Embedding Graphs into Two-Dimensional Simplicial Complexes

Éric Colin de Verdière[1]
Université Paris-Est, LIGM, CNRS, ENPC, ESIEE Paris, UPEM, Marne-la-Vallée
France
eric.colindeverdiere@u-pem.fr

Thomas Magnard[2]
Université Paris-Est, LIGM, CNRS, ENPC, ESIEE Paris, UPEM, Marne-la-Vallée
France
thomas.magnard@u-pem.fr

Bojan Mohar[3]
Department of Mathematics, Simon Fraser University
Burnaby, Canada
mohar@sfu.ca
https://orcid.org/0000-0002-7408-6148

Abstract

We consider the problem of deciding whether an input graph G admits a topological embedding into a two-dimensional simplicial complex $\mathscr{C}$. 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 when $\mathscr{C}$ is part of the input, and we give a polynomial-time algorithm if the complex $\mathscr{C}$ is fixed.

Our strategy is to reduce the problem to 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).

2012 ACM Subject Classification Theory of computation → Computational geometry, Mathematics of computing → Graph algorithms, Mathematics of computing → Graphs and surfaces, Mathematics of computing → Algebraic topology

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

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.27

Related Version A full version of this paper is available at `https://arxiv.org/abs/1803.07032`.

---

[1] Part of this work was done while at Département d'Informatique, École normale supérieure, Paris, France. This author is supported in part by grant ANR-17-CE40-0033 of the French National Research Agency ANR (SoS project).

[2] Part of this work was done while at Département d'Informatique, École normale supérieure, Paris, France. This author is supported in part by grant ANR-17-CE40-0033 of the French National Research Agency ANR (SoS project).

[3] Part of this work was done while this author was invited professor at Département d'Informatique, École normale supérieure, Paris, France. This author is 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).

34th International Symposium on Computational Geometry (SoCG 2018).
Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 27; pp. 27:1–27:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

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 [8, 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 tog1803.07032ether.

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 [25]. The case $d = 3$ is much harder, and has only recently been shown to be decidable by Matoušek, Sedgwick, Tancer, and Wagner [19]. The general problem for arbitrary d has been extensively studied in the last few years, starting with hardness results by Matoušek, Tancer, and Wagner [20], and continuing with some algorithmic results in a series of articles; we refer to Matoušek et al. [19, 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 [23] and algorithms for surface-embedded graphs [7, 11]: The problem is NP-complete, as proved by Thomassen [28], but Mohar [22] 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 [13, 15]). 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 [21]. More recently, Čadek, 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 need not 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, Kratochvìl, 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 $\mathscr{C}$ (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 n vertices and edges, and a 2-complex $\mathscr{C}$ with c simplices.
Question: Does G have a topological embedding into $\mathscr{C}$?

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

▶ **Theorem 1.** *The problem* Embed *is NP-complete.*

▶ **Theorem 2.** *The problem* Embed(n, c) *can be solved in time* $f(c) \cdot n^{O(c)}$*, where* f *is some computable function of* c*.*

As for Theorem 1, it is straightforward that the problem is NP-hard (as the case where $\mathscr{C}$ is a surface is already NP-hard); the interesting part is to provide a certificate checkable in polynomial time when an embedding exists. Note that Theorem 2 shows that, for every fixed complex $\mathscr{C}$, the problem of deciding whether an input graph embeds into $\mathscr{C}$ is polynomial-time solvable. Actually, our algorithm is explicit, in the sense that, if there exists an embedding of G on $\mathscr{C}$, we can provide some representation of such an embedding (in contrast to some results in the theory of graph minors, where the existence of an embedding can be obtained without leading to an explicit construction).

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 $\mathscr{S}$, removing or contracting any edge yields a graph embeddable on $\mathscr{S}$. By Robertson and Seymour's theory, this immediately implies a cubic-time algorithm to test whether a graph G embeds on $\mathscr{S}$, for every fixed surface $\mathscr{S}$ [26]. 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. For example, let $\mathscr{C}$ 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 $\mathscr{C}$, 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 [9], Lazarus and Rivaud [16], and Erickson and Whittlesey [12]). In contrast, the homeomorphism problem for 2-complexes is as hard as graph isomorphism, as shown by Ó Dúnlaing, Watt, and Wilkins [24]. Moreover, the contractibility problem for closed curves on 2-complexes is undecidable; even worse, there exists a fixed 2-complex $\mathscr{C}$ such that the contractibility problem for closed curves on $\mathscr{C}$ 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 [27, Section 9.3]).

