2.

Then, in Section 3, we show that we can make some simple assumptions on the input, and present data structures for representing 2-complexes and graphs embedded on them. If we restrict ourselves to the case where the input 2-complex is homeomorphic to a surface, we essentially consider combinatorial maps of graphs on surfaces, except that the graphs need not be cellularly embedded (such a data structure is called an *extended combinatorial map* [CdVdM14, Section 2.2]). The case of 2-complexes is largely more involved.

In Section 4, we show that if our input graph G has an embedding into our input 2-complex $\mathscr{C}$, then there exists an embedding of G on $\mathscr{C}$ that is *sparse* with respect to a branch decomposition of G. This means that each subgraph of G induced by the leaves of any subtree of the branch decomposition can be separated from the rest of G using a graph embedded on $\mathscr{C}$, called *partitioning graph*, of small complexity. We find that this new structural result, even in the surface case, is interesting and can prove useful in other contexts. If the target space were a surface, we could assume that G is 3-connected and has no loop or multiple edges, which would imply (still with some work) that *any* embedding of G would be sparse, but again the fact that we consider 2-complexes requires additional work.

In Section 5, we present the dynamic programming algorithm, which either determines the existence of an embedding of G on $\mathscr{C}$, or shows that no sparse embedding of G on $\mathscr{C}$ exists (and thus no embedding at all, by the previous paragraph). The idea is to use bottom-up dynamic programming and to consider all regions of the 2-complex in which the subgraph of G (induced by a subtree of the branch decomposition) can be embedded. The complexity depends exponentially on the branchwidth of G.

The previous arguments, most notably in Section 4, implicitly assumed that, if G has an embedding into $\mathscr{C}$, it has a *proper and cellular embedding*, in particular, in which the faces are homeomorphic to disks. In Section 6, we show that we can assume this property. Essentially, we build all 2-complexes "smaller" than $\mathscr{C}$, such that G embeds on $\mathscr{C}$ if and only if it embeds into (at least) one of these 2-complexes, and moreover if it is the case, it has an embedding into (at least) one of these 2-complexes that is proper and cellular. If $\mathscr{C}$ were an orientable surface, we would just consider the surfaces of lower genus; but here a more sophisticated approach is needed.

The above ingredients allow to prove Theorem 1.2 (Section 7).

In Section 8, we show, using an irrelevant vertex method, that we can assume that G has branchwidth polynomial in the size of $\mathscr{C}$ (Theorem 1.3). In a nutshell, if G has large branchwidth, we can compute a subdivision of a large wall, and then (unless G has large genus and is not embeddable on $\mathscr{C}$) compute a large planar part of G containing a large wall; the central vertex of this wall is irrelevant, in the sense that its removal does not affect the embeddability or non-embeddability of the graph into $\mathscr{C}$; iterating, we obtain a graph of branchwidth polynomial in the size of $\mathscr{C}$.

2 Preliminaries

2.1 Graphs and branch decompositions

In this paper, graphs may have loops and multiple edges unless noted otherwise. Let G be a graph; as usual, we denote by $V(G)$ and $E(G)$ the sets of vertices and edges of G. ***Dissolving*** a degree-two vertex v means replacing v and its incident edges vw_1 and vw_2 with an edge w_1w_2.

A ***(rooted) branch decomposition*** of G is a rooted tree $\mathscr{B}$ in which:

- every node has degree either one (it is a ***leaf***) or three (it is an ***internal node***),
- the root is a leaf,
- each non-root leaf is labelled with an edge of G, and this labelling induces a bijection.

The vertices and edges of $\mathscr{B}$ are called ***nodes*** and ***arcs***, respectively.

Each arc α of $\mathscr{B}$ splits the tree $\mathscr{B}$ into two subtrees $\mathscr{B}_1$ and $\mathscr{B}_2$; if, for $i = 1, 2$, we denote by E_i the set of labels appearing in T_i, we see that α naturally induces a partition (E_1, E_2) of the set of edges of G (if α is the arc incident to the root, then one part of the partition is empty). The ***middle set*** associated to α is the set of vertices of G which are the endpoints of at least one edge in E_1 and at least one edge in E_2.

The ***width*** of $\mathscr{B}$ is the maximum size of a middle set associated to an arc of $\mathscr{B}$. The ***branchwidth*** of G is the minimum width of its (rooted) branch decompositions.

The usual definition of a branch decomposition is identical, except that the tree is unrooted, and thus the leaves are in bijection with the edges of G. Our definition turns out to be more convenient to use in the dynamic program. The difference is cosmetic: From any usual branch decomposition, one can trivially obtain a rooted branch decomposition of the same width, by subdividing an arbitrary arc with a new node ν and then connecting ν to a new leaf node ρ, which will serve as the root; the converse operation obviously transforms any rooted branch decomposition into a usual branch decomposition. Since each usual branch decomposition corresponds to a rooted branch decomposition, and both have the same width, we henceforth only work with rooted branch decompositions.

2.2 Surfaces

We will assume some familiarity with surface topology; see, e.g., [MT01, Sti93, CdV18] for suitable introductions under various viewpoints. We recall some basic definitions and properties. A ***surface*** (without boundary) $\mathscr{S}$ is a compact, connected Hausdorff topological space in which every point has an open neighborhood homeomorphic to the open disk. Up to homeomorphism, every surface $\mathscr{S}$ is obtained from a sphere by:

- either removing $g/2$ open disks and attaching a handle (a torus with an open disk removed) to each resulting boundary component, where g is an even, nonnegative integer called the *(Euler)* ***genus*** of $\mathscr{S}$; in this case, $\mathscr{S}$ is ***orientable***;
- or removing g open disks and attaching a Möbius band to each resulting boundary component, for a positive number g called the (non-orientable) ***genus*** of $\mathscr{S}$; in this case, $\mathscr{S}$ is ***non-orientable***.

A ***possibly disconnected surface*** is a disjoint union of surfaces.

A ***surface with boundary*** is obtained from a surface (without boundary) by removing a finite set of interiors of disjoint closed disks. The boundary of each disk forms a ***boundary component*** of $\mathscr{S}$. The ***genus*** of $\mathscr{S}$ is defined as the genus of the original surface without boundary. Equivalently, a surface with boundary is a compact, connected Hausdorff topological space in which every point has an open neighborhood homeomorphic to the open disk or the closed half disk $\{(x, y) \in \mathbb{R}^2 \mid y \geq 0, x^2 + y^2 < 1\}$.

An embedding of a graph G into a surface $\mathscr{S}$ is ***cellular*** if each face of the embedding is homeomorphic to an open disk. If G is cellularly embedded on a surface with genus g and b boundary components, such that the embedding has v vertices, e edges, and f faces, then Euler's formula stipulates that $v - e + f = 2 - g - b$ (these quantities are referred to as the ***Euler characteristic*** of the surface).

2.3 2-complexes

A ***2-complex*** (or two-dimensional simplicial complex) is an abstract simplicial complex of dimension at most two: a finite set of 0-simplices called ***vertices***, 1-simplices called ***edges***, and 2-simplices called ***triangles***. Each edge is a pair of vertices, and each triangle is a triple of vertices; moreover, each subset of size two in a triangle must be an edge.

Each 2-complex $\mathscr{C}$ corresponds naturally to a topological space, obtained as follows: Start with one point per vertex in $\mathscr{C}$; connect them by segments as indicated by the edges in $\mathscr{C}$; similarly, for every triangle in $\mathscr{C}$, create a triangle whose boundary is made of the three edges contained in that triangle. By abuse of language, we identify $\mathscr{C}$ with that topological space.

2.4 Graph embeddings

Each graph has a natural associated topological space (for graphs without loops or multiple edges, this is a specialization of the definition for 2-complexes). An ***embedding*** Γ of a graph G into a 2-complex $\mathscr{C}$ is an injective continuous map from (the topological space associated to) G to (the topological space associated to) $\mathscr{C}$. A ***face*** of Γ is a connected component of the complement of the image of Γ in $\mathscr{C}$.

3 2-complexes and their data structures

3.1 Some preprocessing

A ***3-book*** is a topological space obtained from three triangles by considering one side per triangle and identifying these three sides together into a single edge. We say that a 2-complex $\mathscr{C}$ ***contains a 3-book*** if $\mathscr{C}$ contains three distinct triangles that share a common edge.

Proposition 3.1. *To decide the embeddability of a graph G on a 2-complex $\mathscr{C}$, we can without loss of generality, after a linear-time preprocessing, assume the following properties on the input:*

- *$\mathscr{C}$ has no 3-book and no connected component that is reduced to a single vertex;*
- *G has no connected component reduced to a single vertex, and at most one connected component homeomorphic to a segment.*

Proof. It is known that every graph can be embedded into a 3-book [CdVMM18, Proposition 3.1]. So we can without loss of generality assume that $\mathscr{C}$ contains no 3-book. We remove all the isolated vertices of $\mathscr{C}$, and remove the same number of isolated vertices of G (to the extent possible); this does not affect whether G embeds into $\mathscr{C}$. We then replace each isolated vertex of G with an isolated edge; since $\mathscr{C}$ has no more isolated vertex, this does not affect embeddability of G into $\mathscr{C}$. Finally, for the same reason, if G contains at least two connected components homeomorphic to segments, we replace all these components with a single edge. □

In the rest of this article, without loss of generality, we implicitly assume that $\mathscr{C}$ and G satisfy the properties stated in Proposition 3.1.

3.2 Structure of 2-complexes without 3-book or isolated vertex

Let $\mathscr{C}$ be a 2-complex without 3-book or isolated vertex, and let p be a vertex of $\mathscr{C}$. Following [CdVMM18, Section 2.2], we describe the possible neighborhoods of p in $\mathscr{C}$. A ***cone at p*** is a cyclic sequence of triangles $t_1, \ldots, t_k, t_1$ ($k \geq 3$), all incident to p, such that, for each $i = 1, \ldots, k$, the triangles t_i and t_{i+1} (where $t_{k+1} = t_1$) share an edge incident with p, and any other pair of triangles have only p in common. A ***corner at p*** is a sequence of distinct triangles $t_1, \ldots, t_k$, all incident to p, such that, for each $i = 1, \ldots, k-1$, the triangles t_i and t_{i+1} share an edge incident with p, any other pair of these triangles have only p in common, and no other triangle in $\mathscr{C}$ shares an edge incident with p and belonging to one of $t_1, \ldots, t_k$. An ***isolated segment at p*** is an edge incident to p but not incident to any triangle. The cones, corners, and isolated segments at p form the ***link components*** at p.

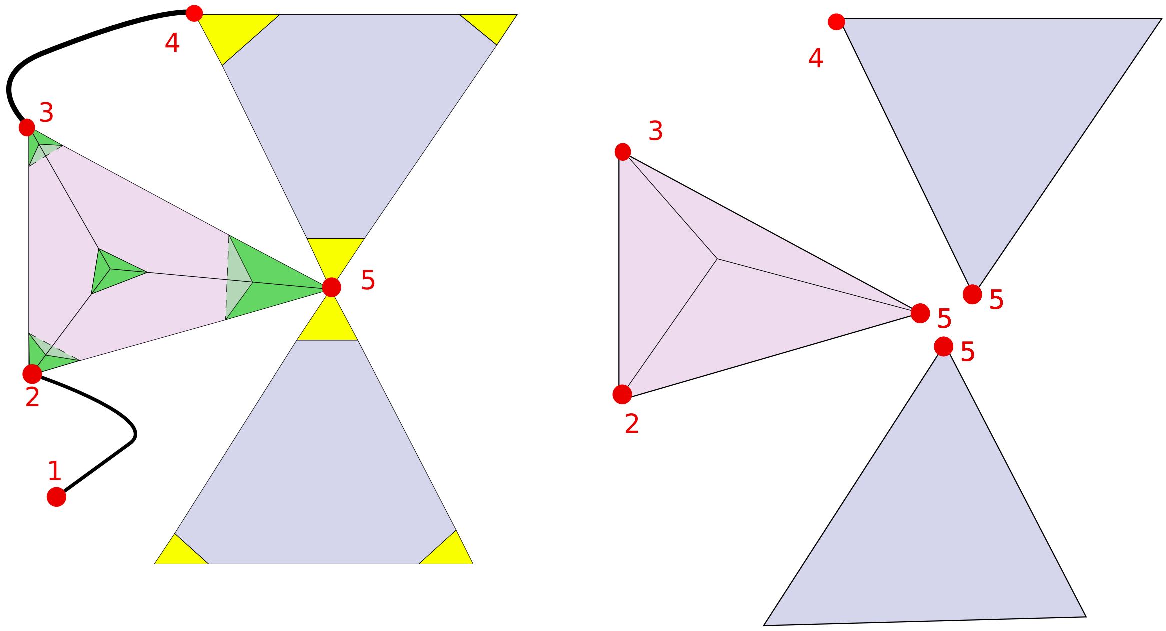

Figure 3.1: On the left: A 2-complex with 5 singular points, numbered from 1 to 5, and 2 isolated edges (one between 3 and 4 and one between 1 and 2) where, at singular points, the cones are in green and the corners in yellow. On the right: the corresponding detached surface.

