



Embedding graphs into two-dimensional simplicial complexes

Éric Colin de Verdière, Thomas Magnard, Bojan Mohar

► To cite this version:

Éric Colin de Verdière, Thomas Magnard, Bojan Mohar. Embedding graphs into two-dimensional simplicial complexes. *Computing in Geometry and Topology*, 2022, 1 (1), pp.1-23. 10.57717/cgt.v1i1.11 . hal-04269443

HAL Id: hal-04269443

<https://hal.science/hal-04269443>

Submitted on 3 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Embedding Graphs into Two-Dimensional Simplicial Complexes

Éric Colin de Verdière 

LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France

Thomas Magnard 

LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France

Bojan Mohar  

Department of Mathematics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada

Abstract

We consider the problem of deciding whether an input graph G admits a topological embedding into an input two-dimensional simplicial complex \mathcal{C} . This problem includes, among others, the embeddability problem of a graph on a surface and the topological crossing number of a graph, but is more general.

The problem is NP-complete in general (if \mathcal{C} is part of the input), and we give an algorithm that runs in polynomial time for any fixed \mathcal{C} .

Our strategy is to reduce the problem into an embedding extension problem on a surface, which has the following form: Given a subgraph H' of a graph G' , and an embedding of H' on a surface S , can that embedding be extended to an embedding of G' on S ? Such problems can be solved, in turn, using a key component in Mohar's algorithm to decide the embeddability of a graph on a fixed surface (STOC 1996, SIAM J. Discr. Math. 1999).

Keywords and phrases computational topology, embedding, simplicial complex, graph, surface

Digital Object Identifier 10.57717/cgt.v1i1.11

Related Version A preliminary version appeared in the *34th Proceedings of the International Symposium on Computational Geometry*, pp.27:1–27:14, 2018.

Funding Part of this work was done while the first two authors were at Département d'Informatique, École normale supérieure, Paris, France, and while the third author was invited professor there. The first two authors are supported by grants ANR-17-CE40-0033 (SoS) and ANR-19-CE40-0014 (MIN-MAX) of the French National Research Agency ANR. The third author was supported in part by the NSERC Discovery Grant R611450 (Canada), by the Canada Research Chairs program, and by the Research Project J1-8130 of ARRS (Slovenia).

1 Introduction

Topological embedding problems. Topological embedding problems are among the most fundamental problems in computational topology, already emphasized since the early developments of this discipline [9, Section 10]. Their general form is as follows: Given topological spaces X and Y , does there exist an embedding (a continuous, injective map) from X to Y ? Since a finite description of X and Y is needed, typically they are represented as finite simplicial complexes, which are topological spaces obtained by attaching simplices (points, segments, triangles, tetrahedra, etc.) of various dimensions together.

The case where the host space Y equals \mathbb{R}^d (or, almost equivalently, \mathbb{S}^d , which can be modeled as a simplicial complex) has been studied the most. The case $d = 2$ corresponds to the planarity testing problem, which has attracted considerable interest [28]. The case $d = 3$ is much harder, and has only recently been shown to be decidable by Matoušek, Sedgwick, Tancer, and Wagner [22]. The general problem for arbitrary d has been extensively studied



in the last few years, starting with hardness results by Matoušek, Tancer, and Wagner [23], and continuing with several algorithmic results in a series of articles; we refer to Matoušek, Sedgwick, Tancer, and Wagner [22, Introduction] for a state of the art.

What about more general choices of Y ? The case where Y is a graph is essentially the subgraph homeomorphism problem, asking if Y contains a subdivision of a graph X . This is hard in general, easy when Y is fixed, and polynomial-time solvable for every fixed X , by using graph minor algorithms. The case where X is a graph and Y a 2-dimensional simplicial complex that is homeomorphic to a surface has been much investigated, also in connection to topological graph theory [26] and algorithms for surface-embedded graphs [12, 8]: The problem is NP-complete, as proved by Thomassen [31], but Mohar [25] has proved that it can be solved in linear time if Y is fixed (in some recent works, the proof has been simplified and the result extended [16, 18]). The case where X is a 2-complex and Y is (a 2-complex homeomorphic to) a surface, essentially boils down to the previous case; see Mohar [24]. More recently, Čadež, Krčál, Matoušek, Vokřínek, and Wagner [4, Theorem 1.4] considered the case where the host complex Y has an arbitrary (but fixed) dimension; they provide a polynomial-time algorithm for the related *map extension problem*, under some assumptions on the dimensions of X and Y ; in particular, Y must have trivial fundamental group (because they manipulate in an essential way the homotopy groups of Y , which have to be Abelian); but the maps they consider do not need to be embeddings.

Another variation on this problem is to try to embed X such that it extends a given partial embedding of X (we shall consider such *embedding extension problems* later). This problem has already been studied in some particular cases; in particular, Angelini, Battista, Frati, Jelínek, Kratochvíl, Patrignani, and Rutter [1, Theorem 4.5] provide a linear-time algorithm to decide the embedding extension problem of a graph in the plane.

Our results. In this article, we study the topological embedding problem when X is an arbitrary graph G , and Y is an arbitrary two-dimensional simplicial complex \mathcal{C} (actually, a simplicial complex of dimension at most two—abbreviated as *2-complex* below). Formally, we consider the following decision problem:

Embed(n, c):

INPUT: A graph G with at most n vertices and edges; a 2-complex \mathcal{C} with at most c simplices.

QUESTION: Does G have a topological embedding into \mathcal{C} ?

(We use the parameters n and c whenever we need to refer to the input size.) Here is our main result:

► **Theorem 1.** *The problem EMBED(n, c), parameterized by c , is in XP: it can be solved in time $f(c) \cdot n^{O(c)}$, where f is some computable function of c .*

Note that Theorem 1 shows that, for every fixed complex \mathcal{C} , the problem of deciding whether an input graph embeds into \mathcal{C} is polynomial-time solvable. Actually, our algorithm is explicit, in the sense that, if there exists an embedding of G on \mathcal{C} , we can provide some representation of such an embedding (in contrast to various results in the theory of graph minors, where the existence of an embedding can be obtained without leading to an explicit construction).

As a byproduct of our techniques, we also prove:

► **Theorem 2.** *The problem EMBED is NP-complete.*

It is actually straightforward to prove that EMBED is NP-hard, because the case where \mathcal{C} is a surface is already NP-hard. The more interesting part, the existence of a certificate checkable in polynomial time when an embedding exists, is not extremely surprising either;

we prove this using the same techniques as in the proof of Theorem 1, so in most of the article we focus on the proof of Theorem 1.

Subsequent work. Very recently, in an extended abstract [7], the first two authors improved over Theorem 1 by proving that the EMBED problem is fixed-parameter tractable in the size of the input complex. However, the techniques are very different and rely on algorithmic graph minor theory.

Why do 2-complexes look harder than surfaces? A key property of the class of graphs embeddable on a fixed surface is that it is minor-closed: Having a graph G embeddable on a surface \mathcal{S} , removing or contracting any edge yields a graph embeddable on \mathcal{S} . By Robertson and Seymour's theory, this immediately implies a cubic-time algorithm to test whether a graph G embeds on \mathcal{S} , for every fixed surface \mathcal{S} [29]. In contrast, the class of graphs embeddable on a fixed 2-complex is, in general, not closed under taking minors, and thus this theory does not apply, at least not directly. For example, let \mathcal{C} be obtained from two tori by connecting them together with a line segment, and let G be obtained from two copies of K_5 by joining them together with a new edge e ; then G embeds into \mathcal{C} , but the minor obtained from G by contracting e does not.

Two-dimensional simplicial complexes are topologically much more complicated than surfaces. For example, there exist linear-time algorithms to decide whether two surfaces are homeomorphic (this amounts to comparing the Euler characteristics, the orientability characters, and, in case of surfaces with boundary, the numbers of boundary components), or to decide whether a closed curve is contractible (see Dey and Guha [10], Lazarus and Rivaud [19], and Erickson and Whittlesey [13]). In contrast, the homeomorphism problem for 2-complexes is as hard as graph isomorphism, as shown by Ó Dúnlaing, Watt, and Wilkins [27]. Moreover, the contractibility problem for closed curves on 2-complexes is undecidable; even worse, there exists a fixed 2-complex \mathcal{C} such that the contractibility problem for closed curves on \mathcal{C} is undecidable. This is because every finitely presented group can be realized as the fundamental group of a 2-complex, and there is such a group in which the word problem is undecidable, by a result of Boone [3]; see also Stillwell [30, Section 9.3].

Despite this stark contrast between surfaces and 2-complexes, if we care only about the polynomiality or non-polynomiality, our results show that the complexities of embedding a graph into a surface or a 2-complex are similar: If the host space is not fixed, the problem is NP-complete; otherwise, it is polynomial-time solvable. Compared to the aforementioned hard problems on general 2-complexes, one feature related to our result is that every graph embeds on a 3-book (a complex made of three triangles sharing a common edge); thus, we only need to consider 2-complexes without 3-books, for otherwise the problem is trivial. This significantly restricts the structure of the 2-complexes to be considered. The problem of whether EMBED admits an algorithm that is fixed-parameter tractable in terms of the parameter c , however, remains open for general complexes, whereas it is the case when restricting to surfaces [25].

