Testing Graph Isotopy on Surfaces

Éric Colin de Verdière · Arnaud de Mesmay

Received: 11 October 2012 / Revised: 2 August 2013 / Accepted: 11 October 2013 /
Published online: 30 October 2013
© Springer Science+Business Media New York 2013

Abstract We investigate the following problem: Given two embeddings G_1 and G_2 of the same abstract graph G on an orientable surface S , decide whether G_1 and G_2 are isotopic; in other words, whether there exists a continuous family of embeddings between G_1 and G_2 .

We provide efficient algorithms to solve this problem in two models. In the first model, the input consists of the arrangement of G_1 (resp., G_2) with a fixed graph cellularly embedded on S ; our algorithm is linear in the input complexity, and thus, optimal. In the second model, G_1 and G_2 are piecewise-linear embeddings in the plane, minus a finite set of points; our algorithm runs in $O(n^{3/2} \log n)$ time, where n is the complexity of the input.

The graph isotopy problem is a natural variation of the homotopy problem for closed curves on surfaces and on the punctured plane, for which algorithms have been given by various authors; we use some of these algorithms as a subroutine.

As a by-product, we reprove the following mathematical characterization, first observed by Ladegaillierie (Topology 23:303–311, 1984): Two graph embeddings are isotopic if and only if they are homotopic and congruent by an oriented homeomorphism.

Keywords Computational topology · Embedded graph · Homotopy · Isotopy · Topological graph theory

É. Colin de Verdière
Département d'informatique, École normale supérieure, CNRS, Paris, France
e-mail: eric.colin.de.verdiere@ens.fr

A. de Mesmay (✉)
Département d'informatique, École normale supérieure, Paris, France
e-mail: arnaud.de.mesmay@ens.fr

1 Introduction

Deciding whether a curve in a space can be deformed to another one is a fundamental topological task. Since the problem is undecidable in general, even in two-dimensional simplicial complexes and 4-manifolds [33, pp. 242–247], one has to put restrictions on the ambient space. The case of surfaces has been much investigated recently (see references below), since it is at the same time non-trivial, interesting, and tractable, and since the underlying mathematics, notably in the field of combinatorial group theory, are well understood.

This paper also studies deformations on surfaces, but instead of curves, we consider graphs drawn on surfaces: Given two embeddings (crossing-free drawings) of a graph G on a surface, can we deform one continuously to the other without introducing intersections between edges during the process?

As a motivating special case, consider a finite set of obstacle points P in the plane and a graph G embedded on $\mathbb{R}^2 \setminus P$ in two different ways, G_1 and G_2 . Does there exist a “morph” between G_1 and G_2 (possibly moving the vertices and bending the edges) that avoids passing over any obstacle? To the best of our knowledge, no algorithm is known for this purpose. This is relevant for morphing applications: To compute a morphing between two images, it is helpful to first build a deformation between compatible graphs representing the most salient features of the images. In such applications, it is sometimes desirable to add some topological requirements on the morphing, e.g., to force some area of the deforming image to always cover a fixed point of the plane during the deformation. Such requirements can be encoded using obstacle points, since a face of the graph containing an obstacle point has to contain it during the whole deformation.

Another motivation for this problem comes from geographic information systems and map simplification. When simplifying a road network, it is crucial that the features of the map (cities, mountains) stay on the same side of the roads. This can be tested by considering these features as obstacle points and testing whether the road networks obtained before and after straightening are isotopic; see, e.g., Cabello et al. [6].

More generally, assume that we have a triangulated surface in $\mathbb{R}^3$, and two embeddings G_1 and G_2 of the same graph G on that surface (not necessarily on the skeleton of the triangulation). Each graph G_i is encoded by its combinatorial arrangement with the triangulation. Can we continuously move G_1 to G_2 ? In this setting, the graphs G_1 and G_2 might represent textures on the surface, and the question is whether one can continuously move one texture so that it coincides with the other.

1.1 Related Work

A *homotopy* between two paths is a continuous family of paths with the same end-points between them. A *homotopy* between two cycles (closed curves) is a continuous family of cycles between them.¹ Dey and Guha [12], and later Lazarus and

¹This is sometimes referred to as *free homotopy* in textbooks, in opposition to *fixed point homotopy*, which we will only use in Sect. 6.

Rivaud [28], as well as Erickson and Whittlesey [17], study the problem of deciding whether two paths or cycles on a surface S are homotopic. Both problems can be solved in optimal linear time if the input curves are represented as walks in a graph embedded on S . Cabello et al. [6] give efficient algorithms for testing homotopy of paths in the special case where the surface S is a punctured plane (a plane, minus a finite set of obstacle points) and the input paths are represented by polygonal paths in the plane.

A related problem, to which a large body of research is devoted, is that of computing shortest homotopic paths or cycles. It was first studied by Hershberger and Snoeyink [23] for a triangulated surface where the vertices lie on the boundary, and revisited by Efrat et al. [13] and Bespamyatnikh [2] for paths in a punctured plane. These algorithms, in particular, allow to decide whether two paths are homotopic, and have a better complexity than the algorithm by Cabello et al. [6] in some cases. Results on the shortest homotopic path/cycle problem are also known for combinatorial surfaces [9, 10]: It is solvable in polynomial time [8].

An *isotopy* between two simple curves (paths or cycles without self-intersections) is a homotopy that keeps the deforming curve simple at each step. If two simple curves are isotopic, they are homotopic. The converse does not hold in general [19], though it holds for simple paths with endpoints on the boundary and for simple cycles not bounding a disk [15]; in these cases, it is equivalent to test homotopy or isotopy. More generally, an *isotopy with fixed vertices* between two graph embeddings is a continuous family of graph embeddings between them, where the vertices are not allowed to move. It is known that one can compute shortest graph embeddings with fixed vertices in polynomial time [7, 8], and similar techniques allow to test whether two graph embeddings are isotopic with fixed vertices in polynomial time.

Two homeomorphisms from S to S are *isotopic* if there is a continuous family of homeomorphisms connecting them. The *mapping class group* of a surface (without boundary) S is, roughly, the set of isotopy classes of all orientation-preserving homeomorphisms from S to S . See, e.g., Farb and Margalit [18] for a recent and exhaustive survey on this topic. Although we use little of this vast theory, it is quite connected to our problem: If G_1 and G_2 are cellularly embedded on S , a homeomorphism of S that maps G_1 to G_2 represents a unique element of the mapping class group, and our problem amounts to testing whether this element is the identity. Hence it is closely related to the *word problem* in mapping class groups, which can be solved in quadratic time [21, 31]. However, these algorithms take as input the collection of Dehn twists corresponding to a mapping class, which is an input incomparable to ours. Note that if S is an n -punctured sphere, the mapping class group of S is the pure braid group with n strands, which has garnered considerable algorithmic attention in recent years, in particular due to its possible applications to cryptography [26].

1.2 Our Results

In this paper, we study the problem of deciding whether two embeddings G_1 and G_2 of the same abstract graph G on a surface S are isotopic; in other words, whether there exists a continuous family of embeddings of G (possibly moving the vertices of G) connecting G_1 with G_2 .

In more detail, the input to the algorithm is a description of the surface S , the graph G , and the graph embeddings G_1 and G_2 . All surfaces are assumed to be compact, connected, and orientable; they may have boundary. An embedding maps each vertex (or edge, or halfedge) of G to the corresponding feature on the surface. In particular, the correspondence between the vertices (or edges, or halfedges) of G_1 and G_2 is given.

Our algorithmic results come in two flavors, depending on the model used. In the general model, G_1 and G_2 are embeddings of G on an arbitrary surface S . To represent them, we use a model similar to the *cross-metric surface model* [8]: A graph embedding is represented by the intersections it forms with a given cellularly embedded graph. More specifically, we assume that S has a fixed graph H cellularly embedded on it (i.e., each face of H is homeomorphic to a disk; for example, H may be a triangulation of S); the input to the algorithm is the combinatorial map of the arrangement of G_1 with H (resp., G_2 with H).² We emphasize that the input does not consider the crossings between G_1 and G_2 . We give a linear-time (and thus optimal) algorithm to decide whether G_1 and G_2 are isotopic:

Theorem 1 *Let S be an orientable surface, possibly with boundary. Let H be a fixed graph cellularly embedded on S . Let G_1 and G_2 be two graph embeddings of the same graph G on S , each in general position with respect to H . Given the combinatorial map of the arrangement of G_1 with H (resp., G_2 with H), of complexity k_1 (resp., k_2), we can determine whether G_1 and G_2 are isotopic in $O(k_1 + k_2)$ time.*

Let us emphasize that the surface S is not fixed in this result; the constant in the $O(\cdot)$ notation does not depend on S .

We also study the complexity of the problem in the case where S is the plane, minus a finite set P of obstacle points, and G_1 and G_2 are piecewise-linear graph embeddings of G in $\mathbb{R}^2 \setminus P$. In this case, the input is the point set P together with the embeddings G_1 and G_2 , where each edge of each embedding is represented as a polygonal path. (Again, the embeddings G_1 and G_2 may intersect arbitrarily.) In this setting, we prove the following result.