The set of edges and triangles incident with a given vertex p of $\mathscr{C}$ are uniquely partitioned into cones, corners, and isolated segments. We say that p is a ***regular point*** if all the edges and triangles incident to p form a single cone or corner; in that case, p has an open neighborhood homeomorphic to a disk or a closed half-disk. Otherwise, p is a ***singular point***. See Figure 3.1, left, for an illustration.

Detaching a singular point p in $\mathscr{C}$ consists of the following operation: replace p with new vertices, one for each cone, corner, and isolated segment at p. Detaching all singular points of a 2-complex (without 3-book) yields a disjoint union of (1) isolated segments and (2) a surface, possibly disconnected, possibly with boundary, called the ***detached surface*** (see Figure 3.1, right). The trace of the singular points on the detached surface are the ***marked points***. Conversely, $\mathscr{C}$ can be obtained from a surface (possibly disconnected, possibly with boundary) and a finite set of segments by choosing finitely many subsets of points and identifying the points in each subset together.

The ***boundary*** of $\mathscr{C}$ is the closure of the set of points of $\mathscr{C}$ that have an open neighborhood homeomorphic to a closed half-plane. Equivalently, it is the union of the edges of $\mathscr{C}$ incident with a single triangle.

3.3 Topological data structure for 2-complexes

We now describe a ***topological data structure*** for 2-complexes without 3-book or isolated vertex that is more appropriate for our purposes. It records only the 2-complex up to homeomorphism, not the combinatorial information given by its simplices. Such a 2-complex $\mathscr{C}$ is obtained from a surface $\mathscr{S}$ (possibly disconnected, possibly with boundary) and a finite set S of segments by identifying together finitely many finite subsets of points. Our data structure stores separately the detached surface $\mathscr{S}$, the set S of isolated segments, and the singular points, and two-way pointers representing incidences between them. In more detail:

- we store the list of the connected components of the detached surface $\mathscr{S}$, and for each such component $\mathscr{S}'$ we store (1) whether it is orientable or not; (2) its genus; (3) a list of pointers to the singular points in the interior of $\mathscr{S}'$; (4) for each boundary component of $\mathscr{S}'$, a cyclically ordered list of pointers to the singular points appearing on that boundary component (if $\mathscr{S}'$ is orientable, the boundary components must be traversed in an order consistent with an arbitrarily chosen orientation of $\mathscr{S}'$);

- we store the list S of isolated segments, and for each of them, two pointers to the singular points at its endpoints;
- conversely, to each singular point is attached a list of pointers to the occurrences of that singular point on the detached surface or as an endpoint of an isolated segment.

The ***size*** of a 2-complex (without 3-book or isolated vertex) is the sum of the number of isolated segments, the number of connected components of the detached surface, the total genus of the detached surface, the total number of boundary components of the detached surface, and the total number of marked points of the detached surface (the occurrences of the singular points). This is, up to a constant factor, the size of the topological data structure indicated above, if the genus is stored in unary.

Given a 2-complex $\mathscr{C}$ without 3-book or isolated vertex, described by vertices, edges, and triangles and their incidence relations, we can easily compute a representation of $\mathscr{C}$ in that data structure, in polynomial time: Indeed, by ignoring the incidences created by vertices, we easily build a triangulation of the surface $\mathscr{S}$ (possibly disconnected, possibly with boundary) and a list of segments S; we then compute the topology of $\mathscr{S}$; finally, we mark the singular points, which are the vertices with several occurrences on $\mathscr{S}$ and/or on S. We remark that the size of the resulting data structure is at most linear in the number of vertices, edges, and triangles of the 2-complex $\mathscr{C}$, because, by Euler's formula, any triangulated surface (possibly with boundary) with k simplices has genus $O(k)$ and a number of boundary components that is $O(k)$. Thus, **in the rest of this article, without loss of generality, we implicitly assume that $\mathscr{C}$ is given in the form of the above topological data structure.** (Conversely, it is not hard to see that every 2-complex is homeomorphic to a 2-complex whose number of simplices that is linear in its size, but we will not need this fact.)

We will need the following lemma.

Lemma 3.2. *Given two 2-complexes $\mathscr{C}$ and $\mathscr{C}'$, given in the topological representation above, of sizes c and c' respectively, we can decide whether $\mathscr{C}$ and $\mathscr{C}'$ are homeomorphic in time $(c+c')^{O(c+c')}$.*

Proof. (We remark that this essentially follows from more general results [ÓDWW00]; the running time of our algorithm might be improvable, but this suffices for our purposes.) As a preprocessing, in the topological data structures of $\mathscr{C}$ and $\mathscr{C}'$, we do the following: whenever a singular point is incident to exactly two isolated segments and is not incident to the detached surface, we dissolve that singular point, removing it and replacing the two incident edges with a single one. Clearly, this does not affect whether $\mathscr{C}$ and $\mathscr{C}'$ are homeomorphic.

After this preprocessing, $\mathscr{C}$ and $\mathscr{C}'$ are homeomorphic if and only if their topological data structures are isomorphic. By this, we mean that there is a bijective correspondence φ from the isolated segments, the connected components of the detached surface, and the boundary components of each connected component of the detached surface of $\mathscr{C}$ to those of $\mathscr{C}'$ that preserves the genus, the orientability, the incidences, and the cyclic ordering of the singular points on each boundary component. More precisely, for the latter point: for each connected component C of the detached surface of $\mathscr{C}$, if C is orientable, then the lists of singular points appearing on each boundary component of C and $\varphi(C)$ are identical up to global reversal of all these cyclic orderings simultaneously, corresponding to a change of the orientation of the connected component; if C is non-orientable, then the lists of singular points appearing on each boundary component of C and $\varphi(C)$ are identical up to the possible individual reversal of some of these cyclic orderings. The proof is tedious but straightforward, and the existence of an isomorphism can obviously be tested in the indicated time. □

3.4 Proper and cellular graph embeddings on 2-complexes

Let $\mathscr{C}$ be a 2-complex with size c, G a graph, and Γ an embedding of G on $\mathscr{C}$. The embedding Γ is ***proper*** if:

- the image of Γ meets the boundary of $\mathscr{C}$ only on singular points;
- the vertices of Γ cover the singular points of $\mathscr{C}$.

The embedding Γ is ***cellular*** if each face of Γ is an open disk plus possibly some part of the boundary of $\mathscr{C}$. We emphasize that this definition slightly departs from the standard one. Moreover, we will only consider cellular embeddings that are proper.

Traditional data structures for graphs on surfaces handle graphs embedded cellularly; *rotation systems* [MT01] constitute one example of such a data structure. In order to have efficient algorithms, refined data structures, e.g., with the gem representation [Epp03, Section 2], are needed. The basic element in the gem representation is the *flag*, an incidence between a vertex, an edge, and a face of the graph. Three involutions allow to move from each flag to a nearby flag. Each flag contains a pointer to the underlying vertex, edge, and face.

One can easily extend such data structures to possibly non-cellular embeddings on surfaces [Cd-VdM14, Section 2.2]. In this framework, one must store the topology of each face, which is not necessarily homeomorphic to a disk. Also, a face may have several boundary components; two-way pointers connect each face to one flag of each boundary component (or to an isolated vertex of the graph, if that boundary component is a single vertex); if a face is orientable and has several boundary components, then these pointers must induce a consistent orientation of these boundary cycles. It is important to remark that this data structure also allows to recover the topology of the underlying surface.

Let Γ be a proper graph embedding of a graph G on a 2-complex $\mathscr{C}$ (under the assumptions of Proposition 3.1). Let $\mathscr{S}$ be the detached surface of $\mathscr{C}$. Because Γ is proper, it naturally induces an embedding Γ', of another graph G', on $\mathscr{S}$; some vertices of G located on singular points of $\mathscr{C}$ are duplicated in G', the vertices of G located in the relative interior of isolated segments are absent from G', and the edges of G not in G' are edges on the isolated segments of $\mathscr{C}$. Our data structure, called ***combinatorial map***, for storing the graph embedding Γ and the 2-complex $\mathscr{C}$ consists of storing (1) the graph embedding Γ' on $\mathscr{S}$, as indicated in the previous paragraph, (2) the isolated segments of $\mathscr{C}$, together with, for each such isolated segment, an ordered list alternating vertices and edges of Γ (or, instead of an edge, a mark indicating the absence of such an edge in the region of the isolated segment between the incident vertices), (3) the identifications of vertices of Γ' that are needed to recover Γ (and thus implicitly $\mathscr{C}$).

Isomorphisms between combinatorial maps are defined in the obvious way, similar to the concept of isomorphism between topological data structures: Two combinatorial maps are isomorphic if there is an isomorphism between the combinatorial maps restricted to the detached surfaces, isomorphisms between the maps on each isolated segments, and such that incidences are preserved on the singular points. We can easily test isomorphism between two combinatorial maps of size k and k', respectively, in $(k+k')^{O(k+k')}$ time.

We will need an algorithm to enumerate all proper embeddings of small graphs on a given 2-complex. This is achieved in the following lemma.

Lemma 3.3. *Let $\mathscr{C}$ be a 2-complex of size c and k an integer. We can enumerate the $(c+k)^{O(c+k)}$ combinatorial maps of graphs with at most k vertices and at most k edges properly embedded on $\mathscr{C}$ in $(c+k)^{O(c+k)}$ time.*

Proof. The strategy is the following. In a first step, we enumerate a set of proper graph embeddings on some 2-complexes, which necessarily contains all the desired combinatorial maps. In a second step, we prune this set to keep only the desired combinatorial maps, by eliminating those that contains too many vertices or edges, or that correspond to an embedding on a 2-complex not homeomorphic to $\mathscr{C}$.

First step. Let Γ be a proper embedding of a graph with at most k vertices and k edges on $\mathscr{C}$. Let $\mathscr{S}$ be the detached surface associated to $\mathscr{C}$; this surface is possibly disconnected and has genus at most c. The image of Γ on $\mathscr{S}$ is a graph with at most $k+c$ vertices and k edges.

We first enumerate, in a possibly redundant way, the set M_1 of combinatorial maps of cellular graph embeddings with at most $k+c$ vertices and k edges on a possibly disconnected surface without boundary. There are $2^{O(c+k)}$ such combinatorial maps, which can be enumerated in $2^{O(c+k)}$ time, for example using rotation systems. Some vertices may be isolated, if the corresponding connected component of $\mathscr{S}$ is a sphere.

By simplifying the surface, every graph embedding can be transformed into a cellular graph embedding: remove each face and paste a disk to each cycle of the graph that was a boundary component

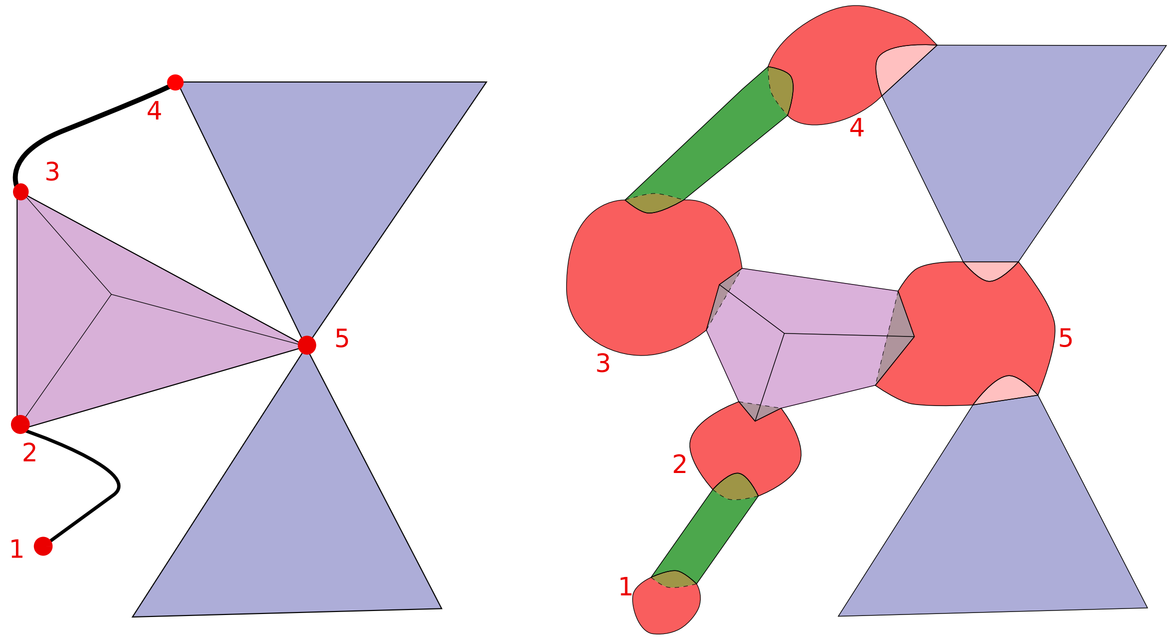

Figure 3.2: On the left: The same 2-complex as Figure 3.1. On the right: the corresponding surface constructed in Lemma 3.4.

of a face. Conversely, every (possibly non-cellular) graph is obtained from some cellular one by (1) connecting some faces together (creating a face of genus zero with several boundary components) and (2) adding some genus and non-orientability to some faces. So, for each map in M_1, we perform all these operations in all possible ways, by putting genus at most c in each face. For each such map in M_1, there are $(c+k)^{O(c+k)}$ possibilities. We thus obtain, in $(c+c)^{O(c+k)}$ time, a set M_2 of $(c+k)^{O(c+k)}$ combinatorial maps on surfaces, and the set M_2 contains all combinatorial maps of (possibly non-cellular) graph embeddings on surfaces of genus at most c.