Why is embedding graphs on 2-complexes interesting? First, let us remark that, if we consider the problem of embedding graphs into simplicial complexes, then the case that we consider, in which the complex has dimension at most two, is the only interesting one, since every graph can be embedded in a 3-dimensional simplex.

We have already noted that the problem we study is more general than the problem of embedding graphs on surfaces. It is indeed quite general, and some other problems studied in

the past can be recast as an instance of EMBED or as variants of it. For example, the *crossing number* of a graph G is the minimum number of crossings in a (topological) drawing of G in the plane. Deciding whether a graph G has crossing number at most k is NP-hard, but fixed-parameter tractable in k , as shown by Kawarabayashi and Reed [17]. This is easily seen to be equivalent to the embeddability of G into the complex obtained by removing k disjoint disks from a sphere and adding, for each resulting boundary component b , two edges with endpoints on b whose cyclic order along b is interlaced. Of course, the embeddability problem on a 2-complex is more general and contains, for example, the problem of deciding whether there is a drawing of a graph G on a surface of genus g with at most k crossings. In topological graph theory, embeddings of graphs on pseudosurfaces (which are special 2-complexes) have been considered; see Archdeacon [2, Section 5.7] for a survey. Slightly more remotely, a *book embedding* of a graph G (see, e.g., Malitz [21]) is also an embedding of G into a particular 2-complex. However, the notion of book embedding incorporates additional constraints on the embedding, so that this concept does not fall under the umbrella of the problem we study.

Strategy of the proof and organization of the paper. For clarity of exposition, in most of the paper, we focus on developing an algorithm for the problem EMBED (Theorem 1). Only at the end (Section 7) we explain why our techniques imply that the problem is in NP.

The idea of the algorithm is to progressively reduce the problem to simpler problems. We first deal with the case where the complex \mathcal{C} contains a 3-book (Section 2.5). From Section 3 onwards, we reduce EMBED to various kinds of *embedding extension problems* (EEPs), namely, decision problems which ask if an embedding of a subgraph can be extended to an embedding of the whole graph in the same complex \mathcal{C} .

More precisely, in Section 3, we reduce EMBED to EEPs on a pure 2-complex (in which every segment of the complex \mathcal{C} is incident to at least one triangle), and in which every singular point of the complex is used by a vertex of the graph, by essentially guessing what happens on the isolated edges and at the singular points. At this stage, the subgraph H of the EEP whose embedding is fixed consists of isolated vertices.

In Section 4, we further reduce it to EEPs on a surface. To do so, we add some edges to H in the neighborhood of the singular points, which are “smoothed” to be made surface-like.

In Section 5, we reduce it to EEPs on a surface in which every face of H is a disk. If the input graph G embeds into the complex, it does not necessarily have a cellular embedding, but we can assume this without loss of generality by adding few edges, which we can guess. By adding suitable paths containing these edges to H , we can assume that H is cellularly embedded. We remark that this is the most technical part, relying on more advanced tools in topological graph theory.

Finally, in Section 6, we show how to solve EEPs of the latter type, as essentially proved by the third author as a subroutine in his article to decide embedability of graphs on surfaces [25].

2 Preliminaries

2.1 Embeddings of graphs into 2-complexes

A *2-complex* is an abstract simplicial complex of dimension at most two: a finite set of 0-simplices called *nodes*, 1-simplices called *segments*, and 2-simplices called *triangles* (we use this terminology to distinguish from that of vertices and edges, which we reserve for graphs); each segment is a pair of nodes, and each triangle is a triple of nodes; moreover, each

subset of size two in a triangle must be a segment. Each 2-complex \mathcal{C} corresponds naturally to a topological space, obtained in the obvious way: Start with one point per node in \mathcal{C} ; connect them by segments as indicated by the segments in \mathcal{C} ; similarly, for every triangle in \mathcal{C} , create a triangle whose boundary is made of the three segments contained in that triangle. By abuse of language, we identify \mathcal{C} with that topological space. To emphasize that we consider the abstract simplicial complex and not only the topological space, we sometimes use the name *triangulation* or *triangulated complex*.

In this paper, graphs are finite, undirected, and may have loops and multiple edges. In a similar way as for 2-complexes, each graph has an associated topological space; an *embedding* of a graph G into a 2-complex \mathcal{C} is an injective continuous map from (the topological space associated to) G to (the topological space associated to) \mathcal{C} .

We will also need the concept of branch of a graph; the number of branches of a graph can be regarded as its “topological size”. Let H be a graph. A path P with vertices u_1, u_2, \dots, u_k or a cycle C with vertices $u_1, u_2, \dots, u_{k-1}, u_k = u_1$ ($k \geq 2$) in H is a *branch* of H if $\deg_H(u_1) \neq 2$, $\deg_H(u_k) \neq 2$, and $\deg_H(u_i) = 2$ for every $i = 2, \dots, k-1$. If H has a component isomorphic to a cycle, we consider such a component as a branch as well. Under this convention we treat an isolated vertex of H as a branch as well. Then every graph is the edge-disjoint union of its branches.

2.2 Structural aspects of 2-complexes

We say that a 2-complex *contains a 3-book* if there are three distinct triangles that share a common segment.

Let p be a node of \mathcal{C} . A *cone at p* is a cyclic sequence of triangles $t_1, \dots, t_k, t_{k+1} = t_1$ ($k \geq 3$), all incident to p , such that, for each $i = 1, \dots, k$, the triangles t_i and t_{i+1} share a segment incident with p , and any other pair of triangles have only p in common. A *corner at p* is an inclusionwise maximal sequence of distinct triangles t_1, \dots, t_k ($k \geq 1$), all incident to p , such that, for each $i = 1, \dots, k-1$, the triangles t_i and t_{i+1} share a segment incident with p , any other pair of these triangles have only p in common. An *isolated segment at p* is a segment incident to p but not incident to any triangle.

If \mathcal{C} contains no 3-books, the set of segments and triangles incident with a given node p of \mathcal{C} are uniquely partitioned into cones, corners, and isolated segments. We say that p is a *regular node* if all the segments and triangles incident to p form a single cone or corner. Otherwise, p is a *singular node*. A 2-complex is *pure* if it contains no isolated segment, and each node is incident to at least one segment.

2.3 Reductions to embedding extension problems

Recall that an embedding extension problem (EEP) is the problem of deciding whether a given embedding Π of a subgraph H of G into a 2-complex \mathcal{C} can be extended to an embedding of G into \mathcal{C} . We will reduce our original problem to more and more specialized embedding extension problems (EEPs). We use the word “reduce” in a somewhat sloppy sense: A decision problem P *reduces* to k instances of the decision problem P' if solving these k instances of P' allows to solve the instance of P in time $O(k)$. We will have to be more precise when we consider the NP-completeness of EMBED in Section 7.

For further reference, the various problems considered in this article are listed in Table 1. We will have to explain how we represent the embedding Π , but this will vary throughout the proof, and we will be more precise about this in subsequent sections. Let us simply remark that, since the complexity of our algorithm is a polynomial of large degree (depending on the

6:6 Embedding Graphs into Two-Dimensional Simplicial Complexes

| Notation | Description |
|---------------|--|
| G | a graph with at most n vertices and edges |
| H | a subgraph of G with at most m branches |
| Π | an embedding of H into \mathcal{C} (or \mathcal{S}) |
| f | a map from P to $V(G) \cup \{\varepsilon\}$, where P is a set of at most m nodes in \mathcal{C} containing all singular nodes of \mathcal{C} . An embedding Γ of G in \mathcal{C} respects f if, for each $p \in P$, the following holds: If $f(p) = \varepsilon$, then p is not in the image of Γ ; otherwise, $\Gamma(f(p)) = p$. |
| \mathcal{C} | a 2-complex with at most c simplices without 3-books (also denoted by \mathcal{S} if it is a surface) |

| Problem | Input | Question |
|------------|--|---|
| EMBED | G, \mathcal{C} (possibly with 3-books) | Does G have an embedding into \mathcal{C} ? |
| EMBED-RESP | G, f, \mathcal{C} | Does G have an embedding into \mathcal{C} respecting f ? |
| EEP-SING | G, H (a set of vertices of G), Π (such that $\Pi(H)$ contains all singular nodes of \mathcal{C}), \mathcal{C} (pure) | Does G have an embedding into \mathcal{C} whose restriction to H is Π ? |
| EEP-SURF | G, H, Π, \mathcal{S} (possibly disconnected) | Does G have an embedding into \mathcal{C} whose restriction to H is Π ? |
| EEP-CONN | G, H (intersecting every connected component of G), Π, \mathcal{S} | Does G have an embedding into \mathcal{C} whose restriction to H is Π ? |
| EEP-CELL | G, H (cellularly embedded, intersecting every connected component of G), Π, \mathcal{S} | Does G have an embedding into \mathcal{C} whose restriction to H is Π ? |

■ **Table 1** The problems to which EMBED is successively reduced.


complex \mathcal{C}) in the size of the input graph, the choice of representation is not very important, because converting between any two reasonable representations is possible in polynomial time.