Theorem 2 *Let P be a set of p points in the plane, and let G_1 and G_2 be two piecewise-linear graph embeddings of the same graph G in $\mathbb{R}^2 \setminus P$, of complexities (number of segments) k_1 and k_2 respectively. We can determine whether G_1 and G_2 are isotopic in $\mathbb{R}^2 \setminus P$ in time $O(n^{3/2} \log n)$ time, where n is the total size of the input. In more detail, the running time is, for any $\varepsilon > 0$,*

$$O((k_1 + p) \log(k_1 + p) + (k_2 + p) \log(k_2 + p) + \min\{(k_1 + k_2)p, p^{1+\varepsilon} + (k_1 + k_2)\sqrt{p} \log p\}).$$

Let us emphasize that the isotopy is a continuous family of topological embeddings; one may assume that all these embeddings are piecewise-linear, but we claim no upper bound on their complexities.

²i.e., the graph obtained by overlaying G_i and H .

Fig. 1 Two embeddings G_1 and G_2 of a one-vertex graph with three loop edges on the sphere with four punctures. These two embeddings are not isotopic, although there exists an oriented homeomorphism mapping one to the other, and the four cycles following the boundaries of the faces are homotopic in G_1 and G_2

1.3 Overview of the Techniques

If two graph embeddings G_1 and G_2 of the same graph G are isotopic, then clearly (1) there is an oriented homeomorphism of the surface that maps G_1 to G_2 ;³ and (2) if γ is a cycle in G (possibly with repeated vertices and edges), then its images in G_1 and G_2 are homotopic.

It was shown by Ladegaillierie [27] that such necessary conditions are, in fact, sufficient. However, the second condition is not algorithmic, since there are infinitely many cycles in G . A close inspection of Ladegaillierie's proof reveals that $O(g + b)$ pairs of cycles need to be tested for homotopy, where g and b denote the genus and number of boundary components of the surface. However, in Ladegaillierie's construction, the complexity of the family of cycles is not explicitly given. With some work, and using some of our techniques, one might be able to obtain an explicit algorithm using his construction, but the running time would certainly be larger, by at least an additional $O(g)$ factor (since Ladegaillierie's decomposition contains a pants decomposition), if not more.

Using a very different method, we reprove Ladegaillierie's characterization in a strengthened form (Theorem 4 below): We provide an explicit set of cycles Λ in G of linear overall complexity that have to be tested for homotopy. Our algorithmic results follow, since one can perform efficiently the homeomorphism test (this essentially amounts to checking equality of two combinatorial maps, see Sect. 2.2), as well as the test of homotopy between two given cycles (using known algorithms or variants of them [6, 12, 28]).

We note that it is not straightforward to find a suitable set Λ in G of overall linear complexity satisfying the above condition. In particular, a natural candidate for Λ would be the set of all facial cycles in G_1 (and thus in G_2). However, Fig. 1 shows that the condition is not fulfilled, even in the case where the surface is the sphere with four punctures.

The paper is organized as follows. After some preliminaries (Sect. 2), we characterize the existence of an isotopy between two *stable* arrangements of cycles, containing no disk with less than four vertices on its boundary (Sect. 3). We then prove

³assuming that G_1 and G_2 are in the interior of S , which is true by our general position assumption.

the aforementioned strengthened form of Ladegaillerie's result (Sect. 4); the idea for the proof is to compute from G_1 and G_2 two stable arrangements of cycles in their tubular neighborhoods, which are isotopic if and only if G_1 and G_2 are isotopic, and to apply the characterization for stable cycles. We then deduce our main computational results in Sect. 5. In Sect. 6, we briefly indicate how our algorithms and results extend to the case where we require some vertices of the graph to be fixed throughout the isotopy.

2 Preliminaries

Here, we review some topological notions. We refer the reader to Hatcher [22], Stillwell [33], or any previous paper on related topics [8] for standard background in algebraic topology used in the present paper, and Farb and Margalit [18] for the connection between covering spaces, hyperbolic geometry, and surface homeomorphisms.

2.1 Background

Curves and Graphs on Surfaces Henceforth, S is a compact, connected, orientable surface with genus g and b boundary components; its *Euler characteristic* is $\chi(S) = 2 - 2g - b$. A *path*, respectively a *cycle*, on a surface S is a continuous map $p : [0, 1] \rightarrow S$, respectively a continuous map $\gamma : S^1 \rightarrow S$, where S^1 is the unit circle. An *arc* is a path intersecting the boundary of S exactly at its endpoints. A curve denotes a path or a cycle. Paths and cycles are *simple* if they are one-to-one. We emphasize that, according to our definition, a cycle in a graph is a closed walk *possibly with repeated vertices and edges* (in contrast to the graph-theoretic notion).

An *embedding* of a graph G on a surface S is, informally, a crossing-free drawing of G on S . A graph embedding is *cellular* if its *faces*, namely, the connected components of the complement of the image of the graph, are homeomorphic to open disks. If G is embedded cellularly on a surface S and if we denote by v , e , and f the respective numbers of its vertices, edges and faces, we have $v - e + f = \chi(S) = 2 - 2g - b$.

A (finite) family of cycles on S is *in general position* if the cycles are in the interior of S , there are finitely many (self-)intersection points, and each intersection is a transverse crossing between exactly two pieces of cycles. Similarly, a (finite) family of graph embeddings on S is *in general position* if all embeddings are in the interior of S , there are finitely many intersection points between two different embeddings, and each intersection is a transverse crossing of the interiors of exactly two edges. Classical approximation techniques, see for example Epstein [15, Appendix], allow us to approximate every edge in a graph embedding by a piecewise linear edge using an ambient isotopy. In particular, all graph embeddings in this paper can be assumed to be piecewise-linear. By doing small perturbations if necessary, this allows us to assume that, moreover, all the graph embeddings we consider are in general position, and we will always make this assumption unless stated otherwise.

The simultaneous drawing of a family of cycles or of graph embeddings in general position on S gives a set of vertices (original vertices or intersection points) and edges (connecting two such vertices); this graph is called the *arrangement* of the family. (Here we also allow edges that are simple cycles without any vertex on them.)

Homeomorphisms and Isotopies At several occasions, we will consider a homeomorphism h that maps a cycle $c : S^1 \rightarrow S$ into another one, $c' : S^1 \rightarrow S$. Unless stated otherwise, this expression means that $h \circ c = c'$, namely, h *pointwise* maps c to c' . However, we will sometimes need weaker concepts. We say that h *maps* c to c' not necessarily pointwise, but only *as sets*, if they do so up to reparameterization; namely, if there is a homeomorphism $\varphi : S^1 \rightarrow S^1$ such that $h \circ c \circ \varphi = c'$. If furthermore the homeomorphism φ is increasing (intuitively, h maps c to the cycle c' not necessarily pointwise, but the orientations of $h(c)$ and c' are the same), we say that h *maps* c to c' not necessarily pointwise, but *preserving the orientations* of the cycles.

Often we will also have an abstract G and two embeddings G_1 and G_2 of G on a surface S ; we say that a homeomorphism h *maps* G_1 to G_2 if it maps each edge of G_1 (pointwise) to the corresponding edge of G_2 .

Two embeddings G_0 and G_1 of the same abstract graph G on S are *isotopic* if there is a continuous family of embeddings $(G_t)_{t \in [0,1]}$ between G_0 and G_1 . In more detail, the data of an embedding is given by the choice of a point for each vertex and a path connecting the appropriate vertices for each edge (with some conditions asserting that no crossing occurs); a family (G_t) of embeddings is continuous if all these maps vary continuously over t . (The vertices may, in particular, move.)

Two homeomorphisms h_0 and h_1 from S to S are *isotopic* if there is a continuous family of homeomorphisms $(h_t)_{t \in [0,1]}$ connecting them; the family (h_t) is called an (*ambient*) *isotopy* between them. By a common abuse of language, a homeomorphism isotopic to the identity is also called an (*ambient*) *isotopy*. For $A \subseteq S$, we say that a homeomorphism h is an ambient isotopy *relatively to* A if there is a continuous family of homeomorphisms between h and the identity such that each homeomorphism is the identity on A .

It holds that two smooth embeddings of the cycle S^1 on S are isotopic if and only if there exists an ambient isotopy of S mapping one to the other (provided the two embeddings are in the interior of S); this follows from tools in differential topology (e.g., vector fields); see for instance Farb and Margalit [18, Proposition 1.11] or Hirsch [24, Theorem 1.3]. In the case of (piecewise-linear) graph embeddings, the same tools might yield the same result, but we did not find any proof in the literature. It is not an issue though, since our proof also implies the statements of Theorems 1 and 2 in which graph isotopy is replaced with ambient isotopy (see Corollaries 3 and 4).

We will need the following basic lemma.

Lemma 1 (Alexander's Lemma; see, e.g., Farb and Margalit [18, Lemma 2.1]) *Let D be a disk and $h : D \rightarrow D$ be a homeomorphism fixed on the boundary of D . Then h is an ambient isotopy relatively to the boundary of D .*

Proof We can assume that D is the unit closed disk in the plane. The continuous family of embeddings between h and the identity can be explicitly defined by

$$F(x, t) = \begin{cases} (1-t)h(\frac{x}{1-t}) & \text{if } 0 \leq |x| \leq 1-t, \\ x & \text{if } 1-t \leq |x| \leq 1, \end{cases}$$

for $0 \leq t < 1$, and $F(x, 1) = x$ for each $x \in D$. □

Covering Spaces and Hyperbolic Geometry If S has a negative Euler characteristic and has no boundary, it can be provided with a hyperbolic metric; this naturally induces a hyperbolic metric on its universal cover $\tilde{S}$, which is then isometric to the open hyperbolic disk D_2 . This open disk is classically compactified with a boundary ∂D_2 [18, Chap. 1], and a non-contractible cycle on S lifts into an arc in $D_2 \cup \partial D_2$. We call its intersections with the boundary the *endpoints* of the lift or, according to the orientation of the lift, its *source* and its *target*. Two such lifts in D_2 have common endpoints if and only if they stay at a bounded distance from each other. If S has boundary, $\tilde{S}$ is isometric to a totally geodesic subspace of D_2 .