Despite this stark contrast between surfaces and 2-complexes, if we care only on 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-book, 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 [22].

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 single tetrahedron.

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 [14]. 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., Malik [18]) is also an embedding of G into a particular 2-complex, with additional constraints on the embedding.

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 2). Only at the end (Section 8) 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 $\mathscr{C}$ contains a 3-book (Section 3). From Section 4 onwards, we reduce Embed to *embedding extension problems* (EEP), similar to the Embed problem except that an embedding of a subgraph H of the input graph G is already specified. In Section 4, we reduce Embed to EEPs on a pure 2-complex (in which every segment of the complex $\mathscr{C}$ is incident to at least one triangle). In Section 5, we further reduce it to EEPs on a surface. In Section 6, we reduce it to EEPs on a surface in which every face of the subgraph H is a disk. Finally, in Section 7, we show how to solve EEPs of the latter type using a key component in an algorithm by the third author [22] to decide embeddability of a graph on a surface.

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 $\mathscr{C}$ corresponds naturally to a topological space, obtained in the obvious way: Start with one point per node in $\mathscr{C}$; connect them by segments as indicated by the segments in $\mathscr{C}$; similarly, for every triangle in $\mathscr{C}$, create a triangle whose boundary is made of the three segments contained in that triangle. By abuse of language, we identify $\mathscr{C}$ 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 $\mathscr{C}$ is an injective continuous map from (the topological space associated to) G to (the topological space associated to) $\mathscr{C}$.

2.2 Structural aspects of 2-complexes

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

Let p be a node of $\mathscr{C}$. A **cone at** p is a cyclic sequence of triangles $t_1, \dots, t_k, t_{k+1} = t_1$, 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 triangles $t_1, \dots, t_k$, 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, and any other pair of triangles have only p in common. An **isolated segment at** p is a segment incident to p but not incident to any triangle.

If $\mathscr{C}$ contains no 3-books, the set of segments and triangles incident with a given node p of $\mathscr{C}$ 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 Embedding extension problems and reductions

An **embedding extension problem** (EEP) is a decision problem defined as follows:

EEP(n, m, c):
Input: A graph G with n vertices and edges, a subgraph H of G with m vertices and edges, and an embedding Π of H into a 2-complex $\mathscr{C}$ with c simplices.
Question: Does G have an embedding into $\mathscr{C}$ whose restriction to H is Π?

To be precise, 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 complex $\mathscr{C}$) 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.

We will reduce our original problem to more and more specialized EEPs. We will 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 8.

2.4 Surfaces

In Section 6, we will assume some familiarity with surface topology; see, e.g., [5, 23, 27] for suitable introductions under various viewpoints. We recall some basic definitions and properties. A **surface** $\mathscr{S}$ is a compact, connected Hausdorff topological space in which every point has a neighborhood homeomorphic to the plane. 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, for an even, nonnegative number g 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 **genus** of $\mathscr{S}$; in this case, $\mathscr{S}$ is **non-orientable**.

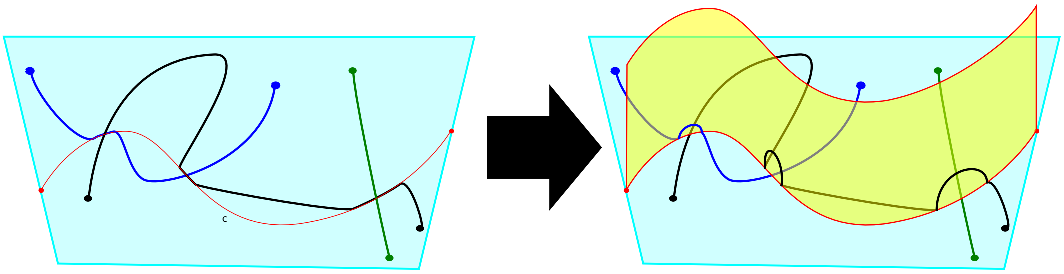

Figure 1 Illustration of the proof of Proposition 3. Left: The drawing of the graph G and the curve c (in thin). Right: The construction of the 3-book and the modification of the drawing.