2.4 Surfaces

In Section 5, we will assume some familiarity with surface topology; see, e.g., [26, 30, 5] for suitable introductions under various viewpoints. We recall some basic definitions and properties. A **surface** \mathcal{S} is a compact, connected Hausdorff topological space in which every point has a neighborhood homeomorphic to the plane. Every surface \mathcal{S} is obtained from a sphere by:

- either removing the interiors of g disjoint closed disks, and identifying their boundaries pairwise, where g is an even, nonnegative integer called the (*Euler*) **genus** of \mathcal{S} ; in this case, \mathcal{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 **genus** of \mathcal{S} ; in this case, \mathcal{S} is **non-orientable**.

We emphasize that we use *Euler* genus throughout this article. A **surface with boundary** is obtained from a surface \mathcal{S} by removing a finite set of interiors of disjoint closed disks. The boundary of each disk forms a **boundary component** of \mathcal{S} . The **genus** of \mathcal{S} is defined as the genus of the ambient surface \mathcal{S} . A **possibly disconnected surface** is a disjoint union of surfaces; its genus is the sum of the genera of its connected components. An embedding



■ **Figure 1** Illustration of the proof of Proposition 3. Left: The drawing of the graph G and the dotted curve c . Right: The construction of the 3-book and the modification of the drawing.

of G into a surface \mathcal{S} , possibly with boundary, 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, with v vertices, e edges, and f faces, then Euler's formula stipulates that $2 - g - b = v - e + f$ (these quantities are referred to as the *Euler characteristic* of the surface).

An *ambient isotopy* of a surface with boundary \mathcal{S} is a continuous family $(h_t)_{t \in [0,1]}$ of self-homeomorphisms of \mathcal{S} such that h_0 is the identity.

2.5 Reduction to complexes containing no 3-book

The following folklore observation allows us to solve the problem trivially if \mathcal{C} contains a 3-book. We include a proof for completeness.

► **Proposition 3.** *If \mathcal{C} contains a 3-book, then every graph embeds into \mathcal{C} .*

Proof. Let G be a graph. We first draw G , possibly with crossings, in general position in the interior of a closed disk D . Let c be a simple curve in D with endpoints on ∂D and passing through all crossing points of the drawing of G . By perturbing c , we can ensure that, in the neighborhood of each crossing point of that drawing, c coincides with the image of one of the two edges involved in the crossing. See Figure 1, left.

Let D' be a closed disk disjoint from D . We attach D' to D by identifying c with a part of the boundary of D' and obtain a complex homeomorphic to a 3-book. Now, in the neighborhood of each crossing of the drawing of G , we push inside D' the part of the edge coinciding with c , keeping its endpoints fixed. See Figure 1, right. This removes the crossings.

So G embeds in the topological space obtained from D by attaching a part of the boundary of D' along c . But this space embeds in \mathcal{C} , because \mathcal{C} contains a 3-book. ◀

3 Reduction to EEPs on a pure 2-complex

Our next task is to reduce the problem EMBED to the following problem on a pure 2-dimensional complex:

EEP-Sing (n, m, c) :

INPUT: A graph G with at most n vertices and edges; a set H of at most m vertices of G ; a pure 2-complex \mathcal{C} with at most c simplices containing no 3-books; an injective map Π from H to the nodes of \mathcal{C} such that $\Pi(H)$ contains all singular nodes of \mathcal{C} .

QUESTION: Does G have an embedding into \mathcal{C} extending Π ?

In this section, we prove the following result.

► **Proposition 4.** *Any instance of EMBED (n, c) reduces to $(cn)^{O(c)}$ instances of EEP-SING $(cn, c, O(c))$.*

First, a definition. Consider a map $f : P \rightarrow V(G) \cup \{\varepsilon\}$, where P is a set of nodes in \mathcal{C} containing all singular nodes of \mathcal{C} . We say that an embedding Γ of G in \mathcal{C} **respects** f if, for each $p \in P$, the following holds: If $f(p) = \varepsilon$, then p is not in the image of Γ ; otherwise, $\Gamma(f(p)) = p$.

In this section, we will need the following intermediate problem:

Embed-Resp (n, m, c) :

INPUT: A graph G with at most n vertices and edges; a 2-complex \mathcal{C} (not necessarily pure) with at most c simplices that contains no 3-books; a map f as above with domain of size at most m .

QUESTION: Does G have an embedding into \mathcal{C} respecting f ?

► **Lemma 5.** *Any instance of $\text{EMBED}(n, c)$ reduces to $(O(cn))^c$ instances of $\text{EMBED-RESP}(cn, c, c)$.*

Proof. Let (G, \mathcal{C}) be an instance of $\text{EMBED}(n, c)$. By Proposition 3, we can without loss of generality assume that \mathcal{C} contains no 3-books. Let G' be the graph obtained from G by subdividing each edge k times, where $k \leq c$ is the number of singular nodes of \mathcal{C} . We claim that G has an embedding into \mathcal{C} if and only if G' has an embedding Γ' into \mathcal{C} such that each singular node of \mathcal{C} in the image of Γ' is the image of a vertex of G' .

Indeed, assume that G has an embedding Γ on \mathcal{C} . Each time an edge of G is mapped, under Γ , to pass through a singular node p of \mathcal{C} , we subdivide this edge and map this new vertex to p ; the image of the embedding is unchanged. This ensures that only vertices are mapped to singular nodes. Moreover, there were at most k subdivisions, one per singular node. So, by further subdividing the edges until each original edge is subdivided k times, we obtain an embedding of G' to \mathcal{C} such that only vertices are mapped on singular nodes. The reverse implication is obvious: If G' has an embedding into \mathcal{C} , then so has G . This proves the claim.

To conclude, for each map f from the set of singular vertices of \mathcal{C} to $V(G') \cup \{\varepsilon\}$, we solve the problem whether G' has an embedding on \mathcal{C} respecting f . The graph G embeds on \mathcal{C} if and only if the outcome is positive for at least one such map f . By construction, there are at most $(kn + 1)^k = (O(cn))^c$ such maps, because $V(G')$ has size at most kn . ◀

► **Lemma 6.** *$\text{EMBED-RESP}(n, m, c)$ reduces to $\text{EEP-SING}(n, m, O(c))$.*

Proof. Let (G, \mathcal{C}, f) be an instance of $\text{EMBED-RESP}(n, m, c)$. Formally, we describe a set of transformations on \mathcal{C} , G , and f . The invariant is that they preserve the existence or non-existence of an embedding of G into \mathcal{C} respecting f .

Step 1. We start by iterating the following procedure: While there exists a degree-two vertex v of G that is not in the image of f and whose two neighbors are distinct, we dissolve v . (By *dissolving a degree-two vertex* v , incident to edges vv_1 and vv_2 , we mean removing v and replacing vv_1 and vv_2 with a single edge v_1v_2 .) It is clear that the original graph G has an embedding in \mathcal{C} respecting f if and only if the new graph (still called G) has an embedding in \mathcal{C} respecting f .

Step 2. If a singular node is not used, then we can remove it without affecting the embeddability of G . However, removing a vertex from a 2-complex does not yield a 2-complex. We thus define a 2-complex that has the same properties. Let p be a singular node of a 2-complex \mathcal{C} . Let T be the set of triangles and segments of \mathcal{C} incident with p , uniquely partitioned into T_1, \dots, T_k , where each T_i is either a cone, a corner, or an isolated segment. The *withdrawal of p from \mathcal{C}* is the complex obtained by doing the following operation for

each $i = 1, \dots, k$: We first create a new node p_i , and then replace, in each triangle and edge of T_i , the node p by p_i .

For every singular node p of \mathcal{C} such that $f(p) = \varepsilon$ and incident to at least two segments, we withdraw p from \mathcal{C} . For each created node p_i , we let $f(p_i) = \varepsilon$. The fact that G embeds, or not, on \mathcal{C} respecting f is preserved: Indeed, if G embeds on the original complex respecting f , then this corresponds to an embedding on the complex obtained by withdrawing p , and avoiding p_1, \dots, p_k , thus respecting f ; conversely, if G embeds on the complex obtained by withdrawing p , respecting f , then it avoids the p_i s, and, after identifying together the nodes p_i , $1 \leq i \leq k$, to a single point p , this embedding avoids p , and thus respects f .

Step 3. At this point, every singular node p with $f(p) = \varepsilon$ is incident to exactly one segment, and to no triangle. For each isolated segment pq of \mathcal{C} such that $u := f(p)$ and $v := f(q)$ are both different from ε , and G contains an edge uv connecting u and v , we remove uv from G and remove pq from \mathcal{C} . We need to prove that this operation does not affect the (non-)existence of an embedding of G respecting f . First, assume that, initially, there was an embedding Γ of G on \mathcal{C} respecting f ; then either segment pq is used by uv , in which case clearly there is still an embedding after this operation, or edge uv does not use segment pq at all, in which case we can first modify Γ by embedding edge uv on segment pq and by moving the edges on pq on the space where uv was before, so we are now in the previous case. Conversely, if after this operation G has an embedding on \mathcal{C} respecting f , trivially it is also the case before.