It is folklore [5, Lemma 1.6.5] that if two cycles on a surface S are homotopic, they have lifts in D_2 with the same endpoints. We also recall that on a hyperbolic surface, every cycle is homotopic to a unique geodesic [18, Proposition 1.3].

2.2 Combinatorial Maps for Non-cellular Embeddings

To each embedding G_1 of an abstract graph G on S corresponds a *combinatorial map*, storing the combinatorial information of that embedding. For graphs cellularly embedded, this notion is well known, and there are several essentially equivalent data structures available to store such combinatorial maps on surfaces [25], like the gem representation [14, 29] and the halfedge data structure of CGAL.⁴ All these data structures have linear complexity in the number of edges of the graph. The gem representation stores four *flags* per edge; intuitively, if the edge is oriented, two flags close to its head, one to its left and one to its right, and similarly for the tail. More formally, a flag represents an incidence between a vertex, an edge, and a face of the embedding. Three involutions allow to move from flag to an “incident” flag in the graph: The first one, i_v , keeps the same edge-face incidence and moves to the opposite vertex; the second one, i_e , keeps the same vertex-face incidence and moves to the opposite edge; the last one, i_f , keeps the same vertex-edge incidence and moves to the opposite face. Also, each flag has a pointer to the underlying vertex, edge, and face of G .

Two combinatorial maps of the same graph G are *isomorphic* if there is a bijection φ between their sets of flags F_1 and F_2 , commuting with the three involutions, and such that the underlying vertex (resp., edge) of a flag f in F_1 is the same as that of $\varphi(f)$; such a bijection is called a *map isomorphism*. It is not hard to test whether two combinatorial maps representing two cellular embeddings G_1 and G_2 of the same abstract graph G on S are isomorphic in linear time. Furthermore, we have the following lemma.

Lemma 2 *The combinatorial maps representing two cellular embeddings G_1 and G_2 of the same abstract graph G on S are isomorphic if and only if there exists a homeomorphism of S mapping G_1 to G_2 .*

Proof If there exists a homeomorphism of S mapping G_1 to G_2 , then obviously the combinatorial maps of G_1 and G_2 are isomorphic. Conversely, any isomorphism

⁴www.cgal.org.

Fig. 2 An example of extended combinatorial map, with the three involutions i_v , i_e , and i_f for a given flag (*in bold*), and the pointers p_1 , p_2 and p_3 between a face and one flag for each incident cycle

of combinatorial maps extends naturally to a homeomorphism between the tubular neighborhoods of the graphs. (Informally, one can build a disk for each vertex and a strip for each edge of G , and attach these disks and strips as prescribed by the combinatorial map of G_1 , so that their union forms a tubular neighborhood of G_1 . And one can similarly do the same for G_2 , the gluings being combinatorially the same since the combinatorial maps are isomorphic.) This homeomorphism between the two tubular neighborhoods extends by radial extension in every disk, and since all the faces are disks, we obtain therefore a homeomorphism of the whole surface S . $\square$

We adapt the gem representation (Fig. 2) to handle graphs that are not cellularly embedded, by adding the following information. (For simplicity of exposition, we only consider graphs without isolated vertices; it is not hard to extend the data structure to handle this case.) First, each face may be incident to several cycles of G , and stores a list containing one flag of each such cycle. Conversely, each flag has a pointer to the face it belongs to. Also, a face may have non-zero genus and contain some boundary components of S , so we store this genus and number of boundary components of S .

Two such *extended combinatorial maps*⁵ are *isomorphic* if, in addition to the conditions above for standard combinatorial maps, the corresponding faces have the same genus and the same number of boundary components of the surface. The corresponding bijection will be called an *extended map isomorphism*. As in the cellular case, one can check whether two extended combinatorial maps of two embeddings G_1 and G_2 of G are isomorphic in linear time in their complexity. And as in the standard case, we have the following lemma.

Lemma 3 *Two extended combinatorial maps of G_1 and G_2 are isomorphic if and only if there exists a homeomorphism of the surface mapping G_1 to G_2 .*

⁵We will drop the adjective “extended” when it is obvious from the context.

Proof The proof is similar to that of Lemma 2. The isomorphism extends to a homeomorphism between the neighborhoods of the graphs G_1 and G_2 ; since each face of G_1 has the same topology as the corresponding one in G_2 , this homeomorphism extends naturally to the whole surface. $\square$

Our construction does not depend on whether S is orientable. If S is oriented, a flag has a natural *orientation*, depending on whether we turn clockwise or counterclockwise around the vertex of the flag when starting on the edge of the flag, close to the vertex, and moving towards the face of the flag. Furthermore, each of the three involutions reverses the orientation. An *orientation* of an extended combinatorial map assigns an orientation “clockwise” or “counterclockwise” to each flag such that each involution reverses the orientation of the flag. There exists an *orientation-preserving isomorphism* between two oriented extended combinatorial maps if they are isomorphic as extended combinatorial maps and the bijection between the flags preserves the orientation. This can also be tested in linear time. Equivalently, there exists an oriented homeomorphism of the surface mapping one graph embedding to the other.

3 Isotopies of Stable Families of Cycles

Let Γ be a family of cycles in general position on S . A k -gon in Γ , for $k \geq 1$, is an open disk on S whose boundary is formed by exactly k subpaths of Γ . A 0-gon is a disk whose boundary is a single simple cycle in Γ . In general, k -gons may contain and may be crossed by other pieces of the arrangement of Γ ; if this is not the case, we say that the k -gon is *empty*.

We say that Γ is *stable* if its arrangement contains no *empty* k -gon for $k \leq 3$. In this section, we prove the following result.

Theorem 3 *Let S be an orientable surface and let $\Gamma_1 = (\gamma_{1,1}, \dots, \gamma_{1,n})$ and $\Gamma_2 = (\gamma_{2,1}, \dots, \gamma_{2,n})$ be two stable families of cycles on S in general position such that:*

1. *There exists an oriented homeomorphism h of S mapping each cycle $\gamma_{1,j}$ of Γ_1 to the corresponding cycle $\gamma_{2,j}$ of Γ_2 not necessarily pointwise, but preserving the orientations of the cycles, and*
2. *Each cycle of Γ_1 is homotopic to the corresponding cycle of Γ_2 .*

Then there is an isotopy of S that maps each cycle of Γ_1 to the corresponding cycle of Γ_2 , not necessarily pointwise, but preserving the orientations of the cycles.

We note that conversely, if there exists an ambient isotopy of S mapping Γ_1 to Γ_2 , then conditions (1) and (2) are satisfied.

As a side remark, we will actually use this result in a setting where we know that h maps each cycle of Γ_1 *pointwise* to the corresponding cycle of Γ_2 ; however, even under this stronger hypothesis, it does not always hold that there exists an isotopy of S that maps each cycle in Γ_1 *pointwise* to the corresponding cycle of Γ_2 ; see Fig. 3. The following corollary adds another hypothesis to ensure that this does not happen. It will be not used directly in this paper, but inspires the techniques used in Sect. 4.

Fig. 3 A stable family comprised of a single cycle $\gamma : S^1 \rightarrow S$, embedded in two different ways on a double torus. Letter A denotes the image of the same point of S^1 in both embeddings, and similarly for B . Both cycles are homotopic, some oriented homeomorphism of S maps one to the other pointwise, and there exists an isotopy of S (namely, the identity) that maps one to the other, preserving their orientation; however, there exists no isotopy of S that maps one to the other pointwise

Corollary 1 *Let S be an orientable surface and let $\Gamma_1 = (\gamma_{1,1}, \dots, \gamma_{1,n})$ and $\Gamma_2 = (\gamma_{2,1}, \dots, \gamma_{2,n})$ be two stable families of cycles on S in general position such that:*

1. *There exists an oriented homeomorphism h of S mapping each cycle $\gamma_{1,j}$ of Γ_1 to the corresponding cycle $\gamma_{2,j}$ of Γ_2 not necessarily pointwise, but preserving the orientations of the cycles, and*
2. *Each cycle of Γ_1 is homotopic to the corresponding cycle of Γ_2 .*

If there is only one orientation-preserving map automorphism of Γ_1 , namely the identity, then there is an ambient isotopy of S mapping each cycle of Γ_1 to the corresponding cycle of Γ_2 pointwise.