A **surface with boundary** is obtained from a surface by removing a finite set of interiors of disjoint closed disks. The boundary of each disk forms a **boundary component** of $\mathscr{S}$. A **possibly disconnected surface** is a disjoint union of surfaces. An embedding of G into a surface $\mathscr{S}$, 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 $2-g-b=v-e+f$.

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

3 Reduction to complexes containing no 3-book

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

▶ **Proposition 3.** *If $\mathscr{C}$ contains a 3-book, then every graph embeds into $\mathscr{C}$.*

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'. 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 $\mathscr{C}$, because $\mathscr{C}$ contains a 3-book. ◀

4 Reduction to EEPs on a pure 2-complex

Our next task is to reduce the problem Embed to a problem on a pure 2-dimensional complex. More precisely, let **EEP-Sing** be the problem EEP, restricted to instances $(G,H,\Pi,\mathscr{C})$ where: $\mathscr{C}$ is a pure 2-complex containing no 3-books; H is a set of vertices of G; and Π is an injective map from H to the nodes of $\mathscr{C}$ such that $\Pi(H)$ contains all singular nodes of $\mathscr{C}$. 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 $f : P \to V(G) \cup \{\varepsilon\}$, where P is a set of nodes in $\mathscr{C}$ containing all singular nodes of $\mathscr{C}$. We say that an embedding Γ of G **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 n vertices and edges, a 2-complex $\mathscr{C}$ (not necessarily pure) containing no 3-books, with c simplices, and a map f as above, with domain of size m.
QUESTION: Does G have an embedding into $\mathscr{C}$ respecting f?

▶ **Lemma 5.** *Any instance of* EMBED(n, c) *reduces to* $(O(cn))^c$ *instances of* EMBED-RESP(cn, c, c).

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

Indeed, assume that G has an embedding Γ on $\mathscr{C}$. Each time an edge of G is mapped, under Γ, to a singular node p of $\mathscr{C}$, 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 $\mathscr{C}$ such that only vertices are mapped on singular nodes. The reverse implication is obvious: If G' has an embedding into $\mathscr{C}$, then so has G. This proves the claim.

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

▶ **Lemma 6.** EMBED-RESP(n, m, c) *reduces to* EEP-SING$(n, m, O(c))$.

Proof. We omit the proof due to space constraints, but the idea is simple:

Because we restrict ourselves to embeddings that respect f, and thus specify which vertex of G is mapped to each of the singular nodes of $\mathscr{C}$, what happens on the isolated segments of $\mathscr{C}$ is essentially determined. ◀

Before giving the proof, we first note:

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

5 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-book 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.

Let **EEP-Surf** be the problem EEP, restricted to instances where the input complex is (homeomorphic to) a possibly disconnected triangulated surface without boundary (which we denote by $\mathscr{S}$ instead of $\mathscr{C}$, for clarity). To represent the embedding Π in such an EEP instance $(G, H, \Pi, \mathscr{S})$, 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

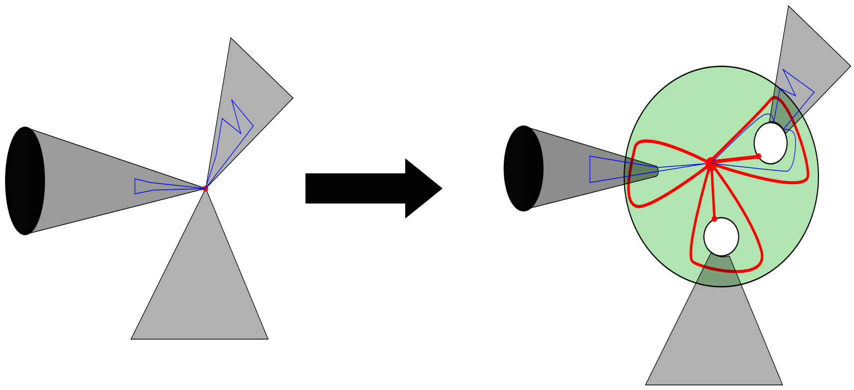

Figure 2 The modification of singular vertices in the proof of Lemma 8. We transform the neighborhood of each singular vertex to make it surface-like. Moreover, we add to H one loop per cone or corner; furthermore, for each corner, we add a vertex and an edge.