Step 4. For each isolated segment pq of \mathcal{C} such that $u := f(p)$ and $v := f(q)$ are both different from ε , but G contains no edge connecting u and v , we do the following. In \mathcal{C} , we remove pq and add a new segment $p'q'$ where p' and q' are new nodes; we also extend f by letting $f(p') = f(q') = \varepsilon$. Finally, if G contains at least one edge of the form ux , where x has degree one and is not in the image of f , we remove a single such edge; similarly, if it contains at least one edge of the form vy , where y has degree one and is not in the image of f , we remove a single such edge. This operation does not affect our invariant, for similar reasons; for example, if initially there was an embedding of G respecting f , then the relative interior of segment pq can only contain (a) connected components of G , each reduced to a single edge of G , which we can re-embed on the new segment $p'q'$, or (b) edges of the form ux or vy , where x and y have degree one and are not in the image of f .

Step 5. For each isolated segment pq where $u := f(p)$ is different from ε but $f(q) = \varepsilon$, we remove pq from \mathcal{C} and add a new segment $p'q'$ where p' and q' are new nodes; we also extend f by letting $f(p') = f(q') = \varepsilon$. Finally, if G contains at least one edge of the form ux , where x has degree one and not in the image of f , we remove a single such edge. The invariant is preserved, for reasons similar to the previous case.

Step 6. Now, every segment of the complex (still called \mathcal{C}) is incident to one or two triangles, except perhaps some segments that are themselves connected components of \mathcal{C} and whose endpoints are not in the domain of f . If there is at least one such segment, we remove all of them from \mathcal{C} , and remove all edges uv from G that are themselves connected components of G and such that u and v are not in the image of f ; as above, this does not affect whether G embeds into \mathcal{C} respecting f .

Step 7. Finally, for each node p of \mathcal{C} incident to no segment, such that $u := f(p)$ is different from ε , we do the following: If u has degree zero, we remove p and u ; otherwise, we immediately return that G does not embed into \mathcal{C} respecting f . For each node p of \mathcal{C} that is incident to no segment and such that $f(p) = \varepsilon$, we remove p , and remove a single degree-zero vertex of G not in the image of f , if one exists.

Conclusion. Now, \mathcal{C} has no node that is itself a connected component; each of its segments is incident to one or two triangles. Also, f maps each singular node of \mathcal{C} to a vertex of G . It may map some other nodes of \mathcal{C} to ε , but such nodes are non-singular and, if an embedding uses them, a slight perturbation will avoid them, so we can remove the nodes p such that $f(p) = \varepsilon$ from the domain of f without affecting the result. Now, we have an EEP as specified in the statement of Proposition 4.

Finally, it is easy to check that, in each of the seven steps above, the numbers of vertices and edges of G do not increase, and the number of simplices of \mathcal{C} increase by at most a multiplicative constant. Moreover, the domain of f also does not increase (it increases when we withdraw singular nodes, but the images of the new nodes are ε , and such nodes are later removed from the domain of f). ◀

Proof of Proposition 4. It follows immediately from Lemmas 5 and 6. ◀

4 **Reduction to an EEP on a possibly disconnected surface**

The previous section led us to an embedding extension problem in a pure 2-complex without 3-books where the images of some vertices are predetermined. Now, we show that solving such an EEP amounts to solving another EEP in which the complex is a surface:

EEP-Surf(n, m, c):

INPUT: A graph G with at most n vertices and edges; a subgraph H of G with at most m branches; an embedding Π of H into a possibly disconnected surface \mathcal{S} without boundary, represented as a triangulation T with at most c simplices, such that Π and T have $O(c)$ crossings.

QUESTION: Does G have an embedding into \mathcal{S} extending Π ?

To represent the embedding Π in such an EEP instance (G, H, Π, \mathcal{S}) , it will be convenient to use the fact that, in all our constructions below, the image of every connected component of H under Π will intersect the 1-skeleton of \mathcal{S} at least once, but only in a finite number of points. (Note that H may use some nodes of \mathcal{S} .) Consider the *overlay* of the triangulation of \mathcal{S} and of Π , the union of the 1-skeleton of \mathcal{S} and of the image of Π ; this overlay is the image of a graph on \mathcal{S} ; each of its edges is either a piece of the image of an edge of H or a piece of a segment of \mathcal{S} ; each of its vertices is the image of a vertex of H and/or a node of \mathcal{S} , or an intersection point between the 1-skeleton of \mathcal{S} and the image of Π . By the assumption above on Π , this overlay is cellularly embedded on \mathcal{S} , and we can represent it by its combinatorial map [20, 11] (possibly on surfaces with boundary, since at intermediary steps of our construction we will have to consider such surfaces).


In this section, we prove the following proposition.

► **Proposition 7.** *Any instance of EEP-SING(n, m, c) reduces to an instance of EEP-SURF($O(n + m + c), O(m + c), O(c)$).*


We will first reduce the original EEP to an intermediary EEP on a surface with boundary.

► **Lemma 8.** *Any instance of EEP-SING(n, m, c) reduces to an instance of EEP-SING($n + O(c), m + O(c), O(c)$) in which the considered 2-complex is a possibly disconnected surface with boundary.*

Proof. Let (G, H, \mathcal{C}, Π) be an instance of EEP-SING(n, m, c). The key property that we will use is that, since \mathcal{C} is pure and contains no 3-books, each singular node is incident to cones and corners only.



■ **Figure 2** The modification of singular vertices in the proof of Lemma 8. We transform the neighborhood of each singular vertex to make it surface-like. Moreover, we add to H one loop per cone or corner, and for each corner, we add a vertex and an edge.



■ **Figure 3** In the proof of Lemma 8, we push some parts of the graph outside S_p by an ambient isotopy of \mathcal{S} that moves a small enough open disk or annulus (in blue), which is disjoint from L_p and E_p , to a larger part.

Figure 2 illustrates the proof. Let (G, H, Π, \mathcal{C}) be the instance of EEP-SING. We first describe the construction of the instance $(G', H', \Pi', \mathcal{S})$ on the possibly disconnected surface with boundary. Let p be a singular node; we modify the complex in the neighborhood of p as follows. Let c_p be the number of cones at p and c'_p be the number of corners at p . We remove a small open neighborhood N_p of p from \mathcal{C} , in such a way that the boundary of N_p is a disjoint union of c_p circles and c'_p arcs. We create a sphere S_p with $c_p + c'_p$ boundary components. Finally, we attach each circle and arc to a different boundary component of S_p , choosing an arbitrary orientation for each gluing; circles are attached to an entire boundary component of S_p , while arcs cover only a part of a boundary component of S_p . Doing this for every singular node p , we obtain a surface (possibly disconnected, possibly with boundary), which we denote by \mathcal{S} .

We now define H' , G' , and Π' from H , G , and Π (again, refer to Figure 2). Let p be a singular node of \mathcal{C} and v_p the vertex of H mapped on p by Π . In H (and thus also G), we add a set L_p of $c_p + c'_p$ loops with vertex v_p . Let q_p be a point in the interior of the punctured sphere S_p ; in Π , we map v_p to q_p , and we map these $c_p + c'_p$ loops on S_p in such a way that each loop encloses a different boundary component of S_p (thus, if we cut S_p along these loops, we obtain $c_p + c'_p$ annuli and one disk).

Finally, we add to H (and thus also to G) a set E_p of c'_p new edges, each connecting v_p to a new vertex. In Π , each new vertex is mapped to the boundary component of S_p corresponding to a corner, but not on the corresponding arc.


Let us call G' and H' the resulting graphs, and Π' the resulting embedding of H' . Note that, from the triangulation of \mathcal{C} with c simplices, we can easily obtain a triangulation of \mathcal{S} with $O(c)$ simplices, and that the image of each edge of H crosses $O(1)$ edges of this triangulation.

It remains to prove that the two EEPs constructed are equivalent.

Let us first prove that any solution Γ of the EEP instance (G, H, Π, \mathcal{C}) yields a solution of the EEP instance $(G', H', \Pi', \mathcal{S})$. Let p be a singular node of \mathcal{C} , and let v_p be the vertex of H mapped to p . We need to modify Γ locally in the neighborhood N_p of p that is removed when transforming \mathcal{C} to \mathcal{S} . Without loss of generality, up to an ambient isotopy of Γ that does not move H , we can assume that the image of Γ intersects N_p exactly in straight line segments having p as one endpoint. To build a solution of $(G', H', \Pi', \mathcal{S})$, we first remove the part of Γ inside N_p , and we reconnect q_p (to which v_p is mapped) to each point of the image of Γ lying on the boundary of N_p , by paths on S_p ; this is certainly possible because each point of S_p not in the image of $E_p \cup L_p$ can be connected to q_p by a path that does not meet the image of $E_p \cup L_p$ (except at q_p). Thus, $(G', H', \Pi', \mathcal{S})$ has a solution.

Conversely, let Γ' be a solution of the EEP $(G', H', \Pi', \mathcal{S})$; we build a solution of (G, H, Π, \mathcal{C}) . As above, let q_p be a point of \mathcal{S} that was obtained from a singular node p of \mathcal{C} . We now show that we can assume that Γ does not enter S_p , except for v_p , E_p , L_p , and the edges in G' incident to v_p .