Finally, we add at most c isolated segments, choose how endpoints of these isolated segments and vertices of the embedding on the detached surface are identified, and decide how each isolated segment is covered by the embedding. There are $(c+k)^{O(c+k)}$ ways to do this. We thus have computed, in $(c+k)^{O(c+k)}$ time, a set M of $(c+k)^{O(c+k)}$ combinatorial maps of graphs on 2-complexes, which contains all the combinatorial maps indicated in the statement of the lemma.

Second step. First, we easily discard the combinatorial maps in M containing more than k vertices or k edges. Then, we discard the maps in M corresponding to a 2-complex different from $\mathscr{C}$. For this purpose, for each map m in M, we iteratively remove the edges of the graph embedding, preserving the underlying 2-complex. When removing an edge from the detached surface, the topology of the incident face(s) change; we preserve this information. Finally, we remove every isolated vertex that does not lie on a singular point of the 2-complex. The data structure that we have now is essentially the one that is described in Section 3.3; we can thus easily decide whether that 2-complex is homeomorphic to $\mathscr{C}$ (Lemma 3.2), and discard m if and only if it is not the case.

Finally, and although this is not strictly needed, we can easily remove the duplicates in the combinatorial maps, by testing pairwise isomorphism between these maps. ☐

3.5 Graphs embeddable on a fixed 2-complex have bounded genus

Lemma 3.4. *Let $\mathscr{C}$ be a 2-complex without 3-book. Let c be either the size of $\mathscr{C}$ or its number of simplices. Every graph embeddable on $\mathscr{C}$ is embeddable on a surface of genus at most* $10c$.

Proof. The strategy is to construct a surface $\mathscr{S}$ of genus at most $10c$ such that every graph embeddable on $\mathscr{C}$ is also embeddable on $\mathscr{S}$. The surface $\mathscr{S}$ is obtained by replacing every isolated segment of $\mathscr{C}$ with a cylinder, and modifying the structure of the 2-complex in the neighborhood of each singular point to make it surface-like; see Figure 3.2.

In more detail, we can obviously assume that $\mathscr{C}$ has no isolated vertex. First, we replace every isolated segment of $\mathscr{C}$ with a cylinder. Then, for every singular point p, we do the following. We remove a small neighborhood of p. We create a sphere with k boundary components, where k is the number of link components at p. Finally, we attach a link component to each of the k boundary components of the sphere:

- for each link component that was a cone, a small neighborhood of p was removed, with boundary a circle; we attach that circle bijectively to a boundary component of the sphere;
- for each link component that was a corner, a small neighborhood of p was removed, with boundary a segment; we attach that segment to a part of a boundary component of the sphere;
- for each link component that was an isolated segment, the isolated segment was replaced with a cylinder; we attach the corresponding boundary component of that cylinder bijectively to a boundary component of the sphere.

(We remark that there are several ways to perform these operations, depending on the orientation of the gluings; any choice will do.) The resulting surface $\mathscr{S}$, which is possibly non-connected and possibly with boundary, has genus at most $10c$. Indeed, if c is the size of $\mathscr{C}$, this follows from Euler's formula (intuitively, the number of "handles" created is at most the number of link components). If c is the number of simplices of $\mathscr{C}$, it follows from the fact that the total genus of the detached surface is at most c, again by Euler's formula, and from the fact that each isolated edge or triangle contributes to an increase of at most six for the genus in the construction above.

Consider an embedding of a graph G on $\mathscr{C}$. It is not hard to transform that embedding into an embedding of G on $\mathscr{S}$: Each cylinder replacing an isolated edge is used only along a single path connecting its two boundary components; if a singular point p is used by the embedding of G on $\mathscr{C}$, we can locally modify the embedding to accommodate the local change from $\mathscr{C}$ to $\mathscr{S}$ at p (see again Figure 3.2). Of course, if G is embeddable on $\mathscr{S}$, it is embeddable on some connected surface of genus at most $10c$. □

4 Partitioning graphs

Let $\mathscr{C}$ be a 2-complex and G a graph, which satisfy the properties of Proposition 3.1. In this section, we lay the structural foundations of the dynamic programming algorithm, described in the next section (Proposition 5.1). The goal, in this section and the following one, is to obtain an algorithm that takes as input $\mathscr{C}$ and G, and, in time FPT in the size of $\mathscr{C}$ and the branchwidth of G, reports correctly one of the following two statements:

- G has no proper cellular embedding on $\mathscr{C}$,
- G has an embedding on $\mathscr{C}$.

This algorithm uses dynamic programming on a rooted branch decomposition of G. When processing a node of the rooted branch decomposition, it considers embeddings of the subgraph of G induced by the edges in the leaves of the subtree rooted at that node in a region of $\mathscr{C}$. This region will be delimited by a *partitioning graph* embedded on $\mathscr{C}$. Our dynamic program will roughly guess the partitioning graph at each node of the rooted branch decomposition. For this purpose, we need that, if G has a proper cellular embedding on $\mathscr{C}$, it has such an embedding that is *sparse*: at each node of the rooted branch decomposition of G, the partitioning graph corresponding to the embedding of the induced subgraph is small (its size is upper-bounded by a function of the branchwidth of G and of the size of $\mathscr{C}$). The goal of this section is to prove that this is indeed the case.

Let $(E_1, \ldots, E_k)$ be an (ordered) partition of the edge set $E(G)$ of G. (We will only use the cases $k = 2$ or $k = 3$.) The ***middle set*** of $(E_1, \ldots, E_k)$ is the set of vertices of G whose incident edges belong to at least two sets E_i.

Let Γ be a proper cellular embedding of G on $\mathscr{C}$. Since Γ is cellular, every boundary of $\mathscr{C}$ is incident to at least one vertex of Γ. Let $\hat{\Gamma}$ be obtained from Γ by adding edges as follows: for any pair

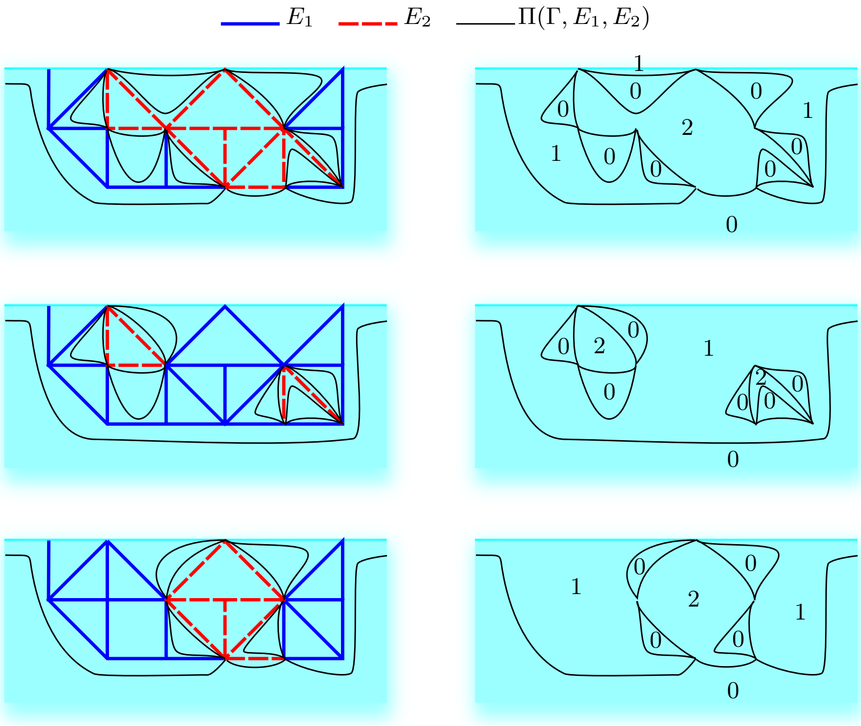

Figure 4.1: Construction of the partitioning graph $\Pi = \Pi(\Gamma, E_1, E_2)$, for three choices of the partition (E_1, E_2) of the same embedding Γ. Only a part of the 2-complex $\mathscr{C}$ is shown, with a boundary at the upper part, and without singular point. Left: The graph embeddings Γ (in thick lines) and Π (in thin lines). Right: The sole graph Π, together with the labelling of its faces.

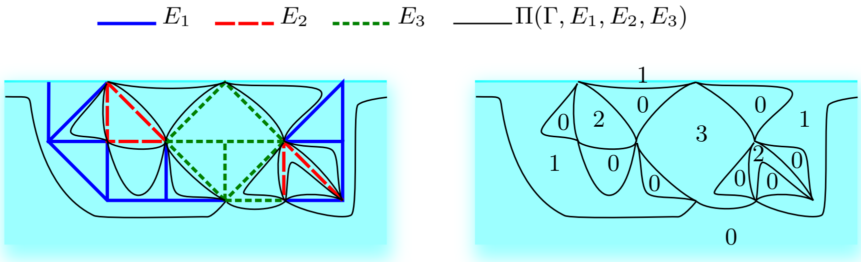

Figure 4.2: The partitioning graph $\Pi = \Pi(\Gamma, E_1, E_2, E_3)$. Left: The graph embeddings Γ (in thick lines) and Π (in thin lines). Right: The sole graph Π, together with the labelling of its faces.

of vertices u and v of Γ consecutive along a given boundary component of $\mathscr{C}$, we connect u and v via a new edge that runs along the boundary component. For each (ordered) partition $(E_1, \ldots, E_k)$ of the edge set of G, we let $\hat{E}_1$ be the union of E_1 and of these new edges, and $\hat{E}_i = E_i$ for each $i \neq 1$; thus, $(\hat{E}_1, \ldots, \hat{E}_k)$ is a partition of the set of edges of $\hat{\Gamma}$.

The ***partitioning graph*** $\Pi(\Gamma, E_1, \ldots, E_k)$ (or more concisely Π) associated to Γ and $(E_1, \ldots, E_k)$ is a graph properly embedded on $\mathscr{C}$ (but possibly non-cellularly), with labels on its faces, defined as follows:

- The vertex set of Π is the union of the singular points of $\mathscr{C}$ and of (the images under Γ of) the middle set of $E_1, \ldots, E_k$.
- The relative interiors of the edges of Π are disjoint from the edges of $\hat{\Gamma}$ and from the isolated segments of $\mathscr{C}$. Let f be a face of $\hat{\Gamma}$ (which is homeomorphic to an open disk plus possibly some points of the boundary of $\mathscr{C}$). Let us describe the edges of Π inside f.

  If, for some $i \in \{1, \ldots, k\}$, the boundary of f is comprised only of edges of $\hat{\Gamma}$ that lie in a single set $\hat{E}_i$, then Π contains no edge inside f. Otherwise, the boundary of f is a succession of edges of $\hat{E}_1$, $\hat{E}_2$, ..., $\hat{E}_k$. The edges of Π inside f run along the boundary of f; for each $i \in \{1, \ldots, k\}$, for each (maximal) group of consecutive edges in $\hat{E}_i$ along the boundary of f, we create an edge of Π that runs along this group, with endpoints the corresponding vertices on the boundary of f (see Figures 4.1 and 4.2). These vertices either are in the middle set of $(E_1, \ldots, E_k)$, or lie on the boundary of $\mathscr{C}$ (and thus on singular points of $\mathscr{C}$).

It follows from the construction that $\hat{\Gamma}$ and Π intersect only at common vertices.

- Each face of Π is labelled by an integer in $\{0, \ldots, k\}$ as follows: faces of Π containing edges in $\hat{E}_i$ are labelled i, and the other ones are labelled 0. By construction of the graph Π, each face of Π contains edges from at most one set $\hat{E}_i$, so this labelling is well defined.

Henceforth, we fix a rooted branch decomposition $\mathscr{B}$ of G, the root of which is denoted by ρ. Every arc α of $\mathscr{B}$ naturally partitions $E(G)$ into two parts E_1 and E_2, in which E_1 is the part on the same side as ρ; this (ordered) partition is the ***edge partition associated to*** α. Recall that Γ is a proper and cellular embedding of G on $\mathscr{C}$; we let $\Pi(\Gamma, \alpha)$ be $\Pi(\Gamma, E_1, E_2)$. Similarly, every node ν of $\mathscr{B}$ naturally partitions $E(G)$ into three parts E_1, E_2, and E_3, in which E_1 is the part on the same side as ρ; this partition is the ***edge partition associated to*** ν; we let $\Pi(\Gamma, \nu)$ be $\Pi(\Gamma, E_1, E_2, E_3)$.

We say that Γ is ***sparse*** (with respect to $\mathscr{B}$) if the following conditions hold, letting c be the size of $\mathscr{C}$ and w the width of $\mathscr{B}$:

- For each arc α of $\mathscr{B}$, the graph $\Pi(\Gamma, \alpha)$ has at most $74c + 26w$ edges;
- similarly, for each internal node ν of $\mathscr{B}$, the graph $\Pi(\Gamma, \nu)$ has at most $3(74c + 26w)$ edges.

The result of this section is the following.