of $\mathscr{S}$ at least once, and finitely many times. (Note that H may use some nodes of $\mathscr{S}$.) Consider the *overlay* of the triangulation of $\mathscr{S}$ and of Π, the union of the 1-skeleton of $\mathscr{S}$ and of the image of Π; this overlay is the image of a graph on $\mathscr{S}$; each of its edges is either a piece of the image of an edge of H or a piece of a segment of $\mathscr{S}$; each of its vertices is the image of a vertex of H and/or a node of $\mathscr{S}$. By the assumption above on Π, this overlay is cellularly embedded on $\mathscr{S}$, and we can represent it by its combinatorial map [10,17] (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$(n + O(c), m + O(c), O(c))$ *in which the considered 2-complex is a possibly disconnected surface with boundary.*

Proof. The key property that we will use is that, since $\mathscr{C}$ is pure and contains no 3-books, each singular node is incident to cones and corners only.

Figure 2 illustrates the proof. Let $(G, H, \Pi, \mathscr{C})$ be the instance of EEP-Sing. We first describe the construction of the instance $(G', H', \Pi', \mathscr{S})$ 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 $\mathscr{C}$, 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 $\mathscr{S}$.

We now define H', G', and Π' from H, G, and Π (again, refer to Figure 2). Let p be a singular node of $\mathscr{C}$ 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 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 $\mathscr{C}$ with c simplices, we can easily obtain a triangulation of $\mathscr{S}$ with $O(c)$ simplices, and that the image of each edge of H crosses $O(1)$ edges of this triangulation.

There remains to prove that these two EEPs are equivalent; we omit the details. ◀

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. Let $(G, H, \Pi, \mathscr{S})$ be an instance of an EEP on a possibly disconnected surface with boundary. We first describe the construction of $(G', H', \Pi', \mathscr{S}')$, the EEP instance on a possibly disconnected surface without boundary.

Let $\mathscr{S}'$ be obtained from $\mathscr{S}$ by gluing a disk D_b along each boundary component. Let b be a boundary component of $\mathscr{S}$. 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. There remains to prove that these two EEPs are equivalent; we omit the details. ◀

Finally:

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

6 Reduction to a cellular EEP on a surface

Let **EEP-Cell** be the problem EEP, restricted to instances $(G, H, \Pi, \mathscr{S})$ where $\mathscr{S}$ is a surface and H is cellularly embedded and intersects each connected component of G.

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$(O(n + m + c), O(n + m + c), c)$.

As will be convenient also for the next section, we do not store an embedding Π of a graph G on a surface $\mathscr{S}$ by its overlay with the triangulation, as was done in the previous section, but 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 necessary 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).

6.1 Reduction to connected surfaces

We first build intermediary EEPs over connected surfaces. Let **EEP-Conn** be the problem EEP, restricted to instances $(G, H, \Pi, \mathscr{S})$ where $\mathscr{S}$ is a surface (connected and without boundary) and H intersects every connected component of G.

▶ **Lemma 11.** *Any instance of* EEP-SURF(n, m, c) *reduces to* $O(m^c)$ *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 1), any instance of EEP-SURF(n, m, c) is equivalent to the disjunction (OR) of $O(m^c)$ instances, each of them being the conjunction (AND) of $O(c)$ instances of EEP-CONN$(n, m + O(c), c)$.

Sketch of proof. We can embed each connected component of G that is planar and disjoint from H anywhere. There remains $O(c)$ connected components of G that are disjoint from H. For each of these, we choose a vertex, and we guess in which face of Π it belongs. We then know which connected component of the surface each connected component of G is mapped into. ◀

6.2 The induction

The strategy for the proof of Proposition 10 is as follows. For each EEP $(G', H', \Pi', \mathscr{S}')$ from the previous lemma, we will extend H' to make it cellular, 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 to prove that this process terminates.

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

$$\mathbf{cd}(H, \Pi, \mathscr{S}) := \sum_{f \text{ face of } \Pi} \text{genus}(f) + \sum_{f \text{ face of } \Pi} (\text{number of boundaries of } f - 1).$$

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 where G is empty. This implies that Π is a cellular embedding if and only if $cd(H, \Pi, \mathscr{S}) = 0$.

The following lemma reduces an EEP to EEPs with a smaller cellularity defect.