Proof Theorem 3 shows that there is an ambient isotopy mapping each cycle of Γ_1 to the corresponding cycle of Γ_2 that is not necessarily pointwise, but preserves the orientation of each cycle. Hence this ambient isotopy induces an orientation-preserving map isomorphism between Γ_1 and Γ_2 . But the extended combinatorial maps of Γ_1 and Γ_2 are the same, since there exists an oriented homeomorphism between them. We thus get an orientation-preserving map automorphism of Γ_1 , and it has to be the identity. This means that each vertex of the arrangement of Γ_1 is mapped to the corresponding vertex of the arrangement of Γ_2 , and similarly for each edge of these arrangements (which are mapped with the same orientation). After an isotopy in a neighborhood of these edges, we get a pointwise ambient isotopy. $\square$

The remaining part of this section is devoted to the proof of Theorem 3 if the surface S has negative Euler characteristic. The proof for the remaining surfaces uses slightly different tools and is deferred to Sect. 7.

3.1 Basic Consequences of Euler's Formula

The following lemmas are simple consequences of Euler's formula.

Lemma 4 *Assume just for this lemma that S is a sphere. Let G be a connected (hence cellularly embedded) graph on S , such that every vertex has degree four. Then there are at least three faces in G with degree smaller than four.*

Proof Let v , e , and f denote the number of vertices, edges, and faces of G . The sum of the degrees of all faces is $2e$, which, by Euler's formula $v - e + f = 2$ and double-counting of the vertex-edge incidences $4v = 2e$, equals $4f - 8$. Since every face has degree at least one, if at most two faces have degree smaller than four, the sum of the degrees of all faces is at least $4f - 6$. This is a contradiction. $\square$

A trivial corollary is the following:

Corollary 2 *Assume just for this corollary that S is a sphere, a disk, or an annulus. Let G be a connected graph embedded on S , such that every vertex has degree four. Then there is at least one face in G with degree smaller than four.*

Proof Let $\bar{S}$ be the sphere obtained from S by attaching a disk to each of the boundary components of S . The graph G is embedded on $\bar{S}$, and by Lemma 4 it contains at least three faces (on $\bar{S}$) with degree smaller than four. At least one of these faces does not contain the disks attached to S , since there are at most two such disks. $\square$

Lemma 5 *Let Γ be a stable family of cycles in a surface S . Then no k -gon can exist in Γ for $k \leq 3$.*

Proof Consider a hypothetical k -gon D for Γ , with $k \leq 3$. We prove that D strictly contains another k' -gon with $k' \leq 3$, which is a contradiction, since by induction this would give an infinite family of k -gons in a finite graph.

Assume first that D has exactly k vertices on its boundary; in other words, no cycle crosses the boundary of D . Since D cannot be an empty k -gon, there must be a connected component Γ' of the arrangement of Γ that is entirely inside D . Corollary 2 implies that Γ' contains empty k' -gons with $k' \leq 3$ inside D , and thus D contains a smaller k' -gon, as desired.

Assume now that D has at least $k + 1$ vertices on its boundary, i.e., it intersects at least another cycle in Γ . Consider the restriction of the arrangement of Γ to the closed disk D , and let G be the connected component of that restriction that contains the boundary of D . To prove the lemma, it suffices to prove that G has an interior face of degree at most three. We now assume that every interior face of G has degree at least four, and will reach a contradiction.

G has two types of edges: e_{ext} external edges lying on the boundary of the k -gon, and e_{int} internal edges. Similarly, it has three types of vertices: v_{int} vertices in the interior of the k -gon, v_{flat} degree-three vertices on the boundary of the k -gon, and $v_{\text{extr}} = k$ degree-two vertices on the boundary of the k -gon. Let f denote the number of interior faces of G in the plane. By double-counting arguments, we obtain:

$$v_{\text{extr}} + v_{\text{flat}} = e_{\text{ext}}, \quad (1)$$

$$4f \leq 2e_{\text{int}} + e_{\text{ext}}, \quad (2)$$

$$4v_{\text{int}} + v_{\text{flat}} = 2e_{\text{int}}. \quad (3)$$

Euler's formula implies

$$4 = 4(v_{\text{extr}} + v_{\text{flat}} + v_{\text{int}}) - 4(e_{\text{ext}} + e_{\text{int}}) + 4f$$

$$\begin{aligned}
&\stackrel{(1)}{=} 4v_{\text{int}} - 4e_{\text{int}} + 4f \\
&\stackrel{(3)}{=} -v_{\text{flat}} - 2e_{\text{int}} + 4f \\
&\stackrel{(2)}{\leq} -v_{\text{flat}} - 2e_{\text{int}} + 2e_{\text{int}} + e_{\text{ext}}.
\end{aligned}$$

With (1), this gives $v_{\text{extr}} \geq 4$ which yields a contradiction and concludes the proof. $\square$

3.2 Following the Geodesics

We recall that $\Gamma_1 = (\gamma_{1,1}, \dots, \gamma_{1,n})$ and $\Gamma_2 = (\gamma_{2,1}, \dots, \gamma_{2,n})$ are two stable families of cycles on S in general position satisfying the hypotheses of Theorem 3. Since we assume that the Euler characteristic of S is negative, we can endow S with a hyperbolic metric and identify $\tilde{S}$ with a subset of the open hyperbolic disk. For each j , if $\gamma_{1,j}$ is not null-homotopic, let g_j be the unique geodesic homotopic to it. (Some g_j 's may be identical.)

The following result uses a proposition by de Graaf and Schrijver [11, Proposition 13], the proof of which is based on a slight refinement of Ringel's theorem [32] for arrangements of curves in the disk. It allows us to push each curve in Γ_1 and Γ_2 in the neighborhood of the corresponding geodesic.

Proposition 1 *No cycle in Γ_1 and Γ_2 is null-homotopic. Furthermore, for $i = 1, 2$, for each $\varepsilon > 0$, up to replacing Γ_i with its image by an ambient isotopy of S , we may assume that each cycle $\gamma_{i,j}$ has a lift that belongs to an ε -neighborhood of a lift of g_j .*

Proof It follows from the aforementioned result by de Graaf and Schrijver [11, Proposition 13] that the cycles $\gamma_{i,j}$ in Γ_i can be moved through Reidemeister moves (not increasing the number of crossings) and ambient isotopies into cycles which have lifts that are ε -close to some lift of a geodesic α_j if $\gamma_{i,j}$ is not null-homotopic, and ε -close to a point of $\tilde{S}$ otherwise. As Γ_i is stable, no Reidemeister move *at all* is possible, so the only possible moves are actually isotopies of the surface.

Now, assume that $\gamma_{i,j}$ is null-homotopic. For ε small enough, $\gamma_{i,j}$ is a contractible cycle in a disk; hence if it is simple, it forms a 0-gon; and if it is non-simple, by Corollary 2, it contains a k -gon for some $k \leq 3$. This contradicts the stability of Γ_1 and Γ_2 , hence no cycle $\gamma_{i,j}$ is null-homotopic.

Now, since lifts of $\gamma_{i,j}$ and α_j are ε -close, they share the same endpoints and thus α_j is the unique geodesic homotopic to $\gamma_{i,j}$, which shows that $\alpha_j = g_j$ and concludes the proof. $\square$

So henceforth we assume that Γ_1 and Γ_2 satisfy the conclusion of the above proposition.

The union of the geodesics g_j forms a graph, possibly with simple cycles without vertices, on the surface S ; we denote by E and V its edges and vertices, see Fig. 4. (Simple cycles without vertex are considered to be closed edges.) Each vertex has

Fig. 4 1. Two geodesics on a double torus, and the associated graph. 2. The corridors associated with them. 3. Cycles in the corridors. In vertex polygons, they intersect in a grid-like pattern

even degree, and each geodesic arriving at a vertex from an edge leaves it via the opposite edge.

Now, following ideas by de Graaf and Schrijver [11], we introduce a polygonal decomposition of an ε -neighborhood of the graph (V, E) ; see Fig. 4. With each edge $e \in E$ we associate an *edge polygon* P_e (actually, a quadrilateral), and with each vertex $v \in V$ of degree $2d$ we associate a *vertex polygon* P_v with $2d$ sides, such that each edge $e = uv$ lies in the interior of $P_e \cup P_u \cup P_v$ and v lies in the interior of P_v , and such that the union of all the polygons forms a tubular neighborhood of the cycles g_i .⁶ We can assume that all the polygons are mutually disjoint, except for P_e and P_u for u a vertex incident to e , which share an edge.

Let us call a *corridor* the polygonal neighborhood of a single geodesic g_i , that is $\bigcup P_e \cup \bigcup P_u$ for all e and u in g_i . Every cycle in Γ_i belongs to a single corridor.

The following proposition is proved using Lemmas 4 and 5 and by standard bigon flipping arguments.

Proposition 2 *For $i = 1, 2$, up to replacing Γ_i with its image under an ambient isotopy of S , we may assume that:*

- *each maximal piece of a cycle in Γ_i within a polygon is simple, and has its endpoints on opposite sides of the polygon;*
- *two such pieces cross at most once; moreover, if they cross, then the four endpoints of these two pieces are all in different sides of the polygon (in particular, the polygon is a vertex polygon).*

Proof If a maximal piece of Γ_i within a polygon is non-simple, it forms a 1-gon, which is impossible (Lemma 5). If two pieces cross twice, they form a 2-gon, which is impossible for the same reason.

⁶In the case where a geodesic coincides with a boundary of the surface, one of the edges of each polygon of that geodesic actually lies on this boundary. Also, for simple cycles without vertex, we introduce an edge polygon with two opposite sides glued together.