First, let D_p be the connected component of S_p minus the image of L_p that is a disk. The part of Γ lying in D_p corresponds to a planar subgraph of G , connected to the rest of G by v_p only; we can re-embed this planar subgraph of G in $S_p \setminus D_p$. Next, consider a cone at p . The part of S_p enclosed by the corresponding loop of L_p is an annulus; the situation is as on Figure 3, top, and, by an ambient isotopy of \mathcal{S} , we can push the part of Γ that lies in the annulus outside it, except for those edges touching v_p . Finally, consider a corner at p , and the annulus that is the part of S_p enclosed by the corresponding loop in L_p . The local picture for this part of S_p is as shown on Figure 3, bottom, and similarly by an ambient isotopy of \mathcal{S} we can push the part of Γ on S_p outside it, except for those edges touching v_p .



■ **Figure 4** The removal of boundary components in the proof of Lemma 9. We attach a disk to every boundary component of the surface. Moreover, if there is a boundary component c containing at least one vertex of H , say v_1, \dots, v_k , we add to H a new vertex v , mapped inside the corresponding disk, and add, in H , one edge between v and each of the v_i s, that edge being mapped inside the disk.

Now, a solution of (G, H, Π, \mathcal{C}) can be obtained by the following procedure, for each of the singular nodes p : (1) Remove the sphere S_p , together with the image of Γ inside S_p ; (2) add the neighborhood N_p of p ; (3) reconnect to p the points on the image of Γ that lie on the boundary of N_p . ◀

We now deduce from the previous EEP the desired EEP on a surface without boundary.

▶ **Lemma 9.** *Any instance of EEP-SING(n, m, c) on a possibly disconnected surface with boundary reduces to an instance of EEP-SURF($n + O(m), O(m), O(c)$).*

Proof. Figure 4 illustrates the proof. Let (G, H, Π, \mathcal{S}) be an instance of EEP-SING on a possibly disconnected surface with boundary. We first describe the construction of $(G', H', \Pi', \mathcal{S}')$, the instance of EEP-SURF.

Let \mathcal{S}' be obtained from \mathcal{S} by gluing a disk D_b along each boundary component. Let b be a boundary component of \mathcal{S} . If Π maps at least one vertex to b , then we add to H (and thus also to G) a new vertex v_b , which we connect, also by a new edge, to each of the vertices mapped to b by Π . We extend Π by mapping vertex v_b and its incident edges inside D_b . Let us call G' and H' the resulting graphs, and Π' the resulting embedding of H' . For each vertex of H on a boundary component, we added to H and G at most one vertex and one edge and thus at most two branches. There remains to prove that these two EEPs are equivalent. Clearly, any solution of (G, H, Π, \mathcal{S}) yields a solution of $(G', H', \Pi', \mathcal{S}')$. Conversely, let Γ' be a solution of $(G', H', \Pi', \mathcal{S}')$; we build a solution of (G, H, Π, \mathcal{S}) . Let b be a boundary of \mathcal{S} . There are two cases:

- If no vertex of H is mapped to b , then Π' maps H' outside the closure of D_b , so, by an ambient isotopy of Γ' that does not move H' , we can push Γ' outside D_b .
- Otherwise, the disk D_b is split into sectors by the edges of $\Pi'(H')$ incident to v_b . The pieces of $\Pi'(H - H')$ in a given sector can be pushed out from that sector by an ambient isotopy that keeps the image of H' fixed.

After doing this for every boundary component b , we obtain that the restriction of Γ' to G is a solution of (G, H, Π, \mathcal{S}) . ◀

Finally:

Proof of Proposition 7. It suffices to apply Lemmas 8 and 9. ◀

5 Reduction to a cellular EEP on a surface

Let us consider the following problem.

EEP-Cell(n, m, g):

INPUT: A graph G with at most n vertices and edges; a subgraph H of G with at most m branches, intersecting every connected component of G ; a cellular embedding Π of H into a surface \mathcal{S} (connected and without boundary) with genus at most g .

QUESTION: Does G have an embedding into \mathcal{S} extending Π ?

In this section, we prove the following proposition.

► **Proposition 10.** *Any instance of EEP-SURF(n, m, c) reduces to $(n+m+c)^{O(m+c)}$ instances of EEP-CELL($O(n+m+c), O(m+c), c$).*

For the proof of Proposition 10, which is given in the rest of Section 5, it will be necessary not to store an embedding Π of a graph G on a surface \mathcal{S} by its overlay with the triangulation, as was done in the previous section. Instead, we forget the triangulation. In other words, we have to store the combinatorial map corresponding to Π , but taking into account the fact that Π is not necessarily cellular. We need to store, for each face of the embedding, whether it is orientable or not, and a pointer to an edge of each of its boundary components (with some orientation information). Such a data structure is known under the name of *extended combinatorial map* [6, Section 2.2] (only orientable surfaces were considered there, but the data structure readily extends to non-orientable surfaces). We start by converting (in polynomial time) the instance of EEP-SURF(n, m, c) to this representation; note that the genus of the surface \mathcal{S} in the instance of EEP-SURF is at most c .

5.1 Reduction to connected surfaces

We first build intermediary EEPs over connected surfaces:

EEP-Conn(n, m, g):

INPUT: A graph G with at most n vertices and edges; a subgraph H of G with at most m branches intersecting every connected component of G ; an embedding Π of H into a surface \mathcal{S} (connected and without boundary) with genus at most g .

QUESTION: Does G have an embedding into \mathcal{S} extending Π ?

► **Lemma 11.** *Any instance of EEP-SURF(n, m, c) reduces to $O(c(m+c)^c)$ instances of EEP-CONN($n, m+c, c$).*

More precisely (and this is a fact that will be useful to prove that EMBED is in NP, see Theorem 2), any instance of EEP-SURF(n, m, c) is equivalent to the disjunction (OR) of at most $(m+c)^c$ instances, each of them being the conjunction (AND) of $O(c)$ instances of EEP-CONN($n, m+c, c$).

Proof. Let (G, H, Π) be an instance of EEP-SURF(n, m, c). We start by removing the connected components of G that are planar and disjoint from H . (Testing planarity takes linear time [14].) This does not change the solution of the EEP, because such connected components can be embedded on an arbitrarily small planar portion of \mathcal{S} (provided \mathcal{S} is non-empty, but otherwise the original EEP-SURF instance can be solved trivially). So without loss of generality, every connected component of G disjoint from H is non-planar. Let V_0 be an arbitrary set of vertices, one per non-planar connected component of G disjoint from H . Without loss of generality, the number of vertices in V_0 is at most the genus of \mathcal{S} , and hence at most c , because otherwise the initial EEP has no solution.

Let $H' := H \cup V_0$; every connected component of G intersects H' . For each vertex of V_0 , we guess the face of Π it has to be embedded in, and extend Π accordingly, by adding the images of V_0 in Π ; let Π' be the resulting embedding of H' . By the previous paragraph, the number of these guesses is at most $(m+c)^{|V_0|} \leq (m+c)^c$. It is clear that the initial EEP has a solution if and only if one of these EEPs $(G, H', \Pi', \mathcal{S})$ has a solution.

These EEPs are almost of the form announced in the lemma, except that \mathcal{S} can be disconnected. However, in any solution of this EEP, we know the connected component of \mathcal{S} each connected component of G has to embed in, because each connected component of G intersects H . We can thus reformulate the EEP as the conjunction of several EEPs, one per connected component of \mathcal{S} . (Of course, we can discard the connected components of \mathcal{S} disjoint from H .) ◀

5.2 Simplifying the faces

The strategy for the proof of Proposition 10 is as follows. For each EEP $(G', H', \Pi', \mathcal{S}')$ from the previous lemma, we will extend H' to make it cellularly embedded in \mathcal{S} by adding either paths connecting two boundary components of a face of H' , or paths with endpoints on the same boundary component of a face of H' in a way that the genus of the face decreases. We first define an invariant that will allow us to control the number of steps needed until this process terminates.

Let Π be an embedding of a graph H on a surface \mathcal{S} . The *cellularity defect* of (H, Π, \mathcal{S}) is the non-negative integer

$$cd(H, \Pi, \mathcal{S}) := \sum_{f \in \mathcal{F}(\Pi)} \text{genus}(f) + \sum_{f \in \mathcal{F}(\Pi)} (\text{number of boundaries of } f - 1)$$

where $\mathcal{F}(\Pi)$ denotes the set of faces of Π .

Some obvious remarks: Π can contain isolated vertices. By convention, each of them counts as a boundary component of the face of Π it lies in. With this convention, every face of H has at least one boundary component, except in the very trivial case when G is empty. This implies that Π is a cellular embedding if and only if $cd(H, \Pi, \mathcal{S}) = 0$. We will also use the following property, usually without mentioning it.

► **Lemma 12.** *Let \mathcal{S} be a surface with genus g and let H be a graph with m branches that is embedded in \mathcal{S} . Then $cd(H, \Pi, \mathcal{S}) < 2m + g$.*

Proof. We have

$$\sum_{f \in \mathcal{F}(\Pi)} \text{genus}(f) \leq \text{genus}(\mathcal{S}) = g.$$

Similarly, every boundary component contains a branch of H and each branch of H participates in at most two boundary components (or twice in the same boundary component). Thus,

$$\sum_{f \in \mathcal{F}(\Pi)} (\text{number of boundaries of } f - 1) \leq 2m. \quad \blacktriangleleft$$

We need some auxiliary definitions and lemmas. The reversal of a walk w in a graph is denoted by \bar{w} . The concatenation of two walks w and w' is denoted by $w \cdot w'$. We use the same notation for paths on a surface. A *lollipop* in a graph is a closed walk of the form $p \cdot q \cdot \bar{p}$, where p is a (simple) path (possibly reduced to a single vertex) and q is a cycle, such

that p and q share exactly one vertex (the end-vertex of p). The **basepoint** of the lollipop is the initial vertex of p . The following lemma reduces an EEP to EEPs with a smaller cellularity defect.