▶ **Lemma 12.** *Any instance of* EEP-CONN$(n, O(n), c)$ *reduces to* $O(n^4)$ *instances* $(G', H', \Pi', \mathscr{S})$ *of* EEP-CONN$(n + O(1), O(n), c)$ *where* $cd(H', \Pi', \mathscr{S}) < cd(H, \Pi, \mathscr{S})$.

The reduction does not depend on the size of H; furthermore, the new graph G' is obtained from the old one by adding exactly one edge and no vertex.

Admitting Lemma 12, the proof of Proposition 10 is straightforward:

Proof of Proposition 10. We first apply Lemma 11, obtaining $O(m^c)$ instances of EEP-CONN$(n, m + c, c)$. To each of these EEPs, we apply recursively Lemma 12 until we obtain cellular EEPs. The cellularity defect of the initial instance $(G, H, \Pi, \mathscr{S})$ is $O(m + c)$, being at most the genus of $\mathscr{S}$ plus $2m$, because each boundary component of a face of Π is incident to at least one edge of H (and each edge accounts for at most two boundary components in this way) or to one isolated vertex of H. Thus, the number of instances of EEP-CELL at the bottom of the recursion tree is $(n + m + c)^{O(m+c)}$, in which the size of the graph is $O(n + m + c)$ and the surface has at most c simplices. ◀

6.3 Proof of Lemma 12

There remains to prove Lemma 12. The proof uses some standard notions in surface topology, homotopy, and homology; we refer to textbooks and surveys [5, 23, 27]. We only consider homology with $\mathbb{Z}/2\mathbb{Z}$ coefficients.

Let f be a surface with a single boundary component and let p be a path with endpoints on the boundary of f. If we contract this boundary component to a single point, the path p becomes a loop, which can be null-homologous or non-null-homologous. We employ the same adjectives null-homologous and non-null-homologous for the path p. Recall that, if p is simple, it separates f if and only if it is null-homologous. The reversal of a path p is denoted by $\bar{p}$. The concatenation of two paths p and q is denoted by $p \cdot q$.

▶ **Lemma 13.** *Let f be a surface with boundary, let a be a point in the interior of f, and let a_1, a_2, and a_3 be points on the boundary of f. For each i, let p_i be a path connecting a_i to a. Let $r_1 = p_2 \cdot \bar{p}_3$, $r_2 = p_3 \cdot \bar{p}_1$, and $r_3 = p_1 \cdot \bar{p}_2$. Let P be a possible property of the paths r_i, among the following ones:*

- *"the endpoints of r_i lie on the same boundary component of f";*
- *"r_i is null-homologous" (if f has a single boundary component).*

Then the following holds: If both r_1 and r_2 have property P, then so does r_3.

Proof. This is a variant on the *3-path condition* from Mohar and Thomassen [23, 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 subpaths of the form $q \cdot \bar{q}$ from a path does not affect homology. ◀

Proof of Lemma 12. Since $cd(H, \Pi, \mathscr{S}) \geq 1$, there must be a face f of H with either (1) several boundary components, or (2) a single boundary component but positive genus. We will consider each of these cases separately, but first introduce some common terminology.

Let F be an arbitrary spanning forest of $G - E(H)$ rooted at $V(H)$. This means that F is 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 path between u and v in G that is the concatenation of $\overline{F(u)}$, edge uv, and $F(v)$.

Case 1: f has several boundary components. Assume that $(G, H, \Pi, \mathscr{S})$ 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 the path $F(uv)$ lies in f and connects two distinct boundary components of f.*

Indeed, let c be a curve drawn in f connecting two different boundary components of f. We can assume that it intersects the boundary of f exactly at its endpoints, at vertices of H. We can deform c 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 c is simple (except perhaps that its endpoints coincide on $\mathscr{S}$). Let $v_1, \ldots, v_k$ be the vertices of G encountered by c, in this order. We denote by $c[i, j]$ the part of c between vertices v_i and v_j. We claim that, for some i, we have that $\overline{F(v_i)} \cdot c[i, i+1] \cdot F(v_{i+1})$ connects two different boundary components of f: Otherwise, by induction on i, applying the first case of Lemma 13 to the three paths $\overline{c[1,i]}$, $F(v_i)$, and $c[i,i+1] \cdot F(v_{i+1})$, we would have that, for each i, $c[1,i] \cdot F(v_i)$ has its endpoints on the same boundary component of f, which is a contradiction for $i = k$ (for which the curve is c). So let i be such that $\overline{F(v_i)} \cdot c[i,i+1] \cdot F(v_{i+1})$ connects two different boundary components of f; letting $u = v_i$ and $v = v_{i+1}$, and embedding edge uv as $c[i,i+1]$, gives the desired embedding of G_{uv}. This proves the claim.