If such a piece has its endpoints on two different sides of a polygon that are not opposite in that polygon, that polygon is a vertex polygon, and the corresponding cycle does not belong to a single corridor, which is impossible by construction.

If such a piece p has its endpoints on the same side s of a polygon, without loss of generality assume that the disk D bounded by s and p contains no other piece with both endpoints on s . Thus the pieces inside D are simple, pairwise disjoint (otherwise these two pieces would form a 3-gon with p), and connect s to p ; so with an isotopy of the surface, we can push p across s , decreasing the total number of intersections between the cycles and the sides of the polygon. After finitely many such operations, no such piece p exists.

There only remains to prove that there cannot be any (self-)intersection among cycles in the same corridor C . We distinguish two cases:

- If C contains at least one vertex polygon, some cycle crosses every cycle in C ; if there were a (self-)intersection in C , there would be a 3-gon.
- On the other hand, if the corridor C contains no vertex polygon, then C is an annulus. Consider the arrangement of the cycles in Γ , and assume there is a crossing. One connected component of this arrangement is a graph where all vertices have degree four. Thus, by Corollary 2, Γ contains a k -gon for $k \leq 3$, contradicting Lemma 5. $\square$

We can now assume that Γ_1 and Γ_2 satisfy the conclusion of Proposition 2. It follows that all the arcs of Γ_i in a given edge polygon P_e are simple, disjoint, and belong to different cycles. Moreover, each polygon P_v is actually a quadrilateral, because otherwise three arcs in P_v coming from six different sides would cross inside P_v , yielding a 3-gon because of the general position assumption, which is impossible (Lemma 5). Finally, within each polygon P_v , the arcs intersect in a grid-like fashion, as in Fig. 4.

3.3 A Technical Result on Corridors

Henceforth, let us choose an arbitrary orientation on S . Let C be a corridor, oriented in the direction of its geodesic g ; let C_l and C_r be the left and right boundaries of C , respectively. We recall that h is the oriented homeomorphism specified in the hypotheses of Theorem 3.

Let Γ_1^C be the subfamily of cycles in Γ_1 that belong to C , and let Γ_2^C be its image by h . Recall that Γ_i^C has no crossing in an edge polygon P_e . For $i = 1, 2$, the ordering of Γ_i^C along C is defined as follows: Consider an arc in an edge polygon P_e , going from C_l to C_r , crossing each cycle in Γ_i^C exactly once, and record the index of the cycles in Γ_i encountered, in this order along the arc. By construction, this ordering does not depend on the choice of the polygon and arc.

Lemma 6 *The orderings of Γ_1^C and Γ_2^C are the same.*

Proof We first claim that the oriented homeomorphism $h : S \rightarrow S$ lifts to an oriented homeomorphism $\tilde{h} : \tilde{S} \rightarrow \tilde{S}$. Indeed, if we denote by π the projection $\pi : \tilde{S} \rightarrow S$

and apply the lifting theorem [22, Proposition 1.33] to $h \circ \pi$, we get⁷ a continuous map $\tilde{h} : \tilde{S} \rightarrow \tilde{S}$ satisfying $\pi \circ \tilde{h} = h \circ \pi$. Let $x \in \tilde{S}$ and $y = \tilde{h}(x)$. Similarly, we lift h^{-1} to a continuous map $\tilde{h}^{-1}$ such that $\tilde{h}^{-1}(y) = x$. Then $\tilde{h}^{-1}$ and $\tilde{h}$ are inverse continuous maps on $\tilde{S}$, so $\tilde{h}$ is a homeomorphism. Furthermore, $\tilde{h}$ is oriented because h is oriented.

Let $\tilde{g}$ be the lift of a geodesic inside a lift $\tilde{C}$ of C . The homeomorphism $\tilde{h}$ maps $\tilde{g}$ into a possibly different lift of C . However, up to composing $\tilde{h}$ with a deck transformation of $\tilde{S}$, we may assume that $\tilde{h}$ maps $\tilde{g}$ into $\tilde{C}$. Furthermore, $\tilde{g}$ and its image by $\tilde{h}$ have the same orientation in $\tilde{C}$, because otherwise the endpoints of each lift of Γ_1^C in $\tilde{C}$ would be exchanged under $\tilde{h}$, which is not the case since h preserves the homotopy classes of the cycles in Γ_1^C . Therefore, $\tilde{h}$ maps $\tilde{g}$ to a path in $\tilde{C}$ with the same source and target. Since $\tilde{h}$ is oriented, the orderings of the cycles in $\tilde{\Gamma}_1^C$ and $\tilde{\Gamma}_2^C$, from left to right in $\tilde{C}$, are the same. It follows that they are also the same in C . $\square$

3.4 End of the Proof

Proof of Theorem 3 Let P_v and P_e be incident vertex and edge polygons, respectively, and let p be the path that is their common boundary. We first build an isotopy of the surface such that, when restricting to p , the image of each cycle in Γ_1 is the same as the corresponding cycle in Γ_2 . For this purpose, note that the restriction of Γ_1 to p is a finite set of points, and similarly for Γ_2 ; furthermore, the numbers of points are the same (by Lemma 6). We can easily push the intersection points on p so that they coincide, by an isotopy of S that is the identity outside a neighborhood of p . Lemma 6 now implies that, after this isotopy, each arc in P_e corresponds to the same cycle in Γ_1 and Γ_2 .

We can do this operation for every intersection p of a vertex and an edge polygon. Now, within each edge polygon, the arcs of Γ_1 are simple, pairwise disjoint and in the same order as the arcs of Γ_2 ; thus, there exists a homeomorphism from P_e to P_e that is the identity on its boundary and maps the image of Γ_1 inside P_e to the image of Γ_2 inside P_e . By Alexander's Lemma, this homeomorphism is an ambient isotopy. Now, within each edge polygon P_e , the images of Γ_1 and Γ_2 are the same, and each arc corresponds to the same cycle in Γ_1 and Γ_2 .

Now, within each vertex polygon P_v , the endpoints of the arcs of Γ_1 and Γ_2 coincide; moreover, they form combinatorially isomorphic arrangements of arcs (namely, grids) inside P_v . The same argument as above shows that an isotopy of P_v maps the arcs of Γ_1 to the arcs of Γ_2 .

Finally, we have found an ambient isotopy i of S that maps each cycle $\gamma_{1,j}$ in Γ_1 to the corresponding cycle $\gamma_{2,j}$ in Γ_2 , as sets but not necessarily pointwise. Furthermore, since $i(\gamma_{1,j})$ is homotopic to $\gamma_{1,j}$, it is also homotopic to $\gamma_{2,j}$, so the ambient isotopy i preserves the orientations of the cycles. $\square$

4 Isotopies of Graph Embeddings

In this section, we prove the following result.

⁷The hypotheses are trivially fulfilled since the universal cover has trivial fundamental group.

Theorem 4 *Let G_1 and G_2 be two graph embeddings of a graph G on an orientable surface S . Assume that there is an oriented homeomorphism h of S mapping G_1 to G_2 . There exists a family Λ of cycles in G such that the following holds: If, for each cycle γ in Λ , the images of γ in G_1 and G_2 are homotopic, then there exists an ambient isotopy of S taking G_1 to G_2 pointwise.*

Furthermore, the cycles in Λ use each edge of G at most four times in total and, given only the combinatorial map of G_1 on S , one can compute the cycles of Λ in linear time in the complexity of that combinatorial map.

Note that, by the homeomorphism condition, the combinatorial maps of G_1 and G_2 on S have to be the same. We also emphasize that in contrast to Theorem 3, the ambient isotopy we obtain is pointwise.

Conversely, if G_1 and G_2 are isotopic (in particular, if there is an ambient isotopy taking G_1 to G_2), there must exist an oriented homeomorphism mapping one to the other, and the images of any cycle of G in G_1 and G_2 are homotopic. Therefore, Theorem 4 implies Ladegaillerie's result [27] stated in the introduction, and also:

Corollary 3 *Let G_1 and G_2 be two graph embeddings of a graph G in the interior of an orientable surface S . Assume that there exists an isotopy between G_1 and G_2 . Then there exists an ambient isotopy of S between G_1 to G_2 .*

In our proof of Theorem 4, if the input graph embeddings are piecewise-linear with respect to a fixed triangulation of S (which we can assume, after an ambient isotopy, by using techniques as in Epstein [15, Appendix]), our ambient isotopy can be chosen so as to be piecewise-linear. In particular:

Corollary 4 *Let G_1 and G_2 be two piecewise-linear graph embeddings of a graph G in the interior of an orientable surface S . Assume that there exists a (not necessarily piecewise-linear) isotopy between G_1 and G_2 . Then there exists a piecewise-linear ambient isotopy of S between G_1 to G_2 .*

For the proof of Theorem 4, the difficulty of the construction resides in the fact that the families Λ_1 and Λ_2 (the images of Λ in G_1 and G_2) must have small complexity. In a sense, Λ_1 forms a topological decomposition of the tubular neighborhood of G_1 using cycles. However, all known topological decompositions of surfaces made of cycles (like pants decompositions [9], octagonal decompositions [8], or systems of loops [16]) have worst-case complexity $\Omega(gn)$, where n is the complexity of the surface; our construction has linear size in the complexity of the object studied. We suspect that our construction can be useful for other purposes as well.