► **Lemma 13.** *Let (G, H, Π, \mathcal{S}) be an instance of $\text{EEP-CONN}(n, m, g)$ such that (H, Π, \mathcal{S}) has positive cellularity defect. Then this instance reduces to $O(m^2 n^2)$ instances $(G', H', \Pi', \mathcal{S})$ of $\text{EEP-CONN}(n+1, m+3, g)$ where $\text{cd}(H', \Pi', \mathcal{S}) < \text{cd}(H, \Pi, \mathcal{S})$. For each such instance, the graph G' is obtained from G by adding exactly one edge, and the subgraph H' is obtained from H by adding a path or a lollipop containing the added edge that is embedded under Π' in one of the faces of the embedding Π of H in such a way that it does not separate that face.*

Admitting Lemma 13, the proof of Proposition 10 is straightforward:

Proof of Proposition 10. Starting with an instance of $\text{EEP-SURF}(n, m, c)$, we first apply Lemma 11, obtaining $O(c(m+c)^c)$ instances of $\text{EEP-CONN}(n, m+c, c)$. To each of these EEPs, we apply recursively Lemma 13 until we obtain cellular EEPs. The cellularity defect of the initial instance (G, H, Π, \mathcal{S}) is $O(m+c)$, by Lemma 12. Thus, the number of instances of EEP-CELL at the bottom of the recursion tree is $(n+m+c)^{O(m+c)}$, in each of which the size of the graph is $O(n+m+c)$, the pre-embedded subgraph has $O(m+c)$ branches, and the surface is the same, thus having genus at most g . ◀

5.3 Proof of Lemma 13

It remains to prove Lemma 13. The proof uses standard notions in surface topology, homotopy, and homology. We refer to textbooks and surveys [26, 30, 5]. We only consider homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$.

Let f be a surface with a single boundary component and let a be a (possibly non-simple) arc in f intersecting the boundary component of f exactly at its endpoints. If we contract the boundary component to a single point, the arc a becomes a loop, which can be null-homologous or non-null-homologous, and one-sided or two-sided. We employ the same adjectives, null-homologous, non-null-homologous, one-sided, and two-sided, for the arc a .

► **Lemma 14.** *Let f be a surface with boundary, let b be a point in the interior of f , and let b_1, b_2 , and b_3 be points on the boundary of f . For each i , let p_i be a (possibly non-simple) path connecting b_i to b . Let $a_1 = p_2 \cdot \bar{p}_3$, $a_2 = p_3 \cdot \bar{p}_1$, and $a_3 = p_1 \cdot \bar{p}_2$. Let \mathcal{P} be any of the following possible properties of the arcs a_i :*

- “the endpoints of a_i lie on the same boundary component of f ”;
- “ a_i is null-homologous” (if b_1, b_2 , and b_3 lie on the same boundary component);
- “ a_i is two-sided” (if b_1, b_2 , and b_3 lie on the same boundary component).

Then the following holds: If both a_1 and a_2 have property \mathcal{P} , then so has a_3 .

(We remark that being null-homologous or two-sided makes only sense for arcs with endpoints on the same boundary component, whence the restriction that f has a single boundary component in the second and third items.)

Proof. This is a variant on the *3-path condition* from Mohar and Thomassen [26, Section 4.3]. The first item is immediate. The second one follows from the fact that homology is an algebraic condition: The concatenation of two null-homologous paths is null-homologous, and removing spurs (subpaths of the form $q \cdot \bar{q}$) from a path does not affect homology. The third one is similar: The concatenation of two two-sided paths is two-sided, and removing spurs does not affect sidedness. ◀

► **Lemma 15.** *Let G be a graph, H a subgraph of G , and Π be an embedding of H on \mathcal{S} . Let K be a subgraph of G that is either a (simple) path intersecting H exactly at its endpoints, or a lollipop intersecting H exactly at its basepoint. Also, assume that there exists an embedding Π' of $H \cup K$ extending Π such that the embedding of K does not separate the face f of Π it belongs to. Let f' be the face of Π' corresponding to f . Then $cd(H \cup K, \Pi', \mathcal{S}) = cd(H, \Pi, \mathcal{S}) - 1$. Moreover:*

1. *If K has its endpoints on two distinct boundary components of f , then f' has the same genus and orientability character as f , and one boundary component less than f .*
2. *If K has its endpoints on the same boundary component of f , then we have the following possibilities:*
 - *if K is two-sided, then the genus of f' equals that of f , minus two, and the number of boundary components of f' equals that of f , plus one;*
 - *if K is one-sided, then f is non-orientable, the genus of f' equals that of f minus one, and the number of boundary components of f' equals that of f . Moreover:*
 - *if the genus of f is one, then f' is orientable;*
 - *if the genus of f is even, then f' is non-orientable.*

Proof. Let us first prove that $cd(H \cup K, \Pi', \mathcal{S}) = cd(H, \Pi, \mathcal{S}) - 1$. By the choice of K , we have that the Euler characteristic of f' exceeds that of f exactly by one. (Proof: The Euler characteristic is a topological invariant. Take any triangulation T of f containing K . If K is a path, cutting T along K to obtain a triangulation of f' duplicates $p + 1$ vertices and p edges of T , where p is the number of edges of K , so the Euler characteristic increases by one. The same statement holds if K is a lollipop.) Moreover, the Euler characteristic of f equals two, minus its genus, minus its number of boundary components, and similarly for f' . The result on the cellularity defect follows.


The other assertions also follow from Euler characteristic arguments. Specifically, we first remark that, in all cases, the claim on the number of boundary components of f' is correct (this only depends on whether K connects the same or different boundary components of f , and whether it is one- or two-sided). Then the genus follows, as above, from Euler characteristic arguments. The claims on the orientability use the following facts: If f is orientable, then so is f' ; if f' has genus zero, then it is orientable; if f' has odd genus, then it is non-orientable. ◀

Proof of Lemma 13. Since $cd(H, \Pi, \mathcal{S}) \geq 1$, there must be a face f of H with either (1) at least two boundary components, or (2) a single boundary component but positive genus. We will consider each of these cases separately, but first we introduce some common terminology.

Let F be an arbitrary spanning forest of $G - E(H)$ rooted at $V(H)$, namely, a subgraph of $G - E(H)$ that is a forest with vertex set $V(G)$ such that each connected component of F contains exactly one vertex of $V(H)$, its *root*. The algorithm starts by computing an arbitrary such forest F in linear time.

For each vertex u of G , let $r(u)$ be the unique root in the same connected component of F as u , and let $F(u)$ be the unique path connecting u to $r(u)$. If u and v are two vertices of G , let G_{uv} be the graph obtained from G by adding one edge, denoted uv , connecting u and v . (This may be a parallel edge if u and v were already adjacent in G , but in such a situation when we talk about edge uv we always mean the new edge.) Let $F(uv)$ be the unique walk in G between u and v that is the concatenation of $\overline{F(u)}$, edge uv , and $F(v)$. Thus, $F(uv)$ intersects H precisely at its endpoints. Note that $F(uv)$ is either a path with its ends in H or a lollipop with its basepoint in H (and otherwise disjoint from H).

Case 1: f has at least two boundary components.



■ **Figure 5** Illustration of the proof of Lemma 13, case 1. Face f of H is an annulus. Only the spanning forest F , together with edge uv , are shown. The thick lines depict $F(uv)$ (a path), with edge uv even thicker.


Assume that (G, H, Π, \mathcal{S}) has a solution Γ . We claim that, for some vertices u and v of G , the embedding Γ extends to an embedding of G_{uv} in which the image of $F(uv)$ lies in f and connects two distinct boundary components of f (see Figure 5).

Indeed, let γ be a curve drawn in f connecting two vertices of H in different boundary components of f and chosen such that it intersects the boundary of f exactly at its endpoints. We can deform γ so that it intersects Γ only at the images of vertices, and never in the relative interior of an edge. We can, moreover, assume that γ is simple (except perhaps that its endpoints may coincide if they join occurrences of the same vertex in different boundary components of f). Let v_1, \dots, v_k be the vertices of G encountered by γ , in this order. We denote by $\gamma[i, j]$ the segment of γ between vertices v_i and v_j . For some i , we have that $\overline{F(v_i)} \cdot \gamma[i, i+1] \cdot F(v_{i+1})$ connects two different boundary components of f : Otherwise, by induction on i , applying the first case of Lemma 14 to the three paths $\gamma[1, i]$, $F(v_i)$, and $\gamma[i, i+1] \cdot F(v_{i+1})$, we would have that, for each i , $\gamma[1, i] \cdot F(v_i)$ has its endpoints on the same boundary component of f , which is a contradiction for $i = k$ (for which the curve is γ). So let i be such that $\overline{F(v_i)} \cdot \gamma[i, i+1] \cdot F(v_{i+1})$ connects two different boundary components of f . By letting $u = v_i$ and $v = v_{i+1}$, and embedding the edge uv as $\gamma[i, i+1]$, gives the desired embedding of G_{uv} . This proves the claim.