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 $(G_{uv}, H \cup F(uv), \Pi', \mathscr{S})$ where Π' is chosen as an appropriate extension of Π. *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 each occurrence of $r(u)$ and $r(v)$ on the boundary of f, we consider the EEP $(G_{uv}, H \cup F(uv), \Pi', \mathscr{S})$ 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 $(G, H, \Pi, \mathscr{S})$ 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 $(G_{uv}, H \cup F(uv), \Pi'', \mathscr{S})$ 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, $(G_{uv}, H \cup F(uv), \Pi', \mathscr{S})$, has a solution as well. These two EEPs are the same except that 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. 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 $(G_{uv}, H \cup F(uv), \Pi'', \mathscr{S})$ to a solution of $(G_{uv}, H \cup F(uv), \Pi', \mathscr{S})$, as desired.

It also follows from the previous paragraph that the cellularity defect decreases by one. To conclude this case, we note that the number of new EEPs is $O(n^4)$: indeed, there are $O(n)$ possibilities for the choice of u (or v), and $O(n)$ possibilities for the choice of the occurrence of $r(u)$ (or $r(v)$) on the boundary of f.

If 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 very similar to the previous case, the main difference being that, instead of paths in f connecting different boundary components of f, we now consider paths in f that are non-null-homologous. ◀

7 Solving a cellular EEP on a surface

▶ **Proposition 14.** *There is a function* $f : \mathbb{N} \to \mathbb{N}$ *such that every instance of* EEP-Cell$(n, O(n), c)$ *can be solved in time* $f(c) \cdot O(n)$*.*

Proof. This is essentially the main result of Mohar [22]. The algorithm in [22] makes reductions to even more specific EEPs, and one feature that is needed for bounding the number of new instances is that the embedded subgraph H satisfies some connectivity assumptions. ◀

8 Proof of Theorems 1 and 2

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

Proof of Theorem 2. This immediately follows from Propositions 3, 4, 7, 10, and 14. ◀

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

Proof of Theorem 1. The problem Embed is NP-hard because deciding whether an input graph embeds into an input surface is NP-hard [28]. It is in NP because (assuming $\mathscr{C}$ contains no 3-books), a certificate that an instance is positive is given by a certificate that some instance of EEP-Surf is positive (see the proof of Proposition 7). Such an instance $(G, H, \Pi, \mathscr{S})$ has a certificate, given by the combinatorial map of a supergraph G' of G cellularly embedded on $\mathscr{S}$ (see Section 6), that can be checked in polynomial time. ◀

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.

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.

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.

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. See `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.

7 Erik Demaine, Shay Mozes, Christian Sommer, and Siamak Tazari. Algorithms for planar graphs and beyond, 2011. Course notes available at `http://courses.csail.mit.edu/6.889/fall11/lectures/`.

8 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.

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

10 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.

11 Jeff Erickson. Combinatorial optimization of cycles and bases. In Afra Zomorodian, editor, *Computational topology*, Proceedings of Symposia in Applied Mathematics. AMS, 2012.

12 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.

13 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.

14 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.

15 Tomasz Kociumaka and Marcin Pilipczuk. Deleting vertices to graphs of bounded genus. arXiv:1706.04065, 2017.

16 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.

17 Sóstenes Lins. Graph-encoded maps. *Journal of Combinatorial Theory, Series B*, 32:171–181, 1982.

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

19 Jiří Matoušek, Eric Sedgwick, Martin Tancer, and Uli Wagner. Embeddability in the 3-sphere is decidable. *Journal of the ACM*, 2018. To appear. Preliminary version in *Symposium on Computational Geometry 2014*.

20 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.

21 Bojan Mohar. On the minimal genus of 2-complexes. *Journal of Graph Theory*, 24(3):281–290, 1997.

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

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

24 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.

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

26 Neil Robertson and Paul D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995.

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

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