4.1 Preprocessing Step

For the proof of Theorem 4, we assume for simplicity of exposition that G_1 and G_2 are known. It is immediate to check that, actually, only the combinatorial map of G_1 on S is needed in the constructions.

Fig. 5 Illustration of the proof of Lemma 7

Proposition 3 *Without loss of generality, we may assume that (1) G_1 has a single face, or none of its faces is a disk, and (2) G_1 has no vertex of degree zero or one. Via the oriented homeomorphism h , the same holds for G_2 .*

Intuitively, the proof is simple: Whenever e is an edge of G bounding two different faces, at least one of which is a disk, removing e in G_1 and G_2 does not change whether G_1 and G_2 are isotopic. Similarly, removing vertices of degree zero has no effect on the existence of an isotopy.

In more detail, we will need the following two lemmas.

Lemma 7 *Let e be an edge of G bounding two different faces, at least one of which is a disk, in the embedding G_1 (and thus also in G_2). Let G'_1 and G'_2 be the embedded graphs obtained after the removal of e . Then G_1 and G_2 are ambient isotopic if and only if G'_1 and G'_2 are ambient isotopic.*

Proof The direct implication is obvious. Now, assume we have an isotopy i mapping G'_1 to G'_2 ; we want to deduce that there is an isotopy mapping e_1 to e_2 (the images of e in G_1 and G_2). By composition with i , we may assume that $G'_1 = G'_2$. By the existence of h , we know that e_1 and e_2 are arcs with the same endpoints in the same face of $G'_1 = G'_2$; furthermore, that face is split into two pieces, one of which is a disk, by e_1 (resp., e_2).

The rest of the proof is illustrated in Fig. 5. Since one of the faces bounded by e_1 is a disk, e_1 can be isotoped (with fixed extremities) to a neighborhood of the curve c closing this disk, and the same goes for e_2 . After this isotopy, consider a disk neighborhood of c containing both e_1 and e_2 . Since these two edges have the same endpoints in this disk, the Jordan–Schönflies theorem implies that there is a homeomorphism of the disk, fixed on the boundary, that maps one to another; then, by Alexander’s Lemma, that homeomorphism can be obtained by an isotopy of the disk. $\square$

Lemma 8 *Let v be a vertex of G of degree one, and let e be its incident edge. Let G'_1 and G'_2 be the embedded graphs obtained after the removal of e and v . Then G_1 and G_2 are ambient isotopic if and only if G'_1 and G'_2 are ambient isotopic.*

Proof Again, one direction is trivial; the converse can be proved using similar ideas as the previous lemma. Here, the topological statement that is used is the following: Let p be a point on the boundary of a disk, and let e_1 and e_2 be two simple paths

having p as an endpoint and intersecting the boundary of the disk exactly at p ; then there is an ambient isotopy of the disk, fixed on its boundary, that maps e_1 to e_2 . This again follows by an application of the Jordan–Schönflies theorem and Alexander’s Lemma (by first extending e_1 and e_2 to simple arcs with the same endpoints). $\square$

Proof of Proposition 3 We show below how to build in linear time a subgraph G'' of G satisfying the desired properties and such that, if G'_1 (resp., G'_2) denotes the restriction of G_1 (resp., G_2) to G'' , then G_1 and G_2 are ambient isotopic if and only if G'_1 and G'_2 are ambient isotopic. This is enough to prove the proposition.

Let $G = (V, E)$. We initially set $E' := E$, and, for each edge of E' in turn, we remove it from E' if and only if it is incident to two distinct faces of (V, E') , at least one of which is a disk (in the embedding G_1 or G_2). This is easy to do in linear time, by initially labeling each face of (V, E') with its topology (genus and number of boundary components) and maintaining this labeling during the process.

Let G'_1 and G'_2 be the embeddings of (V, E') induced by G_1 and G_2 , respectively. Lemma 7 implies that G_1 and G_2 are isotopic if and only if G'_1 and G'_2 are isotopic. Furthermore, if G'_1 has at least two faces, one of which is a disk, there exists an edge in G'_1 incident to a disk and to another face; such an edge would have been removed in the process, which is a contradiction. So the first condition is satisfied.

Moreover, we can, in linear time, iteratively remove all degree-one vertices with their incident edges, until no degree-one vertex remains. (Put all degree-one vertices in any list-type data structure; while the structure is non-empty, extract any vertex; if it still has degree one, remove it with its incident edge; if the opposite vertex on that edge has now degree one, add it to the structure; repeat.) Lemma 8 implies that G_1 and G_2 are isotopic if and only if these new graph embeddings, G''_1 and G''_2 , are isotopic.

Finally, if G''_1 (and G''_2) have isolated vertices, we can safely remove them: Since there is a homeomorphism of S taking G_1 to G_2 , the isolated vertices belong to the same faces in both embeddings. $\square$

Proposition 3 leads us to distinguish two cases, leading to slightly different constructions depending on whether after the preprocessing, G_1 has a single face, or none of its faces are disks. We describe the latter first, as the former will build on it; it will be described in Sect. 4.3.

4.2 Proof of Theorem 4 if no Face is a Disk

In this section, we prove Theorem 4 in the special case where no face of G_1 (or, equivalently, G_2) is a disk. We can assume without loss of generality that G_1 satisfies the properties of Proposition 3.

Construction of the Stable Family Γ We first build a family Γ of cycles in G whose images in G_1 or G_2 are slight perturbations of stable families Γ_1 and Γ_2 . If the images of each cycle in Γ in G_1 and G_2 are homotopic, then this almost implies that G_1 and G_2 are isotopic, which suffices to prove Theorem 4. Unfortunately, this is not entirely true, and we need to test a larger family $\Lambda \supset \Gamma$ for homotopy.

Fig. 6 *Top:* A crossover along an edge of $E' \setminus E''$. *Middle:* An embedded graph G_1 on a genus 3 surface. *Bottom:* The corresponding family Γ_1

Proposition 4 *In linear time, we can construct a family of cycles Γ in G such that:*

- *each edge of G is used at most twice by all the cycles in Γ ;*
- *there exists a stable family Γ_1 on S whose cycles are homotopic (by an arbitrarily small perturbation) to the cycles in the images of Γ in G_1 ;*
- *G_1 does not meet the interior of the faces of the arrangement of Γ_1 that are not disks.*

Proof It is actually simpler to explain the construction of Γ_1 first; see Fig. 6 for an example. For simplicity of notation, we let $G_1 := (V, E)$. Recall that the *cyclomatic number* of a connected graph is the minimum number of edges one needs to delete to obtain a tree. Equivalently, it equals its number of edges, minus its number of vertices, plus one.

Let (V', E') be a connected component of (V, E) . By Proposition 3, we can assume that each vertex has degree at least two; in particular, (V', E') has cyclomatic number at least one.

- If (V', E') has cyclomatic number one, then it must be a single cycle. In this case, we add that cycle to Γ_1 .
- Otherwise, (V', E') has cyclomatic number at least two. We add to Γ_1 the cycles that are the boundaries of a tubular neighborhood of (V', E') . Let E'' be the edge set of a spanning tree of (V', E') . For each edge $e \in E' \setminus E''$, we introduce a “crossover” as in Fig. 6 (top) on the two pieces of Γ_1 that run along edge e : Instead of locally having two pieces of cycles that run along edge e without touching it, we now have two pieces of cycles in the neighborhood of edge e that cross at a single interior point of e . Of course, this operation may change the number of cycles of Γ_1 and create self-intersections. See Fig. 6.

We now prove that Γ_1 is a stable family. By construction, for each connected component, the cycles in Γ_1 do not intersect the chosen spanning tree of (V', E') . Let f be a face of the arrangement of the cycles in Γ_1 . We have to prove that f is not a k -gon with $k \leq 3$. Following the definition of Γ_1 , we observe that f is either an *inner disk*, namely, a disk containing entirely a spanning tree (V', E'') of some connected component of (V, E) , or is contained entirely in a single face of (V, E) .

- Assume first that f is an inner disk, containing the spanning tree (V', E'') . By construction, (V', E') has cyclomatic number at least two, so $|E' \setminus E''| \geq 2$. Each edge $e \in E' \setminus E''$ corresponds to a single crossing between cycles of Γ_1 , and this crossing appears twice along the boundary of f ; so f has $2|E' \setminus E''| \geq 4$ crossings of Γ along its boundary.
- Otherwise, f has the same topology as a face of (V, E) , and therefore cannot be a disk.

Hence Γ_1 is a stable family. It follows from the construction that the computation of Γ_1 takes linear time in the complexity of the combinatorial map of G_1 . The cycles in Γ_1 have been constructed in a tubular neighborhood of the graph G_1 ; more precisely, by construction, they run along a side of the edges of G_1 , swapping side whenever they run along an edge not in a spanning tree. Therefore (by retracting the tubular neighborhood) they naturally correspond to a family of cycles Γ in G , and deducing the family Γ from Γ_1 takes linear time. All these cycles use each edge of G at most twice. Furthermore, also by construction, G_1 is included in Γ_1 and in the faces of Γ_1 that are disks. $\square$

The basic idea of the proof is as follows. Assume that each cycle in Γ_1 is homotopic to the corresponding cycle in Γ_2 . Theorem 3 implies that, after an ambient isotopy of S , we can assume that each cycle in Γ_1 coincides with the corresponding cycle in Γ_2 not necessarily pointwise, but with the same orientation. If this was the case pointwise, then, since G_i is “surrounded” by cycles in Γ_i , this would imply that G_1 and G_2 almost coincide and could be moved one into the other by another isotopy. However, the first isotopy does not necessarily map Γ_1 to Γ_2 pointwise, and we need to test that a few more pairs of cycles are homotopic to ensure that it is the case.