Proposition 4.1. *Let $\mathscr{C}$ be a 2-complex and G a graph, which satisfy the properties of Proposition 3.1. Let $\mathscr{B}$ be a rooted branch decomposition of G. Assume that G has a proper cellular embedding on $\mathscr{C}$. Then it has a proper cellular embedding Γ on $\mathscr{C}$ that is sparse (with respect to $\mathscr{B}$).*

4.1 Monogons and bigons

A ***monogon*** of a graph Π embedded on a 2-complex $\mathscr{C}$ is a face of Π that is an open disk whose boundary is a single edge of Π (a loop). Similarly, a ***bigon*** of Π is a face of Π that is an open disk whose boundary is the concatenation of two edges of Π (possibly the same edge appearing twice). The following general lemma on graphs embedded on surfaces without monogons or bigons will be used; some particular cases have been used before [CCdVE+08, Lemma 2.1].

Lemma 4.2. *Let $\mathscr{S}$ be a surface of genus g without boundary. Let Π be a graph embedded on $\mathscr{S}$, not necessarily cellularly. Assume that Π has no monogon or bigon. Then $|E(\Pi)| \leq \max\{0, 3g + 3|V(\Pi)| - 6\}$.*

Proof. We begin by adding edges to Π as long as it is possible to do so, without introducing any new vertex, monogon, or bigon. Let Π' be the resulting embedded graph. We claim that every face of Π' is a triangle (a disk incident with three edges), except in the following cases:

1. Π' is the empty graph;
2. $\mathscr{S}$ is a sphere, and Π' has two vertices and no edge;
3. $\mathscr{S}$ is a sphere, and Π' has a single vertex and no edge;
4. $\mathscr{S}$ is a projective plane, and Π' has a single vertex and no edge.

Indeed, let f be a face of Π'. If f has no boundary component, then since $\mathscr{S}$ is connected, we are in Case 1 (an isolated vertex would account for a boundary component). Assume that f has at least two boundary components. We add an edge in f connecting vertices on these two boundary components. This cannot create any monogon or bigon, except if the two boundary components are both reduced to

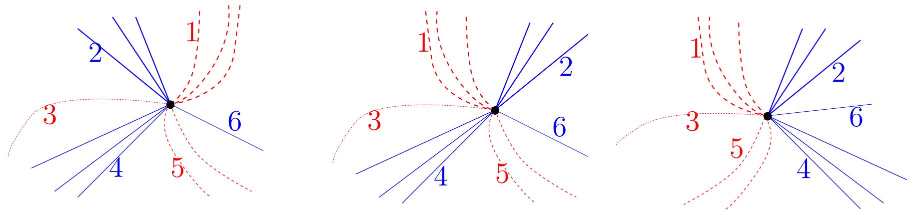

Figure 4.3: Left: A vertex with 6 intervals, numbered from 1 to 6. Middle: The cyclic order obtained by applying the first type of simplification operation on intervals 1 and 2. After the simplification, the intervals 1 and 3 are merged into a single one, and similarly for the intervals 2 and 6. Right: The cyclic order obtained by applying the second type of simplification to the configuration on the left, on pairs of intervals $\{1, 2\}$ and $\{4, 5\}$. After the simplification, the intervals 1, 3, and 5 are merged, and similarly for the intervals 2, 6, and 4.

a single vertex and $\mathscr{S}$ is a sphere (Case 2). So we can assume that f has a single boundary component. If f is orientable and has genus zero, then either we are in Case 3, or $\mathscr{S}$ is a disk of degree at least four, in which case we can add an edge to split it into smaller disks without creating any monogon or bigon. If f is orientable and has Euler genus at least two (i.e., orientable genus at least one), we can add an edge that forms a non-separating arc (relatively to the boundary) in f; it does not form any monogon or bigon. If f is non-orientable and has (non-orientable) genus at least two, we can add an edge that forms a separating arc in f, cutting that surface into two non-orientable surfaces of genus at least one; it does not form any monogon or bigon. Finally, if f is non-orientable and has (non-orientable) genus one, either the boundary component is reduced to a single vertex, so we are in Case 4, or this face has degree at least one; then it is a Möbius band with at least one vertex and one edge on its boundary, and we can add a loop that is a non-contractible arc (relatively to the boundary) in f, without forming any monogon or bigon. The only remaining possibility is that f is a triangle.

It is clear that the statement of the lemma holds whenever we are in one of the four above cases. So we can assume without loss of generality that each face of Π' is a triangle. Since $V(\Pi) = V(\Pi')$ and $|E(\Pi)| \leq |E(\Pi')|$, it suffices to prove the result for Π' instead of Π. Double-counting the incidences between edges and faces implies that the number of triangles τ satisfies $3\tau = 2|E(\Pi')|$; plugging this into Euler's formula implies that $|V(\Pi'| - |E(\Pi')|/3 = 2 - g$, so $|E(\Pi')| = 3g + 3|V(\Pi')| - 6$, as desired. ☐

4.2 Vertex simplifications

The proof of Proposition 4.1 starts with any proper cellular embedding of Γ, and iteratively changes the cyclic ordering of edges around vertices in a specific way. Let (E_1, E_2) be an (ordered) partition of $E(G)$, let v be a vertex of G, and let C be a link component at v (if the image of v under Γ is a singular point, there may be several such link components). We restrict our attention to the edges of $\hat{\Gamma}$ incident to v and belonging to C, in cyclic order around v. For $i = 1, 2$, an ***interval*** (at v, relatively to $(\hat{E}_1, \hat{E}_2)$) is a maximal contiguous subsequence of edges in this cyclic ordering that all belong to $\hat{E}_i$; the interval is labelled i. ***Simplifying*** v (with respect to (E_1, E_2)) means changing the cyclic ordering of the edges of $\hat{\Gamma}$ incident to v in C by one of the two following operations (Figure 4.3):

1. either exchanging two consecutive intervals in that ordering, in a way that the ordering of the edges in each interval is preserved; this operation is allowed only if v is incident to at least four intervals;

2. or performing the previous operation twice, on two disjoint pairs of consecutive intervals in that ordering; this is allowed only if v is incident to at least six intervals.

We will rely on the following lemma.

Lemma 4.3. *Let Γ be a proper cellular embedding of G on $\mathscr{C}$, and let (E_1, E_2) be an (ordered) partition of $E(G)$. Let Γ' be another proper cellular embedding of G, obtained from Γ by simplifying one or two vertices with respect to (E_1, E_2), while keeping the other cyclic orderings unchanged. Then:*

1. $|E(\Pi(\Gamma', E_1, E_2))| < |E(\Pi(\Gamma, E_1, E_2))|$*;*

2. for each (ordered) partition $(\tilde{E}_1, \tilde{E}_2)$ *of* $E(G)$ *such that* $\hat{\tilde{E}}_i \subseteq \hat{E}_j$ *for some* $i, j \in \{1, 2\}$*, we have* $|E(\Pi(\Gamma', \tilde{E}_1, \tilde{E}_2))| \leq |E(\Pi(\Gamma, \tilde{E}_1, \tilde{E}_2))|$.

Proof. The proof is based on the following easy but key observations (the second one will be reused later):

- A simplification of v strictly decreases the number of intervals at v;
- the number of half-edges of $\Pi(\Gamma, E_1, E_2)$ at v in the link component C equals twice the number of intervals associated to $(\hat{E}_1, \hat{E}_2)$ at v in C.

The first point of the lemma immediately follows. For the second point, let us consider, in the cyclic ordering around v in C, a maximal contiguous sequence of edges in $\hat{\tilde{E}}_i$. Since $\hat{\tilde{E}}_i \subseteq \hat{E}_j$, when simplifying with respect to (E_1, E_2), this sequence is still contiguous in the new embedding Γ'. It follows that the number of intervals associated to $(\hat{\tilde{E}}_1, \hat{\tilde{E}}_2)$ does not increase when replacing Γ with Γ'. □

4.3 Rearranging Γ with respect to an edge partition

We can now prove the following lemma:

Lemma 4.4. *Let* Γ *be a proper cellular embedding of* G *on* $\mathscr{C}$*, and let* (E_1, E_2) *be an (ordered) partition of* $E(G)$*. There exists a proper cellular embedding* Γ' *of* G *such that:*

- $|E(\Pi(\Gamma', E_1, E_2))| \leq 74c + 26w$*, where* w *is the size of the middle set of* (E_1, E_2)*;*
- *for each (ordered) partition* $(\tilde{E}_1, \tilde{E}_2)$ *of* $E(G)$ *such that* $\hat{\tilde{E}}_i \subseteq \hat{E}_j$ *for some* $i, j \in \{1, 2\}$*, we have* $|E(\Pi(\Gamma', \tilde{E}_1, \tilde{E}_2))| \leq |E(\Pi(\Gamma, \tilde{E}_1, \tilde{E}_2))|$.

Proof. Here is an overview of the proof. Let $\Pi := \Pi(\Gamma, E_1, E_2)$. We will assume that Π has "many monogons or bigons" (in a sense made precise below) and show that there is another cellular embedding Γ' of G such that:

- $|E(\Pi(\Gamma', E_1, E_2))| < |E(\Pi(\Gamma, E_1, E_2))|$;
- for each (ordered) partition $(\tilde{E}_1, \tilde{E}_2)$ of $E(G)$ such that $\hat{\tilde{E}}_i \subseteq \hat{E}_j$ for some $i, j \in \{1, 2\}$, we have $|E(\Pi(\Gamma', \tilde{E}_1, \tilde{E}_2))| \leq |E(\Pi(\Gamma, \tilde{E}_1, \tilde{E}_2))|$.

By repeatedly iterating this argument, and up to replacing Γ with Γ', this implies that we can assume without loss of generality that Π has "not too many monogons or bigons". We will then show that this latter property implies that Π has at most $74c + 26w$ edges, which concludes.

First, let v be a vertex of Π, and let C be a link component of $\mathscr{C}$ at v in Π. Assume that v has at least 8 incident half-edges in C (and thus at least four intervals), and that Π has a monogon incident to v in C. By construction of Π, the monogon is not labelled 0. Thus, a non-empty subgraph of Γ lies inside the monogon, attached to the rest of Γ only by v, and corresponds to an interval s of Γ at v in C. In Γ, we move the part of Γ that lies inside the monogon on the other side of the edges comprising an adjacent interval s'; see Figure 4.4. This simplifies v by swapping s with s', because Γ has at least 4 intervals at v in C. Note that there is no singular point in the interior of the monogon, because there would be a vertex of Π located on the singular point. The resulting graph embedding Γ' is still proper and cellular, and satisfies the desired properties, by Lemma 4.3.

Now, let us assume that Π contains a sequence of bigons $B_1, \ldots, B_8$ such that B_i and B_{i+1} share an edge for each i. So without loss of generality, we can assume that B_1 and B_5 are labelled 1, B_3 and B_7 are labelled 2, and the other bigons are labelled 0. We modify Γ by exchanging the parts of Γ inside B_3 and B_5; see Figure 4.5. The resulting embedding Γ' is also proper and cellular. This operation simplifies the endpoints u and v of these bigons. (If $u = v$, this is Case 2 from the definition of a simplification.)

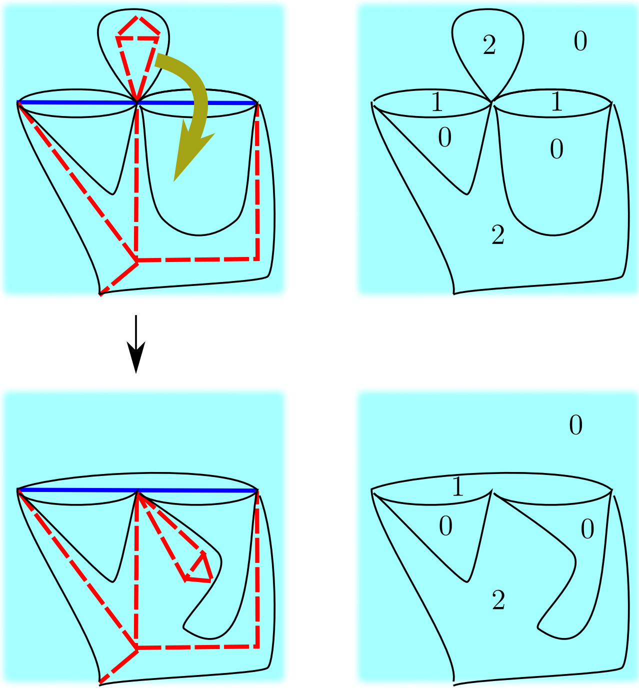

Figure 4.4: Decreasing the number of monogons in Π.