By definition of $F(uv)$, and since it connects two distinct boundary components of f , it is actually a path (without repeated vertices) in f . (Its endpoints may coincide on \mathcal{S} .) The strategy now is to guess the vertices u and v and the way the path $F(uv)$ is drawn in f , and to solve a set of EEPs $(G_{uv}, H \cup F(uv), \Pi', \mathcal{S})$ where Π' is chosen as an appropriate extension of Π .

Case 1a: Let us first assume that f is orientable. One subtlety is that, given u and v , there can be several essentially different ways of embedding $F(uv)$ inside f , if there is more than one occurrence of $r(u)$ and $r(v)$ on the boundary of f . So we reduce our EEP to the following set of EEPs: For each choice of vertices u and v of G , and any occurrences of $r(u)$ and $r(v)$ on the boundary of f , we consider the EEP $(G_{uv}, H \cup F(uv), \Pi', \mathcal{S})$ where Π' extends Π and maps $F(uv)$ to an arbitrary path in f connecting the chosen occurrences of $r(u)$ and $r(v)$ on the boundary of f .

It is clear that, if one of these new EEPs has a solution, the original EEP has a solution. Conversely, let us assume that the original EEP (G, H, Π, \mathcal{S}) has a solution; we now prove that one of these new EEPs has a solution. By our claim above, for some choice of u and v , some EEP $(G_{uv}, H \cup F(uv), \Pi'', \mathcal{S})$ has a solution, for some Π'' mapping $F(uv)$ inside f and connecting different boundary components of f . In that mapping, $F(uv)$ connects two occurrences of $r(u)$ and $r(v)$ inside f . We prove that, for these choices of occurrences of $r(u)$ and $r(v)$, the corresponding EEP described in the previous paragraph, $(G_{uv}, H \cup F(uv), \Pi', \mathcal{S})$, has a solution as well. These two EEPs are the same except that



■ **Figure 6** Illustration of the proof of Lemma 13, case 2. Face f of H is a torus with a single boundary component. Only the spanning forest F , together with edge uv , are shown. The thick lines depict $F(uv)$ (a lollipop), with edge uv even thicker.

the path $F(uv)$ may be drawn differently in Π' and Π'' , although they connect the same occurrences of $r(u)$ and $r(v)$ on the boundary of f . By Lemma 15, under Π' , the face f is transformed into a face f' that has the same genus and orientability character as f , but one boundary component less. The same holds, of course, for Π'' . Moreover, the ordering of the vertices on the boundary components of the new face is the same in Π' and Π'' . Thus, there is a homeomorphism h of f that keeps the boundary of f fixed pointwise and such that $h \circ \Pi''|_{F(uv)} = \Pi'|_{F(uv)}$. This homeomorphism, extended to the identity outside f , maps any solution of $(G_{uv}, H \cup F(uv), \Pi'', \mathcal{S})$ to a solution of $(G_{uv}, H \cup F(uv), \Pi', \mathcal{S})$, as desired.

It also follows from Lemma 15 that the cellularity defect decreases by one. To conclude this case, we note that the number of new EEPs is $O(m^2n^2)$: indeed, there are $O(n)$ possibilities for the choice of each of u and v , and $O(m)$ possibilities for the choice of each of the occurrence of $r(u)$ and $r(v)$ on the boundary of f .

Case 1b: Let us assume that f is non-orientable. The same argument works, except that there are two possibilities for the cyclic ordering of the vertices along the new boundary component of the new face: If we walk along one of the boundary components of f (in an arbitrary direction), use p , and walk along the other boundary component of f , we do not know in which direction this second boundary component is visited. So we actually need to consider two EEPs for each choice of u , v , and occurrences of $r(u)$ and $r(v)$, instead of one. The rest is unchanged.

Case 2: f has a single boundary component and positive genus. The proof is similar to Case 1. The main difference is that, instead of curves connecting different boundary components of f , we now consider curves in f that are non-null-homologous. Another difference is that the walks $F(uv)$ may repeat vertices and edges; however, by construction, $F(uv)$ is either a (simple) path or a lollipop.

Case 2a: Let us assume that f is orientable. Assume that (G, H, Π, \mathcal{S}) has a solution Γ . We claim that, for some vertices u and v of G (allowing the possibility that $u = v$), the embedding Γ extends to an embedding of G_{uv} in which the image of $F(uv)$ lies in f and is non-null-homologous. See Figure 6. The proof is similar in spirit to the corresponding claim in Case 1: We let γ be a non-null-homologous curve in f intersecting the boundary of f exactly at its endpoints; we can assume similarly as before that it is simple and intersects only vertices of Γ , in the order v_1, \dots, v_k . For some i , $\overline{F(v_i)} \cdot \gamma[i, i+1] \cdot F(v_{i+1})$ must be non-null-homologous, by induction and by Lemma 14; this gives an embedding of G_{uv} .

So we reduce the original EEP to the following set of EEPs: For each choice of vertices

u and v of G , and each occurrence of $r(u)$ and $r(v)$ on the boundary of f , we consider the EEP $(G_{uv}, H \cup F(uv), \Pi', \mathcal{S})$ where Π' extends Π and embeds $F(uv)$ in f in such a way that (1) $F(uv)$ is non-null-homologous, and (2) $F(uv)$ connects the chosen occurrences of $r(u)$ and $r(v)$ on the boundary of f . As before, the only subtlety is to prove that, if we have two EEPs $(G_{uv}, H \cup F(uv), \Pi', \mathcal{S})$ and $(G_{uv}, H \cup F(uv), \Pi'', \mathcal{S})$ such that $F(uv)$ are not embedded exactly in the same way in Π' and Π'' , but are non-null-homologous in f and connect the same occurrences of $r(u)$ and $r(v)$ on the boundary of f , then these EEPs are equivalent. This follows from Lemma 15: The image of $F(uv)$ in Π' is non-null-homologous, thus non-separating, and thus cuts f into a face that is an orientable surface with two boundary components and with (Euler) genus that of f minus two; the same holds in Π'' . Moreover, the ordering of the vertices along the boundary components of the new face is the same in both Π' and Π'' . Thus, as in the previously treated case, there is a homeomorphism h of f that keeps the boundary of f fixed pointwise and such that $h \circ \Pi''|_{F(uv)} = \Pi'|_{F(uv)}$. We complete the proof in the same way as before. It also follows from Lemma 15 that the cellularity defect decreases by one. The number of these new EEPs is $O(m^2n^2)$.

Case 2b: Let us assume that f is non-orientable. Also assume that (G, H, Π, \mathcal{S}) has a solution Γ . We claim that, for some vertices u and v of G (allowing the possibility that $u = v$), the embedding Γ extends to an embedding of G_{uv} in which the image of $F(uv)$ lies in f and is one-sided. The proof is as in Case 1a, replacing “non-null-homologous” with “one-sided”.

Then, the genus of f is even or equal to 1, then a similar argument can be used: Regardless of the way we embed $F(uv)$ in a way that (1) it is one-sided and (2) it connects the chosen occurrences $r(u)$ and $r(v)$ on the boundary of f , cutting f along $F(uv)$ results in a surface whose topology is uniquely determined, by Lemma 15. (Note that $F(uv)$ is non-separating because it is one-sided.) Moreover, the ordering of the vertices along the boundary of the new face is uniquely determined. The cellularity defect also decreases by one, and the same argument as above concludes.

Finally, if the genus of f is odd and at least three, then cutting f along $F(uv)$ results in a surface in which the ordering of the vertices along the single boundary component is uniquely determined, but this surface, with one boundary component and with genus that of f minus one, can be orientable or not. Thus, for each choice of vertices u and v of G , and each occurrence of $r(u)$ and $r(v)$ on the boundary of f , we actually need to consider two EEPs, one in which $F(uv)$ is embedded as a one-sided curve in a way that it cuts f into an orientable surface, and one in which $F(uv)$ is embedded in a way that it cuts f into a non-orientable surface. The rest of the argument is unchanged. \blacktriangleleft

6 Solving a cellular EEP on a surface

► **Theorem 16** ([25]). *There is a computable function $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that every instance of EEP-CELL(n, m, c) can be solved in time $f(m, c) \cdot O(n)$.*

When solving EEPs on surfaces, a useful property of the embedded subgraph is the following one. We say that a subgraph H of G has *property (E)* if H has no local bridges. This means that every branch of H forms an induced subgraph of G , and for every connected component of $G - V(H)$, its neighbors in H are not all contained in a single branch of H . In fact, what we need is a weaker version of property (E) where we first prescribe a subset V_0 of vertices of H , where V_0 contains all vertices whose degree in H is different from two, and possibly a constant number of vertices whose degree in H is equal to two. To each vertex in V_0 , whose degree is 2 in H , we add an edge to make it a degree-3 vertex. Let \hat{H} be

the resulting graph. Then we say that H has property (E) with respect to V_0 if the above property holds with respect to the branches of \hat{H} .