Fixing the Map Automorphism We now prove:

Proposition 5 *In linear time, we can construct a family of cycles Λ in G such that:*

- *each edge of G is used at most thrice by all the cycles in Λ .*
- *if we denote by Λ_1 and Λ_2 the images of Λ in G_1 and G_2 , if every cycle in Λ_1 is homotopic to its counterpart in Λ_2 , then an ambient isotopy of S maps Γ_1 to Γ_2 pointwise.*

We will need the following rather independent lemma in the course of the proof.

Lemma 9 *Let C be a family of simple cycles on S , pairwise disjoint except at a single point p , where two cycles may or may not cross. Assume that no component of $S \setminus C$ is a disk bounded by one or two cycles. Then the cycles in C are pairwise (freely) non-homotopic.*

Proof We will use the fact that two simple homotopic cycles cross transversely an even number of times (because they form bigons [18, Proposition 1.7]).

First, no cycle in C is contractible; otherwise, it would bound a disk on the surface. The cycles inside that disk are all contractible, and therefore do not cross at p because of the aforementioned fact. Taking an innermost such cycle, we obtain a component of $S \setminus C$ bounded by one cycle, contradicting the assumption.

Assume now for the sake of a contradiction that two cycles c_1 and c_2 are homotopic. They meet at point p without crossing transversely. After a local perturbation, these cycles become disjoint, and therefore bound an annulus [15, Lemma 2.4]. Thus, the unperturbed cycles can be viewed as two loops ℓ_1 and ℓ_2 based at p that bound a disk. There may be other cycles inside that disk, but in all cases a face inside it is a disk bounded by one or two cycles, which is impossible. $\square$

Proof of Proposition 5 The family Λ is the union of the stable family Γ defined in Proposition 4 and of the family Φ defined as follows. Recall that in the proof of Proposition 4, we considered each connected component (V', E') of the graph $G = (V, E)$ in turn. If (V', E') had cyclomatic number at least two, we considered the edge set E'' of a spanning tree of (V', E') . A *fundamental cycle* of (V', E') is a simple cycle in (V', E') containing exactly one edge in $E' \setminus E''$. We put in Φ an arbitrary fundamental cycle for each connected component (V', E') of cyclomatic number at least two. The fundamental cycles of a connected component (V', E') can be extended towards an arbitrary root p of the spanning tree (V', E'') and then slightly perturbed on S so that they become simple and pairwise disjoint except at p , where they may or may not cross. The faces of this new family C of perturbed cycles correspond to the faces of (V', E') . Moreover, C satisfies the hypotheses of Lemma 9: Indeed, if there is a disk in $S \setminus C$ bounded by one or two cycles, there must be at least one connected component of G_1 inside it because no face of G_1 is a disk; but then this connected component is contractible, which is absurd since the preprocessing removed all the contractible components of G_1 . Hence, the fundamental cycles of any given connected component (V', E') are pairwise non-homotopic.

Clearly the family $\Lambda = \Gamma \cup \Phi$ can be computed in linear time and uses each edge of G at most thrice. Assume that, for each cycle λ in Λ , the images of λ in G_1 and G_2 are homotopic. There remains to prove that some isotopy of S maps Γ_1 to Γ_2 pointwise.

By Proposition 4, Γ_1 is a stable family, and of course $\Gamma_2 := h(\Gamma_1)$ as well. Since each cycle in Γ_1 is homotopic to the corresponding cycle in Γ_2 , Theorem 3 implies that some isotopy of S takes Γ_1 to Γ_2 , not necessarily pointwise, but preserving the orientations of the cycles.⁸ Therefore, up to composing with this isotopy, we can assume that each cycle in Γ_1 coincides, as a set, with the corresponding cycle in Γ_2 , and with the same orientation. Hence, this isotopy induces an orientation-preserving map isomorphism i between Γ_1 and Γ_2 , and since Γ_1 and Γ_2 have the same extended combinatorial maps, i can be viewed as an orientation-preserving map automorphism of Γ_1 . As each cycle in Γ_1 is isotoped to the corresponding cycle in Γ_2 , i maps each

⁸Actually, if S has nonnegative Euler characteristic, the results in Sect. 7 show that the isotopy can be chosen so as to be pointwise, which concludes the proof of this proposition.

Fig. 7 A graph Γ_1 drawn on a sphere with four holes. The inner disk is the darker part, and the orientations at each vertex are pictured according to the orientation of the single cycle. If $i(v) = v'$, φ is sent to φ' which is not homotopic to it

connected component of Γ_1 to itself, and it maps each crossing of Γ_1 to a crossing of Γ_1 .

We now want to ensure that i is the identity map automorphism, which would imply the existence of a pointwise ambient isotopy between Γ_1 and Γ_2 (as in the proof of Corollary 1). However, this is not necessarily the case, as was pictured in Fig. 3.

Let (V', E') be a connected component of G ; let Γ'_1 be the arrangement of the cycles of Γ_1 corresponding to that connected component. If (V', E') has cyclomatic number one, by construction, Γ'_1 is just a cycle, which i maps to itself, preserving its orientation; so i is the identity map automorphism on Γ'_1 .

Otherwise, (V', E') has cyclomatic number at least two. We first note that Γ'_1 is connected; indeed, the inner disk D of (V', E') is bounded by all cycles in Γ'_1 . Moreover, the faces of Γ_1 that are disks are exactly the inner disks; so i maps inner disks to inner disks, and therefore maps D to itself.

Let v be the vertex of Γ'_1 corresponding to the cycle φ in $\Phi = \Lambda \setminus \Gamma$ (see Fig. 7). If i maps v to another vertex v' of Γ'_1 , as the inner disk is mapped to itself, i necessarily maps φ to a cycle in the inner disk crossing v' once, i.e. another fundamental cycle in (V', E') , which is, as shown above, not homotopic to φ ; this is a contradiction. So i maps v to itself. Furthermore, if we orient the four edges incident to v with the orientation of the corresponding cycles, v has two outgoing edges, consecutive in the cyclic order around v , and two incoming edges, also consecutive. Since i maps each edge to another edge with the same orientation, it maps v to v , and it is an orientation-preserving map automorphism, it must thus map each edge incident to v to itself, with the same orientation. Since Γ'_1 is connected, by propagation we deduce that i is the identity map automorphism on Γ'_1 , which concludes the proof. $\square$

End of Proof of Theorem 4 We now conclude the proof of Theorem 4 if none of the faces of G_1 are disks.

According to the hypotheses, for all the cycles $\gamma \in \Lambda$, the images of γ in G_1 and G_2 are homotopic, which implies by Proposition 5 that we can assume that $\Gamma_1 = \Gamma_2$ pointwise. Then, each face of Γ_1 is mapped by h to itself, because h is an oriented homeomorphism.

In particular, $G_1 \cap \Gamma_1 = G_2 \cap \Gamma_2$. In every disk of $S \setminus \Gamma_1 = S \setminus \Gamma_2$, the oriented homeomorphism h is the identity on the boundary; therefore, by Alexander's Lemma, it is an ambient isotopy relatively to the boundary. This gives us an isotopy between G_1 and G_2 on every such disk, relatively to $\Gamma_1 = \Gamma_2$. By gluing these isotopies together along their boundaries, we get an isotopy of S mapping G_1 to G_2 , because G_1 and G_2 are included in the closures of the faces of $\Gamma_1 = \Gamma_2$ that are disks (Proposition 4).

Since the family Λ covers each edge of G at most thrice, it has linear complexity. As it can be computed in linear time, this concludes the proof of Theorem 4 if none of the faces of G_1 are disks.

4.3 Proof of Theorem 4 if the Only Face of G_1 is a Disk

By Proposition 3, either (1) G_1 has no face that is a disk, or (2) G_1 has a single face, and that face is a disk. We proved Theorem 4 in case (1) in the previous section, and shall now deal with case (2). In other words, we assume that G_1 is a *cut graph*.

In that case, the above construction does not seem to work: Since inner disks are not the only faces of Γ_1 that are disks, there is no guarantee that an inner disk is mapped to itself in the proof of Proposition 5. To circumvent this issue, the high-level idea is the following: We remove one cycle from G_1 so that the only face of G_1 is not a disk anymore but a cylinder, in which case the results from the previous section apply. We then check that the remaining cycle and its counterpart in G_2 are homotopic, and prove that this guarantees the existence of an isotopy between G_1 and G_2 .

Since G_1 is a cut graph, S has no boundary. Moreover, G is connected, and G_1 is made of a spanning tree $T = (V', E')$ (as in Sect. 4.2) and $2g$ additional edges. Let e be one of these edges (chosen arbitrarily); let γ_e be the fundamental cycle with respect to T corresponding to edge e . Let G' be the graph G with edge e removed, and let G'_1 and G'_2 be the restrictions of G_1 and G_2 to G' . Note that G'_1 has a single face, which is a cylinder.