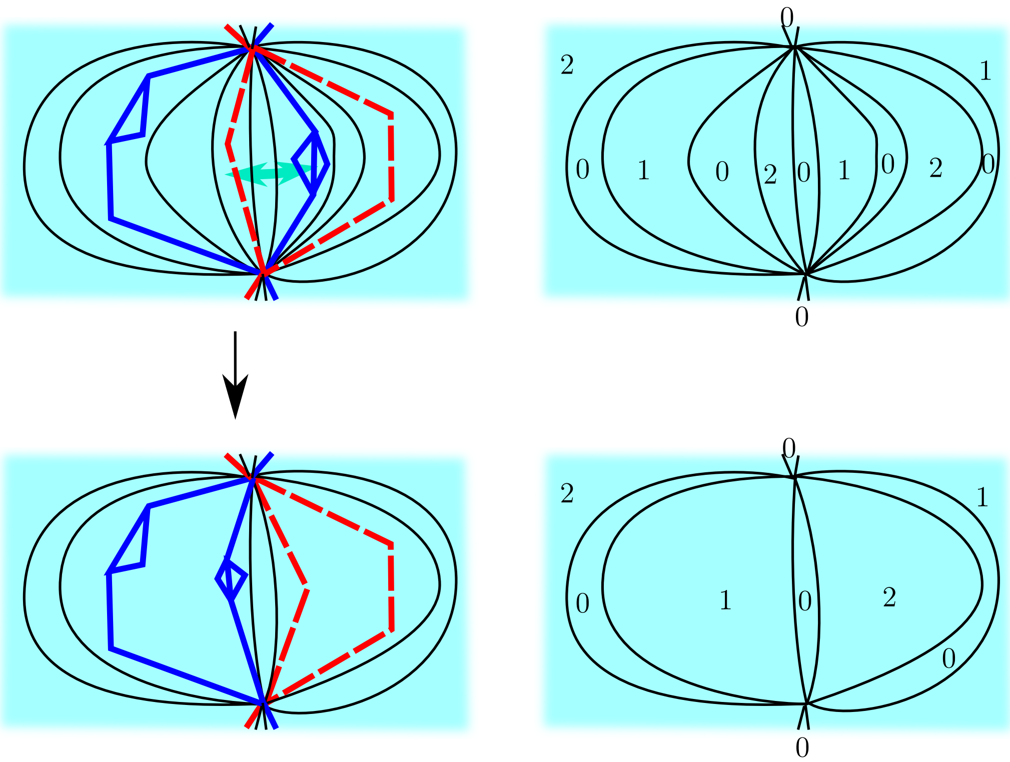

Figure 4.5: Decreasing the number of bigons in Π.

We can iterate the above procedures, but only finitely many times because $|E(\Pi(\Gamma, E_1, E_2))|$ strictly decreases at each step. We have proved that, up to changing our initial embedding Γ, we can assume without loss of generality the following: (i) Let v be a vertex of Π that is incident to a monogon of Π in a link component C; then v is incident to at most 7 half-edges in C, and thus at most 4, since this number is a multiple of 4 by construction; (ii) Π has no sequence of 8 consecutive bigons as above. To conclude, it suffices to prove that any graph Π satisfying these conditions has at most $74c+26w$ edges.

We modify Π by removing all monogons, and then by iteratively replacing each bigon with a single edge, when the edges bounding the bigon are distinct. The removal of monogons does not create any sequence of 8 consecutive bigons, because monogons are attached to vertices of degree at most 4 in their link component. So in the first step, for each vertex v and each link component C of v, at most two monogons in C attached to v are removed; the number of such monogons is at most $2w$ (for the vertices of Π not on a singular point of $\mathscr{C}$) plus $2c$ (for the vertices of Π on singular points of $\mathscr{C}$). In the second step, the number of edges is divided by at most 8. Thus, if Π' denotes the new embedding, we have:

$$|E(\Pi)| \leq 2(c+w) + 8|E(\Pi')|. \tag{1}$$

We now bound the number of edges of Π'. For this purpose, let $\mathscr{S}$ be the detached surface of $\mathscr{C}$, and let Π'' be the graph naturally corresponding to Π' on $\mathscr{S}$ (see Section 3.4). Any bigon of Π'' whose boundary consists of the same edge repeated twice is itself a connected component of $\mathscr{S}$: either a sphere, in which case the corresponding connected component of Π'' is made of two vertices and a single edge, or a projective plane, in which case the corresponding connected component of Π'' is made of a single vertex and a single edge. Thus, in these connected components, the number of edges of Π'' is at most the number of vertices. Let $\mathscr{S}_0$ be obtained from $\mathscr{S}$ by removing these connected components, and Π''_0 the restriction of Π'' to $\mathscr{S}_0$.

Π''_0 has no monogon or bigon. Let $\bar{\mathscr{S}_0}$ be obtained from $\mathscr{S}_0$ by attaching a handle to each boundary component; it has a natural cellular graph embedding with at most $2c$ edges, and thus genus at most $2c$; the graph Π''_0 corresponds to an embedding of a graph $\bar{\Pi}''_0$ on $\bar{\mathscr{S}_0}$, with no monogon or bigon. Lemma 4.2 applied to the restriction of $\bar{\Pi}''_0$ to each connected component of $\bar{\mathscr{S}_0}$ implies that $|E(\bar{\Pi}''_0)| \leq 6c + 3|V(\bar{\Pi}''_0)|$.

Thus, $|E(\Pi')| = |E(\Pi'')| \leq 6c+3|V(\Pi'')|$. Moreover, $|V(\Pi'')| \leq c+w$. Now, Inequality (1) implies that $|E(\Pi)| \leq 2(c+w) + 8(6c+3(c+w)) = 74c+26w$, as desired. □

4.4 Proof of Proposition 4.1

Proof of Proposition 4.1. Let $\mathscr{B}$ be a rooted branch decomposition of G, and let Γ be a proper cellular embedding of G on $\mathscr{C}$. We consider each arc α of the rooted branch decomposition in turn, in an arbitrary order. For each such arc, we modify Γ by applying Lemma 4.4. We first claim that after these iterations, for each arc α of $\mathscr{B}$, we have $|E(\Pi(\Gamma,\alpha))| \leq 74c+26w$.

First, immediately after applying the above procedure to an arc $\tilde{\alpha}$ of $\mathscr{B}$, corresponding to the (ordered) partition $(\tilde{E}_1, \tilde{E}_2)$ of $E(G)$, we have $|E(\Pi(\Gamma, \tilde{E}_1, \tilde{E}_2))| \leq 74c+26w$. We now prove that subsequent applications of Lemma 4.4 to other arcs of the rooted branch decomposition do not increase this number of edges. Indeed, let α be another arc, corresponding to the (ordered) partition (E_1, E_2) of $E(G)$, to which we apply Lemma 4.4. The arc α partitions the nodes of the tree $\mathscr{B}$ into two sets N_1 and N_2, and similarly $\tilde{\alpha}$ partitions the nodes of the tree $\mathscr{B}$ into two sets $\tilde{N}_1$ and $\tilde{N}_2$. Because $\mathscr{B}$ is a tree, we have $\tilde{N}_i \subseteq N_j$ for some $i, j \in \{1,2\}$. This implies that $\hat{\tilde{E}}_i \subseteq \hat{E}_j$ for some $i, j \in \{1,2\}$; so the second item of Lemma 4.4 implies that the number of edges of $\Pi(\Gamma, \tilde{E}_1, \tilde{E}_2)$ does not increase when processing arc α. This proves the claim.

Finally, there remains to prove that, for each internal node ν of $\mathscr{B}$, the graph $\Pi(\Gamma,\nu)$ has at most $3(74c+26w)$ edges. Let (E_1, E_2, E_3) be the edge partition associated with ν. By the claim we have just proved, it suffices to prove that the number of edges of this graph is at most the sums of the numbers of edges of $\Pi(\Gamma, E_1, E_2 \cup E_3)$, $\Pi(\Gamma, E_1 \cup E_2, E_3)$, and $\Pi(\Gamma, E_1 \cup E_3, E_2)$.

Let v and C be as above. We look at the cyclic ordering, around v at C, of the half-edges of $\hat{\Gamma}$ (ignoring the possible boundary component of $\mathscr{C}$ that may arise once in this cyclic ordering). Between two such consecutive half-edges, there are either zero or two half-edges in $\Pi(\Gamma,\nu)$. In the latter case, this means that these two consecutive half-edges of Γ are in two different sets E_i and E_j. Thus, between these two consecutive half-edges of Γ, necessarily two half-edges appear in at least one of $\Pi(\Gamma, E_1, E_2 \cup E_3)$, $\Pi(\Gamma, E_1 \cup E_2, E_3)$, and $\Pi(\Gamma, E_1 \cup E_3, E_2)$. This concludes the proof. □

5 Dynamic programming algorithm

The result of this section is the following proposition.

Proposition 5.1. *Let $\mathscr{C}$ be a 2-complex and G a graph, which satisfy the properties of Proposition 3.1. Let c be the size of $\mathscr{C}$ and n the number of vertices and edges of G. Let $\mathscr{B}$ be a rooted branch decomposition of G of width w. In $(c+w)^{O(c+w)}n$ time, one can report one of the following statements, which is true:*

- *G has no sparse proper cellular embedding into $\mathscr{C}$;*
- *G has an embedding into $\mathscr{C}$.*

(Proposition 4.1 implies that we can remove the adjective "sparse" in the above proposition.)

5.1 Bounding graphs

Let $\mathscr{B}$ be a rooted branch decomposition of G of width w. Recall (see Section 2.1) that the root ρ of $\mathscr{B}$ is a leaf associated to no edge of G. Our algorithm will use dynamic programming in the rooted branch decomposition. For each arc α of $\mathscr{B}$, let G_α be the subgraph of G induced by the edges of G corresponding to the leaves of the subtree of $\mathscr{B}$ rooted at α. The general idea is that we compute all possible relevant embeddings of G_α in subregions of $\mathscr{C}$. Such subregions will be delimited by a graph embedded on $\mathscr{C}$ of small complexity. For the dynamic program to work, we also need to keep track of the location of the vertices in the middle set of α. More precisely, a ***bounding graph*** for G_α is a proper labelled graph embedding Π on $\mathscr{C}$ (but possibly non-cellular), such that:

- some vertices of Π are labelled; these labels are exactly the vertices of the middle set associated with α, and each label appears exactly once;
- each unlabelled vertex of Π is mapped to a singular point of $\mathscr{C}$;
- each face of Π is labelled 0, 1, or 2;
- G_α has an embedding Γ_α that ***respects*** Π: each vertex of Π labelled v is mapped, under Π, to the image of v in Γ_α; moreover, the relative interior of each edge of Γ_α lies in the interior of a face of Π labelled 2.

(It may seem strange to require that each singular point of $\mathscr{C}$ be covered by a vertex of Π; however, it is necessary to cover at least the singular points on the boundary of $\mathscr{C}$, and this more general requirement, albeit slightly artificial, simplifies the argumentation. Also, it is slightly simpler to have three labels for the faces of a bounding graph, although two would suffice.)

A bounding graph for G_α is ***sparse*** if it has at most $74c+26w$ edges. Remark that, if Γ is a sparse proper cellular embedding of G on $\mathscr{C}$ (as defined in Section 4), then $\Pi(\Gamma, \alpha)$ is a sparse bounding graph for the restriction of Γ to G_α.

Henceforth, we regard two (labelled) properly embedded graphs as equal if and only if their (labelled) combinatorial maps are isomorphic. This convenient abuse of language is legitimate because whenever two properly embedded graphs have the same combinatorial map, there is an ambient self-homeomorphism of $\mathscr{C}$ that maps one into the other.

A list $\mathscr{L}_\alpha$ of sparse bounding graphs for G_α is ***exhaustive*** if the following condition holds: If G has a sparse proper cellular embedding on $\mathscr{C}$, then for each such embedding Γ, the (combinatorial map of the) graph $\Pi(\Gamma, \alpha)$ is in $\mathscr{L}_\alpha$.

The induction step for the dynamic programming algorithm is the following.

Proposition 5.2. *Let ν be a non-root node of $\mathscr{B}$ and α be the arc of $\mathscr{B}$ incident to ν that is the closest to the root ρ. Assume that, for each arc $\beta \neq \alpha$ of $\mathscr{B}$ incident to ν, we are given an exhaustive list of sparse bounding graphs for G_β. Then we can, in $(c+w)^{O(c+w)}$ time, compute an exhaustive list of $(c+w)^{O(c+w)}$ sparse bounding graphs for G_α.*

Assuming Proposition 5.2, the proof of which is deferred to the next subsection, it is easy to prove Proposition 5.1:

Proof of Proposition 5.1, assuming Proposition 5.2. We apply the algorithm of Proposition 5.2 in a bottom-up manner in the rooted branch decomposition $\mathscr{B}$. Let α be the arc of $\mathscr{B}$ incident with the root node ρ. We end up with an exhaustive list of sparse bounding graphs for $G_\alpha = G$. By definition of a bounding graph, if this list is non-empty, then G has an embedding on $\mathscr{C}$. On the other hand, by definition of an exhaustive list, if this list is empty, then G has no sparse proper cellular embedding on $\mathscr{C}$.

There are $O(n)$ recursive calls, each of which takes $(c+w)^{O(c+w)}$ time. □

5.2 The induction step: Proof of Proposition 5.2

Proof of Proposition 5.2. First case. Let us first assume that ν is a (non-root) leaf of $\mathscr{B}$; thus, G_α is a single edge uv. We will compute *all* the labelled combinatorial maps of sparse bounding graphs for G_α. It is clear that this will be an exhaustive list. Indeed, assume that G has a sparse proper cellular embedding Γ on $\mathscr{C}$; by sparsity, $\Pi(\Gamma, \alpha)$ has at most $74c + 26w$ edges; thus, one of the labelled combinatorial maps computed will be equal to that of $\Pi(\Gamma, \alpha)$.