Proof of Theorem 16. This is essentially the main result from [25]. The algorithm from [25] first reduces the problem to an instance $(G', H', \Pi', \mathcal{S})$ such that H' satisfies property (E) with respect to a subset V_0 that contains all vertices of H of degree different from two (all these are also vertices in H' of degree different from two) plus a constant number of vertices of degree two in H' , where this constant number is bounded from above in terms of the genus of \mathcal{S} , which is itself bounded from above by c . For this purpose, [25] relies on another paper [15]. After achieving this property, the paper [25] reduces the EEP to a constant number (where the constant depends on c) of “simple” extension problems [25, Section 4], which are then solved in [25, Theorem 5.4]. ◀

7 Proof of Theorems 1 and 2

We can finally prove our main results. First, let us prove that we have an algorithm with complexity $f(c) \cdot n^{O(c)}$ for some computable function f .

Proof of Theorem 1. This immediately follows from Propositions 3, 4, 7, 10, and Theorem 16. Consider an instance of $\text{EMBED}(n, c)$. Proposition 3 allows to discard the 2-complexes containing a 3-book. Proposition 4 reduces the problem to $(cn)^{O(c)}$ instances of $\text{EEP-SING}(cn, c, O(c))$. Proposition 7 reduces each such instance into an instance of $\text{EEP-SURF}(O(cn), O(c), O(c))$. Proposition 10 reduces that instance into $O(cn)^{O(c)}$ instances of $\text{EEP-CELL}(O(cn), O(c), O(c))$. Theorem 16 shows that each such instance can be solved in time $f(O(c), O(c)) \cdot O(cn) \leq f_0(c)n$ for an appropriate computable function $f_0 : \mathbb{N} \rightarrow \mathbb{N}$. ◀

Finally, we prove that the problem EMBED is NP-complete:

Proof of Theorem 2. Let us first prove that the problem is NP-hard. The following problem GRAPH-GENUS is NP-hard: Given a graph G and an integer g , decide whether G embeds on the orientable surface of genus g [31]. This almost immediately implies that EMBED is NP-hard; the only subtlety is that in GRAPH-GENUS , g is specified in binary, thus more compactly than a triangulated surface of genus g (and thus $\Omega(g)$ triangles). To be very precise, given an instance (G, g) of GRAPH-GENUS , we transform it in polynomial time into an equivalent instance of EMBED as follows: If G has at most g edges, then we transform it into a constant-size positive instance of EMBED (every graph with g edges embeds on the orientable surface of genus g); otherwise, we consider the instance (G, \mathcal{C}) where \mathcal{C} is a 2-complex that is an orientable surface of genus g ; since G has at least g edges, the transformation takes polynomial time in the size of (G, g) .

We now prove that the problem EMBED belongs to NP. The case where \mathcal{C} contains a 3-book is trivial; let us assume that it is not the case. The proof of Proposition 7 shows that an EMBED instance is positive if and only if at least one instance of EEP-SURF , among $(cn)^{O(c)}$ of them, is positive. The certificate indicates which of these instances is positive (this requires a polynomial number of bits), together with a certificate that this instance is indeed positive (see below). To check this certificate, the algorithm builds the corresponding instance of EEP-SURF (as done in Section 4—this takes polynomial time) and checks the certificate.

Here is a way to provide a certificate for an instance of EEP-SURF . In Section 5, we have proved that, if we have an instance (G, H, Π, \mathcal{S}) of EEP-SURF , then there exists a cellular embedding Γ' (in the form of a combinatorial map) of a graph G' containing G , and such

that Γ' extends Π . Moreover, G' is obtained from G by adding a number of edges that is $O(c)$, where c is the size of the original complex. (Recall that in the instance of EEP-SURF, the size of H is $O(c)$.) The cellular embedding Γ' of G' , given as a combinatorial map, is the certificate that (G, H, Π, \mathcal{S}) is positive: Given (G, H, Π, \mathcal{S}) and this certificate, we can in polynomial time check that G' contains G , that the restriction of Γ' to H is indeed Π , and that the combinatorial map of Γ' is indeed an embedding on \mathcal{S} . ◀

Acknowledgement

We would like to thank the referees whose detailed comments allowed us to improve the presentation of this article.

References

- 1 Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Vít Jelínek, Jan Kratochvíl, Maurizio Patrignani, and Ignaz Rutter. Testing planarity of partially embedded graphs. *ACM Transactions on Algorithms (TALG)*, 11(4):32, 2015. doi:10.1145/2629341.
- 2 Dan Archdeacon. Topological graph theory. A survey. *Congressus Numerantium*, 115:5–54, 1996.
- 3 William W. Boone. The word problem. *Annals of Mathematics*, 70:207–265, 1959. URL: <https://www.jstor.org/stable/1970103>.
- 4 Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner. Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. *SIAM Journal on Computing*, 43(5):1728–1780, 2014. doi:10.1137/120899029.
- 5 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. URL: <http://www.arxiv.org/abs/1702.05358>.
- 6 Éric Colin de Verdière and Arnaud de Mesmay. Testing graph isotopy on surfaces. *Discrete & Computational Geometry*, 51(1):171–206, 2014. doi:10.1007/s00454-013-9555-4.
- 7 Éric Colin de Verdière and Thomas Magnard. An FPT algorithm for the embeddability of graphs into two-dimensional simplicial complexes. In *Proceedings of the 29th European Symposium on Algorithms (ESA)*, pages 32:1–32:17, 2021. doi:10.4230/LIPIcs.ESA.2021.32.
- 8 Erik Demaine, Shay Mozes, Christian Sommer, and Siamak Tazari. Algorithms for planar graphs and beyond, 2011. Course notes. URL: <http://courses.csail.mit.edu/6.889/fall11/lectures/>.
- 9 Tamal K. Dey, Herbert Edelsbrunner, and Sumanta Guha. Computational topology. In Bernard Chazelle, Jacob E. Goodman, and Richard Pollack, editors, *Advances in Discrete and Computational Geometry – Proc. 1996 AMS-IMS-SIAM Joint Summer Research Conf. Discrete and Computational Geometry: Ten Years Later*, number 223 in Contemporary Mathematics, pages 109–143. AMS, 1999. URL: <http://pub.ist.ac.at/~edels/Papers/1999-B-01-ComputationalTopology.pdf>.
- 10 Tamal K. Dey and Sumanta Guha. Transforming curves on surfaces. *Journal of Computer and System Sciences*, 58:297–325, 1999. doi:10.1006/jcss.1998.1619.
- 11 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. URL: <https://dl.acm.org/doi/10.5555/644108.644208>.
- 12 Jeff Erickson. Combinatorial optimization of cycles and bases. In Afra Zomorodian, editor, *Computational topology*, Proceedings of Symposia in Applied Mathematics. AMS, 2012. URL: <https://jeffe.cs.illinois.edu/pubs/pdf/optcycles.pdf>.

- 13 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. doi:10.1137/1.9781611973105.118.
- 14 John Hopcroft and Robert Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974. doi:10.1145/321850.321852.
- 15 Martin Juvan, Jože Marinček, and Bojan Mohar. Elimination of local bridges. *Mathematica Slovaca*, 47:85–92, 1997. URL: <https://www.sfu.ca/~mohar/Papers/Local.pdf>.
- 16 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. doi:10.1109/FOCS.2008.53.
- 17 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. doi:10.1145/1250790.1250848.
- 18 Tomasz Kociumaka and Marcin Pilipczuk. Deleting vertices to graphs of bounded genus. *Algorithmica*, 81:3655–3691, 2019. doi:10.1007/s00453-019-00592-7.
- 19 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. doi:10.1109/FOCS.2012.12.
- 20 Sóstenes Lins. Graph-encoded maps. *Journal of Combinatorial Theory, Series B*, 32:171–181, 1982. doi:10.1016/0095-8956(82)90033-8.
- 21 Seth M. Malitz. Genus g graphs have pagenumbers $O(\sqrt{g})$. *Journal of Algorithms*, 17:85–109, 1994. doi:10.1006/jagm.1994.1028.
- 22 Jiří Matoušek, Eric Sedgwick, Martin Tancer, and Uli Wagner. Embeddability in the 3-sphere is decidable. *Journal of the ACM*, 65(1):5:1–5:49, 2018. doi:10.1145/3078632.
- 23 Jiří Matoušek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in R^d . *Journal of the European Mathematical Society*, 13(2):259–295, 2011. doi:10.4171/JEMS/252.
- 24 Bojan Mohar. On the minimal genus of 2-complexes. *Journal of Graph Theory*, 24(3):281–290, 1997.
- 25 Bojan Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM Journal on Discrete Mathematics*, 12(1):6–26, 1999. doi:10.1137/S089548019529248X.
- 26 Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 2001.
- 27 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. doi:10.1142/S0218195900000279.
- 28 Maurizio Patrignani. Planarity testing and embedding. In Roberto Tamassia, editor, *Handbook of graph drawing and visualization*. Chapman and Hall, 2006.
- 29 Neil Robertson and Paul D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.
- 30 John Stillwell. *Classical topology and combinatorial group theory*. Springer-Verlag, New York, second edition, 1993. URL: <https://link.springer.com/book/10.1007/978-1-4612-4372-4>.
- 31 Carsten Thomassen. The graph genus problem is NP-complete. *Journal of Algorithms*, 10(4):568–576, 1989. doi:10.1016/0196-6774(89)90006-0.