So let us describe how to enumerate all the labelled combinatorial maps of bounding graphs for G_α. Using Lemma 3.3, we enumerate all possible labelled (combinatorial maps of) proper graph embeddings Π on $\mathscr{C}$ such that:

- two vertices of Π are labelled u and v; the other vertices are unlabelled; the singular points of $\mathscr{C}$ are covered by the vertices of Π; conversely, every vertex of Π, except perhaps u and/or v, is mapped to singular points of $\mathscr{C}$;
- Π has at most $74c + 26w$ edges;
- each face of Π is labelled 0, 1, or 2;
- Π has a face labelled 2 whose boundary contains both vertices u and v.

It is clear that these labelled combinatorial maps represent all the sparse bounding graphs for G_α.

Second case. Let us now assume that ν is an internal node of $\mathscr{B}$. As above, let α be the arc of $\mathscr{B}$ incident to ν that is the closest to the root ρ. Let β and γ be the arcs different from α incident to ν. Let $\mathscr{L}_\beta$ and $\mathscr{L}_\gamma$ be exhaustive lists of bounding graphs for G_β and G_γ, respectively. Intuitively, every pair of bounding graphs in $\mathscr{L}_\beta$ and $\mathscr{L}_\gamma$ that are compatible, in the sense that the regions labelled 2 in each of these two graphs are disjoint, will lead to a bounding graph in $\mathscr{L}_\alpha$. This is the motivating idea to our approach. More precisely, we will enumerate labelled combinatorial maps Π, each of which can be "restricted" to two compatible graphs, which are possible bounding graphs for G_β and G_γ. If these two restrictions lie in $\mathscr{L}_\alpha$ and $\mathscr{L}_\beta$, this leads to a graph that is added to $\mathscr{L}_\alpha$.

We first introduce some terminology. Let Π be a graph properly embedded on $\mathscr{C}$ (possibly non-cellularly), with faces labelled 0, 1, 2, or 3, and with labels on some vertices. We denote by Π^- the map obtained from Π by replacing each label 3 on a face by a 2. Let i, j, k be integers such that $\{i, j, k\} = \{1, 2, 3\}$. We will define a graph embedding $\Pi_{i,j}$ obtained from Π by somehow "merging" faces i and j. First, for an illustration, refer back to Figures 4.1 and 4.2: If Π is the graph embedding depicted on the right of Figure 4.2, then the configurations shown on the right of Figure 4.1 correspond, from top to bottom, to $\Pi_{2,3}$, $\Pi_{1,3}$, and $(\Pi_{1,2})^-$.

Formally, $\Pi_{i,j}$ is defined as follows. First, let us replace all face labels j by i. Now, for each face f of Π that is homeomorphic to a disk and labelled 0, we do the following. The boundary of f is made of edges of Π; for the sake of the discussion, let us temporarily label each such edge by the label of the face on the other side of f. If all edges on the boundary of f all labelled i, then we remove all these edges, and f becomes part of a larger face labelled i. Otherwise, for each maximal subsequence $e_1, \ldots, e_\ell$ of edges along the boundary of f that are all labelled i, we remove each of $e_1, \ldots, e_\ell$, and replace them with an edge inside f from the source of e_1 to the target of e_ℓ. Finally, we remove all isolated vertices that do not coincide with a singular point of $\mathscr{C}$, and all vertices in the relative interior of an isolated segment that are incident to two faces with the same label.

The easy but key properties of this construction are the following:

(i) Assume that $\Pi_{1,3}$ is a bounding graph for G_β and $(\Pi_{1,2})^-$ is a bounding graph for G_γ. Then $\Pi_{2,3}$ is a bounding graph for G_α.

(ii) The node ν naturally partitions the edge set of G into three parts, which we denote by E_1 (on the side of α), E_2 (on the side of β), and E_3 (on the side of γ). Assume that G has a sparse proper cellular embedding Γ on $\mathscr{C}$ and that $\Pi = \Pi(\Gamma, E_1, E_2, E_3)$. Then:

- $\Pi(\Gamma, \alpha) = \Pi(\Gamma, E_1, E_2 \cup E_3) = \Pi_{2,3}$;
- $\Pi(\Gamma, \beta) = \Pi(\Gamma, E_1 \cup E_3, E_2) = \Pi_{1,3}$;
- $\Pi(\Gamma, \gamma) = \Pi(\Gamma, E_1 \cup E_2, E_3) = (\Pi_{1,2})^-$.

Property (ii) is, again, illustrated by Figures 4.1 and 4.2: If (E_1, E_2, E_3) is the edge partition depicted on Figure 4.2, then the edge partitions depicted on Figure 4.1, left, are, respectively, $(E_1, E_2 \cup E_3)$, $(E_1 \cup E_3, E_2)$, and $(E_1 \cup E_2, E_3)$. As shown above, the corresponding partitioning graphs are respectively $\Pi_{2,3}$, $\Pi_{1,3}$, and $\Pi^-_{1,2}$.

We compute our exhaustive list $\mathscr{L}_\alpha$ of sparse bounding graphs for G_α as follows. Initially, let this list be empty. Using Lemma 3.3, we enumerate all combinatorial maps Π of graphs with at most $c+3w$ vertices and $3(74c+26w)$ edges properly embedded on $\mathscr{C}$ (possibly non-cellularly), with faces labelled 0, 1, 2, or 3, and such that the labels appearing on the vertices are exactly the vertices of the middle set of α, β, or γ (and each label appears exactly once). This takes $(c+w)^{O(c+w)}$ time. Whenever $\Pi_{1,3} \in \mathscr{L}_\beta$ and $(\Pi_{1,2})^- \in \mathscr{L}_\gamma$, we add $\Pi_{2,3}$ to $\mathscr{L}_\alpha$. Finally, we eliminate duplicates by testing pairwise isomorphism between the labelled combinatorial maps in $\mathscr{L}_\alpha$, and remove the graphs that are not sparse or contain vertices that bear a label not in the middle set of α.

$\mathscr{L}_\alpha$ contains only sparse bounding graphs for G_α, by (i) above. Moreover, let Γ be a sparse proper cellular graph embedding of G on $\mathscr{C}$. By sparsity, one of the graphs Π enumerated in the previous paragraph is $\Pi(\Gamma, \nu)$. By definition of $\mathscr{L}_\beta$ and $\mathscr{L}_\gamma$, we have that $\Pi(\Gamma, \beta) \in \mathscr{L}_\beta$ and $\Pi(\Gamma, \gamma) \in \mathscr{L}_\gamma$, so by (ii) above, $\Pi(\Gamma, \alpha) \in \mathscr{L}_\alpha$, which implies that $\mathscr{L}_\alpha$ is exhaustive. $\square$

6 Reduction to proper cellular embeddings

This section is devoted to proving the following result:

Proposition 6.1. *Let $\mathscr{C}$ be a 2-complex with at most c simplices, and G a graph with at most n vertices and edges and branchwidth at most w. Assume that G and $\mathscr{C}$ satisfy the properties of Proposition 3.1. In $c^{O(c)} + O(cn)$ time, one can compute a graph G', and $c^{O(c)}$ 2-complexes $\mathscr{C}_i$, such that:*

1. *each $\mathscr{C}_i$ and G' satisfy the properties of Proposition 3.1;*
2. *G' has at most $5cn$ vertices and $5cn$ edges, and branchwidth at most w;*
3. *each $\mathscr{C}_i$ has size at most c;*
4. *if, for some i, G' embeds into $\mathscr{C}_i$, then G embeds into $\mathscr{C}$;*
5. *if G embeds into $\mathscr{C}$, then for some i, G' has a proper cellular embedding into $\mathscr{C}_i$.*

We start with auxiliary results. Let $\mathscr{S}$ be a surface (possibly disconnected, possibly with boundary). A ***cutting operation*** on $\mathscr{S}$ consists of cutting it along a simple closed curve, and attaching a disk to the resulting boundary component(s). A cutting operation is ***essential*** if the simple closed curve is non-contractible.

The following result is not hard and essentially folklore (a related but slightly weaker result is provided by Matoušek et al. [MSTW16, Lemma 3.1]), but we could not find a precise reference.

Lemma 6.2. *Let $\mathscr{S}$ be a (connected) surface with genus g. The number of possibly disconnected surfaces, up to homeomorphism, that can be obtained from $\mathscr{S}$ by a cutting operation is at most $g+3$, and we can compute them in linear time. Moreover, this cutting operation leads either to a single surface with genus strictly smaller than g, or to two surfaces, the sum of the genera of which equals g, and the size of the surface (sum of the number of connected components, total genus, and number of boundary components) increases by at most one.*

Proof. This basically follows from the classification of surfaces together with Euler's formula. A cutting operation of $\mathscr{S}$ along a closed curve γ falls into exactly one of the following three categories:

1. *Case where γ is separating.* The cutting operation on $\mathscr{S}$ results in two surfaces $\mathscr{S}_1$ and $\mathscr{S}_2$, in which their respective genera g_1 and g_2 satisfy $g = g_1 + g_2$. Moreover, if $\mathscr{S}$ is non-orientable, then at least one of $\mathscr{S}_1$ or $\mathscr{S}_2$ is non-orientable. Finally, all pairs of surfaces $(\mathscr{S}_1, \mathscr{S}_2)$ satisfying these constraints can be obtained as the result of a cutting operation on $\mathscr{S}$.

2. *Case where γ is non-separating but two-sided.* This is only possible if $g \geq 2$. The cutting operation on $\mathscr{S}$ results in a single surface $\mathscr{S}'$ with genus $g-2$. If $\mathscr{S}$ is orientable, then so is $\mathscr{S}'$; otherwise, $\mathscr{S}'$ is either orientable or non-orientable (unless of course $g = 2$, in which case it is necessarily orientable, or g is odd, in which case it is necessarily non-orientable). All surfaces $\mathscr{S}'$ satisfying these constraints can be obtained.

3. *Case where γ is one-sided.* This is only possible if $\mathscr{S}$ is non-orientable and $g \geq 1$. The cutting operation on $\mathscr{S}$ results in a single surface $\mathscr{S}'$ with genus $g-1$, orientable or not (unless of course $g = 1$, in which case it is orientable, or g is even, in which case it is non-orientable). All surfaces $\mathscr{S}'$ satisfying these constraints can be obtained. ☐

Lemma 6.3. *Let $\mathscr{S}$ be a surface with k connected components, total genus g, and with b boundary components in total. In $(k+g+b)^{O(g+b)}$ time, we can enumerate all $(k+g+b)^{O(g+b)}$ possibly disconnected surfaces with boundary, up to homeomorphism, arising from $\mathscr{S}$ by one or several successive* essential *cutting operations. These surfaces have $O(k+g+b)$ connected components and size $O(k+g+b)$.*

Proof. It is useful to organize the set of all surfaces (possibly disconnected, possibly with boundary) arising by essential cutting operations in a tree with root $\mathscr{S}$, in which the children of a node result from a single essential cutting operation. We prove that (1) the depth of the tree is $O(g+b)$ and that (2) each node of the tree has $O((k+1)(g+1)(b+1))$ children, which concludes (because by Lemma 6.2, the size of a surface increases by at most one by a cutting operation).

Let $\mathscr{S}'$ be a (possibly disconnected, possibly with boundary) surface resulting from a sequence of essential cutting operations on $\mathscr{S}$. By Lemma 6.2, the total genus of $\mathscr{S}'$ is at most g. Moreover, since we consider only essential cutting operations, each connected component of $\mathscr{S}'$ either has positive genus or contains at least one boundary component, unless it was itself a connected component of $\mathscr{S}$; so the number of connected components of $\mathscr{S}'$ is at most $k+g+b$.

Let $\varphi(\mathscr{S}')$ be equal to twice the total genus of $\mathscr{S}'$ minus its number of connected components. By Lemma 6.2, this potential function strictly decreases at each cutting operation. Moreover, we have $\varphi(\mathscr{S}) = 2g - k$, and by the previous paragraph $\varphi(\mathscr{S}')$ is at least $-(k+g+b)$. This proves (1).

By Lemma 6.2, for any surface of genus g without boundary, there are at most $g+3$ ways of performing a cutting operation up to homeomorphism. After a sequence of essential cutting operations, we have a surface $\mathscr{S}'$ with at most $k+g+b$ connected components, with total genus at most g, and with b boundary components. The number of surfaces that can be obtained from $\mathscr{S}'$ by a cutting operation is at most $k(g+3)(b+1)$, since we first choose which connected component to cut along, the way to cut it ignoring the boundary components, and the number of boundary components in each connected component (if the cut is separating). ☐

We can now conclude the proof of this section.

Proof of Proposition 6.1. First, if the detached surface of $\mathscr{C}$ is non-empty, then we test the planarity of each connected component of G [HT74] in $O(n)$ time, and remove every connected component of G that is planar; obviously, this does not affect the embeddability of G on $\mathscr{C}$.

In a second step, we split each isolated segment of $\mathscr{C}$ into five isolated segments. For each subset S of the isolated segments, we build a new 2-complex obtained from $\mathscr{C}$ by removing S. We obtain $2^{O(c)}$ 2-complexes, each of size $O(c)$. The input graph G embeds on $\mathscr{C}$ if and only if it embeds into one of these 2-complexes; moreover, if G embeds on $\mathscr{C}$, it embeds into one of the 2-complexes in a way that every isolated segment is covered by the embedding. (Indeed, remember that G has at most one connected component homeomorphic to a segment.)

We now iteratively dissolve every degree-two vertex of G, and then subdivide $5c$ times each edge of G. This new graph G' has at most $5cn$ vertices and edges, and branchwidth at most w. Clearly, G embeds on $\mathscr{C}$ if and only if G' embeds in one of the 2-complexes defined in the previous paragraph; moreover, if G embeds on $\mathscr{C}$, then G' has an embedding on one of these 2-complexes in which the relative interior of every edge is distinct from any singular point (and, as above, such that every isolated segment is covered by the embedding).

Each singular point p of each of these 2-complexes is incident to at least two link components. For each such singular point p and for each partition of the link components at p, we replace p with new vertices, one for each element in the partition; two link components at p stay adjacent via one of these new vertices if and only if these link components are in the same part. We obtain $c^{O(c)}$ 2-complexes, each of size $O(c)$. The input graph G embeds on $\mathscr{C}$ if and only if G' embeds in one of these 2-complexes; moreover, if G embeds on $\mathscr{C}$, then G' has an embedding into one of these 2-complexes in which every link component of each singular point p is used by an edge of G connected to p in that link component (and, as above, such that every isolated segment is covered by the embedding, and such that the relative interior of every edge is distinct from any singular point).

Every embedding of a graph into a 2-complex can be perturbed so that it avoids the boundary of the 2-complex, except possibly at singular points. This means that, if G embeds on $\mathscr{C}$, then G' has a proper cellular embedding into one of the 2-complexes built in the previous paragraph, except that the faces of G' may fail to be disks, but are more general (connected) surfaces with boundary.

To dispense ourselves from this latter exception, we need to build more 2-complexes. This case occurs only if the detached surface is non-empty, so by our earlier preprocessing, we can assume that G' contains no planar connected component, and so has $O(c)$ connected components (because, by Lemma 3.4, in order for G' to be embeddable on a 2-complex of size $O(c)$, it must have genus $O(c)$).

The detached surface $\mathscr{S}$ is a surface (possibly disconnected, possibly with boundary); the trace of the set of singular points on $\mathscr{S}$ corresponds to marked points, some in the interior of $\mathscr{S}$, some on the boundary. Henceforth, we regard the former as small boundary components. For each 2-complex obtained above, we consider, up to homeomorphism, all 2-complexes arising from zero, one, or several essential cutting operations on $\mathscr{S}$, and then by removing an arbitrary subset of the connected components of the resulting surface. Up to homeomorphism, by Lemma 6.3, there are $c^{O(c)}$ ways of cutting $\mathscr{S}$; since we consider all 2-complexes obtained up to homeomorphism, we need to consider each boundary component of the detached surface $\mathscr{S}$ as labelled (which is not the case in Lemma 6.3); however, this only adds a factor of $c^{O(c)}$. In total, we obtain, in $c^{O(c)}$ time, $c^{O(c)}$ 2-complexes, each of size $O(c)$. Then, for each such 2-complex, we consider all possible ways of removing an arbitrary subset of connected components of the 2-complex; the number of the resulting 2-complexes is still $c^{O(c)}$. By construction, the input graph G embeds on $\mathscr{C}$ if and only if G' embeds in one of these 2-complexes. Moreover, assume that it is the case; as shown above, G' has a proper embedding into one of the 2-complexes in the previous paragraph, except that faces of G' are (connected) surfaces, not necessarily disks. Whenever a face has non-empty boundary and is not homeomorphic to a disk, we perform an essential cutting operation of that face along a closed curve inside that face; the closed curve along which we cut is non-contractible in $\mathscr{S}$, because otherwise it would bound a disk in $\mathscr{S}$, which would itself contain a planar connected component of G', and we have shown above that we may assume that no such component exists. After iterating this operation as much as possible, every face of G' in the resulting 2-complex is either is a disk or has empty boundary; in the latter case, G' avoids the corresponding connected component, so we can simply remove it. Eventually, after a number of essential cutting operations of $\mathscr{S}$ and removing some connected components of the 2-complex, the embedding of G' is cellular in one of the $c^{O(c)}$ 2-complexes of size $O(c)$ enumerated above. □

7 Algorithm for bounded branchwidth: Proof of Theorem 1.2

Proof of Theorem 1.2. By Proposition 3.1, we can assume that $\mathscr{C}$ has no 3-book and no connected component that is reduced to a single vertex, and that G has no connected component reduced to a single vertex and at most one connected component homeomorphic to a segment. If necessary, we convert the combinatorial description of $\mathscr{C}$ into the topological data structure (Section 3.3).

We apply Proposition 6.1. In $c^{O(c)} + O(cn)$ time, we obtain a graph G' and a set of $c^{O(c)}$ 2-complexes $\mathscr{C}_i$ such that:

1. each $\mathscr{C}_i$ and G' satisfy the properties of Proposition 3.1;
2. G' has at most $5cn$ vertices and $5cn$ edges, and branchwidth at most w;
3. each $\mathscr{C}_i$ has size at most c;
4. if, for some i, G' embeds into $\mathscr{C}_i$, then G embeds into $\mathscr{C}$;
5. if G embeds into $\mathscr{C}$, then for some i, G' has a proper cellular embedding into $\mathscr{C}_i$.

We then run the algorithm from Proposition 5.1 in each of the instances $(\mathscr{C}_i, G')$, in total time $(c+w)^{O(c+w)}n$. This algorithm correctly reports either that G' has no sparse proper cellular embedding into $\mathscr{C}_i$ or that G' has an embedding into $\mathscr{C}$. If for at least one of these instances, the algorithm reports that G' embeds into $\mathscr{C}_i$, then we report that G embeds into $\mathscr{C}$. Otherwise, we report that G does not embed into $\mathscr{C}$.

There remains to prove that the algorithm is correct. If our algorithm reports that G embeds into $\mathscr{C}$, then it is obviously indeed the case (Property (4) above). Conversely, let us assume that G has an embedding into $\mathscr{C}$. Thus, by Property (5) above, let i be such that G' has a proper cellular embedding into $\mathscr{C}_i$. By Proposition 4.1, G' also has such an embedding into $\mathscr{C}_i$ that is sparse. Thus, the algorithm in Proposition 5.1 (correctly) reports that G' has an embedding into $\mathscr{C}_i$, and finally our overall algorithm reports that G has an embedding into $\mathscr{C}$. □

8 Reduction to bounded branchwidth: Proof of Theorem 1.3

This section is devoted to the proof of Theorem 1.3. The deterministic and randomized algorithms differ only in the algorithm that we use as a subroutine to compute a large grid minor. The proof technique is based on an irrelevant vertex method; we borrow ingredients to Kociumaka and Pilipczuk [KP19, Section 5], but some new arguments are needed, in particular in the beginning of the proof of Proposition 8.1.

8.1 Finding a large planar part

A ***wall*** of size $k \times k$ is a subgraph of the $(k \times k)$-grid obtained by removing alternately the vertical edges of even (resp. odd) x-coordinate in each even (resp. odd) line, and then the degree-one vertices; see Figure 8.1.

As an intermediate goal towards the proof of Theorem 1.3, we will prove in this subsection:

Proposition 8.1. *Let G be a graph with n vertices and edges and $g \geq 2$ be an integer. We can do one of the following:*

1. *compute a rooted branch decomposition of G of width $g^{O(1)}$;*
2. *correctly report that G has genus at least g;*

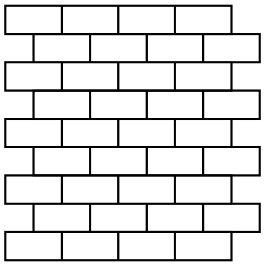

Figure 8.1: A (10×10)-wall.

3. *or compute a cycle γ of G such that one connected component of $G - \gamma$ is planar and contains a subdivision of the $(g \times g)$-wall, which is also computed.*

either in deterministic $f(c) \cdot n^2$ time for some computable function f, or in expected polynomial time.

We will use the following lemma.

Lemma 8.2. *Let G be a graph with n vertices and edges, and $k \geq 2$. Then, one can compute either a rooted branch decomposition of G of width $k^{O(1)}$ or a $(k \times k)$-grid minor of G, in deterministic $f(k) \cdot n^2$ time, where f is a computable function, or in expected polynomial time.*

Proof. We can obviously assume that k is at most the number of vertices of G, for otherwise any rooted branch decomposition of G has width at most k. Let d be a large enough universal constant, which can be computed from the hidden constants in the $O(\cdot)$ notation of the results used below.

We first use any algorithm to approximate the treewidth, e.g., Fomin et al. [FLS+18, Theorem 1.1]: Provided d is large enough, in polynomial time, we either compute a tree decomposition of width at most $(dk^d)^2$, and thus immediately obtain a (rooted) branch decomposition of width $(dk^d)^2$ [RS91, Theorem 5.1], as desired, or correctly report that the treewidth is at least dk^d.

In the latter case, provided d is large enough, G has a $(k \times k)$-grid minor, by a result by Chekuri and Chuzhoy [CC16]. We can compute it in expected polynomial time, by an algorithm from the same article [CC16]. Alternatively, we can compute it in deterministic $f(k) \cdot n^2$ time, for some computable function f, by a result by Robertson and Seymour [RS95, Algorithm 4.4]. □

We remark that the above proof contains some bottlenecks in the running time in Theorems 1.3 and 1.1. Specifically:

- the randomness is solely due to the above use of the algorithm in the article by Chekuri and Chuzhoy [CC16];
- the fact that the dependence in the size of the 2-complex is not specified is solely due to the use of Robertson and Seymour [RS95, Algorithm 4.4] in computing a grid minor, but it can in principle be made explicit.

Proof of Proposition 8.1. We can again obviously assume that g is at most the number of vertices of G. We apply Lemma 8.2 with $k = 60\lceil\sqrt{g}\rceil^4$. If the outcome is a rooted branch decomposition, then the algorithm returns it (Case 1). Otherwise, we have computed a $(60\lceil\sqrt{g}\rceil^4 \times 60\lceil\sqrt{g}\rceil^4)$-grid minor of G, and thus a subgraph $\dot{W}$ of G that is a subdivision of a $(60\lceil\sqrt{g}\rceil^4 \times 60\lceil\sqrt{g}\rceil^4)$-wall W.

We first compute disjoint non-adjacent $(50g\lceil\sqrt{g}\rceil \times 50g\lceil\sqrt{g}\rceil)$-walls $W_1, \ldots, W_g$ of W, and the corresponding subdivisions $\dot{W}_1, \ldots, \dot{W}_g$ that are subgraphs of $\dot{W}$, in a way that $W - W_i$ is connected for each i. For each i, we consider the subgraph G_i of G induced by the vertices v of G such that: (1) there exists a path from v to $\dot{W}_i$; (2) every path from v to $\dot{W}_j$, for some $j \neq i$, uses at least one vertex from $\dot{W}_i$. The graphs G_i are pairwise disjoint. We test the planarity of each of them in linear time [HT74]. If all of them are non-planar, then G has genus at least g, so we correctly report this (Case 2). So without loss of generality, one of these graphs, say G_1, is planar, and our algorithm computes it.

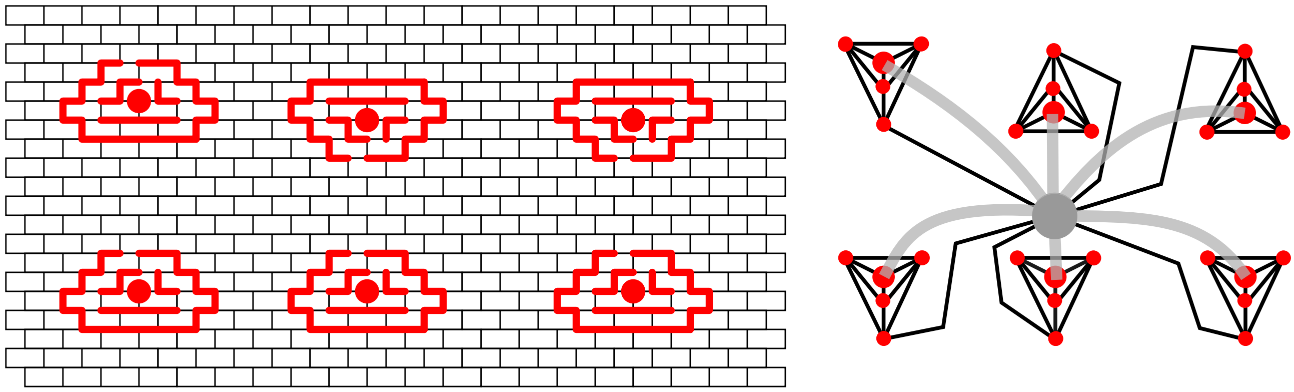

Figure 8.2: Illustration of the proof of the claim in Proposition 8.1. The wall W_1 (left), and a schematic view of the J_p minor (right), illustrated with $g = 6$. Let K_5^- be the graph K_5 with one edge removed. If W_1 has many vertices connected to the outside, then a set U of at least g vertices connected to the outside (represented as big disks on the left) are pairwise distant. In the neighborhood of these vertices, we build a K_5^- minor for each vertex in U, in which the vertex of U is the "central" vertex (the thick paths need to be contracted to obtain copies of K_5^-). Using the fact that the neighborhoods are far apart, we build a J_g-minor of G (right), in which the apex results from the contraction of the subdivided wall $\dot{W}_1$ minus the union of the K_5^- minors, and each path connecting $u \in U$ to $\dot{W} - \dot{W}_1$ (which is connected by assumption) is contracted, except its first edge, to connect u to the apex.

By 3-connectivity, W_1, and thus also $\dot{W}_1$, has a unique combinatorial embedding in the plane, up to symmetry and up to the choice of the outer (infinite) face; we consider the natural embedding of W_1 in which the outer face has the largest degree. An *inner vertex* of W_1 is one that is at distance at least 6 from the outer face in the natural embedding of W_1. Remark that each vertex of W_1 is a vertex of $\dot{W}_1$. We say that a vertex u of W_1 is *connected to the outside* if there is, in G, a path whose vertices are, in this order, u, possibly some vertices of $\dot{W}_1$ but not of W_1, possibly some vertices of $G - \dot{W}_1$, and finally one vertex in $\dot{W} - \dot{W}_1$.

We claim the following: *If at least* $1000g$ *inner vertices of* W_1 *are connected to the outside, then* G *has genus at least* g. The strategy is similar to the argument in Kociumaka and Pilipczuk [KP19, Lemma 5.3]; we summarize the proof. If at least $1000g$ inner vertices of W_1 are connected to the outside, then a set U of g inner vertices of W_1 are connected to the outside, and at pairwise distance at least 16 in W_1. This implies that G contains, as a minor, the graph J_g obtained from g copies of K_5 by subdividing an edge from each copy with a degree-two vertex and identifying these g new vertices into a single vertex, the *apex* of J_g; see Figure 8.2. This graph has genus at least g, by the main result of an article by Miller [Mil87, Theorem 1]. This proves the claim.

In polynomial time, we can compute the inner vertices of W_1 connected to the outside. If there are at least $1000g$ of these, we report that G has genus at least g (Case 2), which is correct by the above claim. Otherwise, from the $(50g\lceil\sqrt{g}\rceil \times 50g\lceil\sqrt{g}\rceil)$-wall W_1, we can compute a cycle γ in W_1, enclosing a $(g \times g)$-wall W_1' in W_1 (in the natural embedding of W_1) such that no vertex of W_1 inside γ (in the natural embedding of W_1) is connected to the outside. Let $\dot{\gamma}$ be the cycle of $\dot{W}_1$ corresponding to γ, and let H be the connected component of $G - \dot{\gamma}$ containing the vertices of this $(g \times g)$-wall. There remains to prove that H is planar, and since G_1 is planar, it suffices to prove that H is a subgraph of G_1. If it were not the case, H would contain a vertex of $\dot{W}_j$ for some $j \neq 1$; but that would imply that a vertex of W_1' is connected to the outside, which is not the case. We can thus correctly report γ (Case 3). □

8.2 Finding an irrelevant vertex

The following proposition will imply that if the third possibility in the statement of Proposition 8.1 holds (for some g large enough), then one has an *irrelevant vertex* for the embedding instance.

Proposition 8.3. *Let $\mathscr{C}$ be a 2-complex with $c \geq 1$ simplices. Let G be a graph and γ be a cycle in G*

such that one connected component of $G-\gamma$ is planar and contains a subdivision of the $(1000c\times 1000c)$-wall. Let v be the central vertex of this wall. Then G is embeddable on $\mathscr{C}$ if and only if $G-v$ is.

The proof of Proposition 8.3 builds upon the article by Kociumaka and Pilipczuk [KP19, Section 5.3]: The strategy is the same, up to simple variations to take into account the fact that we consider a 2-complex, not a surface. Rather than repeating the same arguments, we summarize the proof, emphasizing the differences.

A ***circular wall*** W [KP19, Figure 8] of height h and circumference ℓ is a 3-regular graph that, in some embedding of W in the plane, is represented as the union of h vertex-disjoint cycles, called ***circles***, organized in a concentric way, such that any two consecutive circles are connected by ℓ ***radial edges***; in successive layers, the radial edges are interleaved. A ***ring*** R of this circular wall W is a subgraph contained in the (closed) annulus between two circles at distance four in the radial order; a ring thus contains five vertex-disjoint circles, consecutive in this concentric order, together with edges connecting them; the ***central circle*** of R is the third one in that order. A ***central brick*** of R is the boundary of a face of W incident to its central circle.

If G is embeddable on $\mathscr{C}$, then obviously $G-v$ is also embeddable on $\mathscr{C}$; the hard part is the reverse direction. In the given subdivision of the $(1000c\times 1000c)$-wall in G, we first compute a subdivision $\dot{W}$ of a circular wall W of height $45c+10$ and circumference three, so that v is located inside the inner circle of this circular wall in the planar embedding of W; see Kociumaka and Pilipczuk [KP19, Figures 8 and 9].

Let us consider an embedding ψ of $G-v$ on $\mathscr{C}$. This induces embeddings of $\dot{W}-v$ and $W-v$ on $\mathscr{C}$, also denoted by ψ. Each ring of W corresponds naturally to a subdivided ring of $\dot{W}$. Such a subdivided ring is ***embedded plainly*** in ψ if it lies in the detached surface $\mathscr{S}$ of $\mathscr{C}$ and each central brick of the ring is a two-sided closed curve in $\mathscr{S}$ and bounds a disk on $\mathscr{S}$, the interior of which does not contain any vertex or edge of $\dot{W}$.

We have the following lemma.

Lemma 8.4. *Some subdivided ring of $\dot{W}$ is embedded plainly in ψ.*

Proof. At most c vertices or edges of W are mapped, under ψ, to a singular point of $\mathscr{C}$. At most c edges of W are mapped, under ψ, to an isolated segment of $\mathscr{C}$. Our circular wall of height $5(9c+2)$ and circumference 3 contains $9c+2$ vertex-disjoint rings of circumference 3. Thus, there exists a set $\mathcal{R}$ of $7c+2$ vertex-disjoint rings of W of circumference 3, the vertices and edges of which are mapped, under ψ, to the complement of the singular points and the isolated segments of $\mathscr{C}$.

In other words, ψ maps each ring of $\mathcal{R}$ to the detached surface $\mathscr{S}$ of $\mathscr{C}$, and thus also to the surface $\bar{\mathscr{S}}$ obtained by attaching a handle to every boundary component of $\mathscr{S}$. The resulting surface has a natural cellular graph embedding with at most $7c$ edges, and thus has genus at most $7c$. Since $\mathcal{R}$ contains $7c+2$ rings, this implies that some ring in $\mathcal{R}$ is embedded plainly, by ψ, in $\bar{\mathscr{S}}$; see the proof of Kociumaka and Pilipczuk [KP19, Lemma 5.4] for details. Hence it is also embedded plainly, by ψ, in $\mathscr{C}$. ☐

The rest of the proof of Proposition 8.3 uses the same arguments as in Kociumaka and Pilipczuk [KP19], so we only summarize it:

Proof of Proposition 8.3. Let R be the ring obtained from Lemma 8.4. Its central cycle γ is mapped, under ψ, to a two-sided cycle in the detached surface $\mathscr{S}$, and moreover, the edges of R incident to γ are partitioned, under ψ, between the sides of γ exactly as in the natural embedding [KP19, Corollary 5.5]. Since R is embedded plainly in the detached surface, one can embed the (planar) part of G that lies inside R in a neighborhood of γ, and then insert v without interfering with the rest of the embedding. ☐

8.3 Proof of Theorem 1.3

Proof of Theorem 1.3. We first apply Proposition 3.1: without loss of generality, $\mathscr{C}$ has no 3-book and no connected component that is reduced to a single vertex, and G has no connected component reduced to a single vertex, and at most one connected component homeomorphic to a segment. Let n be the number of vertices and edges of the input graph G, and c be the number of simplices of $\mathscr{C}$. We apply Proposition 8.1 to the graph G, letting $g = 1000c$. In deterministic $f(c) \cdot n^2$ time for some computable function f, or in expected time polynomial in n and c, we obtain one of the following outcomes:

1. a rooted branch decomposition of G of width $O(g)^{O(1)}$;
2. that G has genus at least g, and is thus not embeddable on $\mathscr{C}$ (Lemma 3.4);
3. a cycle γ of G such that one connected component of $G - \gamma$ is planar and contains a subdivision of the $(g \times g)$-wall.

In the first two cases, we are done. In the third case, by applying Proposition 8.3, we obtain a vertex v such that G embeds on $\mathscr{C}$ if and only if $G - v$ does. By iterating the same procedure a number of times that is at most the number of vertices of G, we necessarily reach case (1) or (2), which concludes. □

We remark that the proof goes through if the input 2-complex is given in the form of the topological data structure, and c denotes its size, instead of the number of simplices of $\mathscr{C}$.

As mentioned above, the proof of Theorem 1.1 follows immediately from Theorems 1.2 and 1.3.

Acknowledgment

We would like to thank Arnaud de Mesmay for useful discussions.

References

[AGK08] Isolde Adler, Martin Grohe, and Stephan Kreutzer. Computing excluded minors. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 641–650, 2008.

[BCG+06] Christoph Buchheim, Markus Chimani, Carsten Gutwenger, Michael Jünger, and Petra Mutzel. Crossings and planarization. In Roberto Tamassia, editor, *Handbook of graph drawing and visualization*. Chapman and Hall, 2006.

[CC16] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM*, 63(5):Article 40, 2016.

[CCdVE+08] Erin W. Chambers, Éric Colin de Verdière, Jeff Erickson, Francis Lazarus, and Kim Whittlesey. Splitting (complicated) surfaces is hard. *Computational Geometry: Theory and Applications*, 41(1–2):94–110, 2008.

[CdV18] Éric Colin de Verdière. Computational topology of graphs on surfaces. In Jacob E. Goodman, Joseph O'Rourke, and Csaba Toth, editors, *Handbook of Discrete and Computational Geometry*, chapter 23. CRC Press LLC, third edition, 2018.

[CdVdM14] Éric Colin de Verdière and Arnaud de Mesmay. Testing graph isotopy on surfaces. *Discrete & Computational Geometry*, 51(1):171–206, 2014.

[CdVMM18] Éric Colin de Verdière, Thomas Magnard, and Bojan Mohar. Embedding graphs into two-dimensional simplicial complexes. In *Proceedings of the 34th International Symposium on Computational Geometry (SOCG)*, pages 27:1–27:14, 2018.

[CFK+15] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*. Springer-Verlag, 2015.

[DG99] Tamal K. Dey and Sumanta Guha. Transforming curves on surfaces. *Journal of Computer and System Sciences*, 58:297–325, 1999.

[Epp03] David Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 599–608, 2003.

[EW13] Jeff Erickson and Kim Whittlesey. Transforming curves on surfaces redux. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1646–1655, 2013.

[FLS+18] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, Michał Pilipczuk, and Marcin Wrochna. Fully polynomial-time parameterized computations for graphs and matrices of low treewidth. *ACM Transactions on Algorithms*, 14(3):Article 34, 2018.

[HT74] John Hopcroft and Robert Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974.

[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 *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 771–780, 2008.

[KP19] Tomasz Kociumaka and Marcin Pilipczuk. Deleting vertices to graphs of bounded genus. *Algorithmica*, 81:3655–3691, 2019.

[KR07] Ken-ichi Kawarabayashi and Bruce Reed. Computing crossing number in linear time. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 382–390, 2007.

[LR12] Francis Lazarus and Julien Rivaud. On the homotopy test on surfaces. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 440–449, 2012.

[Mal94] Seth M. Malitz. Genus g graphs have pagenumber $O(\sqrt{g})$. *Journal of Algorithms*, 17:85–109, 1994.

[Mil87] Gary L Miller. An additivity theorem for the genus of a graph. *Journal of Combinatorial Theory, Series B*, 43(1):25–47, 1987.

[Moh99] Bojan Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM Journal on Discrete Mathematics*, 12(1):6–26, 1999.

[MSTW16] Jiří Matoušek, Eric Sedgwick, Martin Tancer, and Uli Wagner. Untangling two systems of noncrossing curves. *Israel Journal of Mathematics*, 212:37–79, 2016.

[MT01] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces.* Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 2001.

[ÓDWW00] Colm Ó Dúnlaing, Colum Watt, and David Wilkins. Homeomorphism of 2-complexes is equivalent to graph isomorphism. *International Journal of Computational Geometry & Applications*, 10:453–476, 2000.

[Pat06] Maurizio Patrignani. Planarity testing and embedding. In Roberto Tamassia, editor, *Handbook of graph drawing and visualization.* Chapman and Hall, 2006.

[RS91] Neil Robertson and Paul D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52:153–190, 1991.

[RS95] Neil Robertson and Paul D Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of combinatorial theory, Series B*, 63(1):65–110, 1995.

[RS04] Neil Robertson and Paul D. Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92:325–357, 2004.

[Sch13] Marcus Schaefer. Toward a theory of planarity: Hanani–Tutte and planarity variants. *Journal of Graph Algorithms and Applications*, 17(4):367–440, 2013.

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

[Tho89] Carsten Thomassen. The graph genus problem is NP-complete. *Journal of Algorithms*, 10(4):568–576, 1989.