We can now apply the result of the previous section to G'_1 and G'_2 : We obtain a family Λ' of cycles in G' with the property that, if their images in G'_1 and G'_2 are homotopic, then G'_1 and G'_2 are isotopic. Furthermore, Λ' can be computed in linear time, and uses each edge of G' at most thrice.

To prove Theorem 4 for our graph G , we take $\Lambda := \Lambda' \cup \{\gamma_e\}$. Obviously, Λ can be computed in linear time and uses each edge of G at most four times in total. Assume now that the images of each cycle of Λ in G_1 and G_2 are homotopic. It suffices to prove that, under this condition, some ambient isotopy of S takes G_1 to G_2 . Since the images of each cycle in Λ' in G'_1 and G'_2 are homotopic, we may assume that $G'_1 = G'_2$. There remains to prove that an isotopy of the surface allows to push the image e_1 of e in G_1 to the image e_2 of e in G_2 .

Fig. 8 1. The graph $G'_1 = G'_2$ and the edges e_1 and e_2 . 2. The cylinder C obtained after cutting along G'_1 , with the arcs a_1 and a_2 corresponding to e_1 and e_2 , and the cycles b and b' . If a_1 and a_2 are not isotopic, the images of γ_e in G_1 and G_2 cannot be (freely) homotopic

The surface obtained after cutting S along $G'_1 = G'_2$ is a cylinder C , and the images of e on C become arcs a_1 and a_2 with the same endpoints, one on each boundary, as shown on Fig. 8. (Indeed, if both endpoints were on the same boundary, a_1 would bound two faces on C which would correspond to two faces on S , reaching a contradiction.) Now, to conclude, we only need to show that a_1 and a_2 are isotopic relatively to the boundary of this cylinder.

The end of the proof uses some elementary notions of homology, we refer to [22] for the background. Assume, for the sake of a contradiction, that a_1 and a_2 are non-isotopic arcs relatively to the boundary of C . This implies that they are non-homotopic on C [15, Theorem 3.1]. Hence there exists an integer $n \neq 0$ such that a_2 is homotopic to $a_1 \cdot b^n$, where b is a loop that is a boundary of the cylinder C . Since the images of γ_e in G_1 and G_2 are (freely) homotopic on S , they are $\mathbb{Z}$ -homologous. This implies that b^n , and thus b , has zero $\mathbb{Z}$ -homology. By translating along the cylinder, b is homotopic on S to a simple cycle b' that crosses e_1 exactly once and crosses G_1 nowhere else, as pictured on Fig. 8. Hence b' is a simple cycle on S that crosses the image of γ_e in G_1 exactly once; thus b' is non-separating, and therefore cannot have zero $\mathbb{Z}$ -homology. This contradiction completes the proof of Theorem 4.

5 Algorithms

For Theorems 1 and 2, it suffices to be able to test the existence of an oriented homeomorphism, and of homotopies between the cycles in Λ , as computed by Theorem 4, in the indicated amount of time. We prove Theorems 1 and 2 in Sects. 5.1 and 5.2, respectively.

5.1 Surfaces: Proof of Theorem 1

Recall that, in Theorem 1, the input of the algorithm consists of a fixed graph H cellularly embedded on a fixed surface S , and of embeddings G_1 and G_2 of a graph G . Furthermore, k_1 (resp., k_2) denotes the complexity of the combinatorial map of the arrangement of G_1 (resp., G_2) with H .

Homeomorphism Test For the case of graphs on surfaces, the existence of an oriented homeomorphism that maps G_1 to G_2 can be checked in $O(k_1 + k_2)$ time.

Indeed, let us choose an arbitrary orientation on S ; this induces an orientation of the combinatorial map of H , and hence an orientation of the combinatorial map of the arrangement of G_i and H , for $i = 1, 2$. Computing the number of boundary components of S in each face, as well as the genus of each face (using the Euler characteristic), and “erasing” the graph H in both arrangements gives us oriented combinatorial maps for each G_i . They are isomorphic if and only if there exists an oriented isomorphism of S between G_1 and G_2 (Lemma 3), and this can be checked in linear time.

Homotopy Tests The homotopy tests can also be performed in $O(k_1 + k_2)$ time. Indeed, recall that the input to the algorithm consists of the combinatorial maps of the arrangement of G_1 and H on one hand, and of G_2 and H on the other hand, where H is a fixed cellular graph embedding; k_1 and k_2 denote the complexities of these two maps. In $O(k_1 + k_2)$ total time, we can compute the cyclically ordered list of edges of H crossed by each cycle of Λ in the embeddings G_1 and G_2 . This gives us a set of pairs of cycles in the dual graph H^* of H that have to be tested for homotopy. The total complexity of these cycles is $O(k_1 + k_2)$, and H^* has complexity $O(k_1 + k_2)$ as well. Lazarus and Rivaud, and later Erickson and Whittlesey [17, 28] prove that, after a preprocessing linear in the complexity of the cellular graph H^* , one can test homotopy of cycles in H^* in time linear in the complexities of these cycles. (An earlier paper by Dey and Guha [12] claims a similar result, except for some low-genus surfaces, but Lazarus and Rivaud point out some problems in their proof.) These papers address only the case of surfaces without boundary, but the case of surfaces with boundary is easier, as the fundamental group is free; alternatively, homotopy tests for cycles on surfaces with boundary can be performed using an algorithm for surfaces without boundary by first attaching a handle to each boundary component, which does not change the outcomes of the homotopy tests.

This concludes the proof of Theorem 1.

5.2 Punctured Plane: Proof of Theorem 2

We now give our algorithm for the punctured plane model; so let G_1 and G_2 be two embeddings of a graph G in the punctured plane $\mathbb{R}^2 \setminus P$. Let $\mathcal{P}$ be a set of disjoint open polygons (for example squares), one around each point of P , that avoid G_1 and G_2 ; also, let B be a large closed square such that G_1 , G_2 , and the closure of $\mathcal{P}$, are in the interior of B . By compactness, any isotopy between G_1 and G_2 , if it exists, must avoid neighborhoods of P and stay in a bounded area of the plane; therefore, such an isotopy exists if and only if such an isotopy exists in $B \setminus \mathcal{P}$. In other words, since $B \setminus \mathcal{P}$ is a surface with boundary, we are exactly in the topological setting of the previous sections, except that the input to the algorithm is given in a different form.

As above, our algorithm relies on two subroutines: a test for the existence of an oriented homeomorphism, and a test for homotopy between cycles. We actually give two algorithms for the latter problem, because, depending on the ratio between $k_1 + k_2$ and p , one is faster than the other.

Homeomorphism Test To test whether there exists an oriented homeomorphism of the plane that maps G_1 to G_2 in $\mathbb{R}^2 \setminus P$, we compute the oriented combinatorial map of G_1 in the punctured plane using a sweep-line algorithm for $G_1 \cup P$ in $O((k_1 + p) \times \log(k_1 + p))$ time [1]. Then we apply the same procedure with G_2 instead of G_1 , and check that the two resulting oriented combinatorial maps are isomorphic.

First Algorithm for Homotopy Tests We transform the input into the surface model. For this purpose, we compute a triangulation T of B in $O(p \log p)$ time (for example, a Delaunay triangulation). We can then easily determine the arrangement of G_1 with T in $O(k_1 p)$ time, because each segment in G_1 has $O(p)$ crossings with T ; that arrangement has complexity $O(k_1 p)$. We can apply the same procedure to G_2 . After a slight modification of T that does not affect its complexity, we may assume that T is a triangulation of the bounding box, minus a set of small square obstacles. We can then test homotopy of cycles in time linear in the number of their crossings with T , either by computing and comparing their cyclically reduced crossing words with T (since the fundamental group is a free group) or by applying the algorithm by Lazarus and Rivaud or the one by Erickson and Whittlesey. This takes $O((k_1 + k_2)p)$ time.

Second Algorithm for Homotopy Tests To get a subquadratic running time in the input size, we improve the homotopy test by adapting an algorithm by Cabello et al. [6, Sect. 4]. Their algorithm tests homotopy for paths in the punctured plane, not homotopy for cycles; however, it can be modified to handle this case also. More precisely, we show below that, after $O(p^{1+\varepsilon})$ preprocessing time (for any $\varepsilon > 0$), one can test homotopy of two (possibly non-simple) cycles γ_1 and γ_2 of complexities m_1 and m_2 , respectively, in $O((m_1 + m_2)\sqrt{p} \log p)$ time.

The main idea is to replace the triangulation T in the first algorithm above with a cellular decomposition of B that has a nicer property: Each line in the plane crosses at most $O(\sqrt{p})$ segments of that decomposition. Cabello et al. [6] show how to compute such a cellular decomposition in $O(p^{1+\varepsilon})$ time. The two input cycles γ_1 and γ_2 cross this decomposition $O((m_1 + m_2)\sqrt{p})$ times. Then, computing the cyclically ordered lists of edges of the decomposition crossed by these two cycles takes $O((m_1 + m_2)\sqrt{p} \log p)$ time using ray shooting, as done also in the paper by Cabello et al. We conclude using the same method as in the first algorithm, with the cellular decomposition in place of the triangulation. This proves Theorem 2.

6 Graph Isotopies with Fixed Vertices

In this section, we briefly indicate how the previous techniques extend to the graph isotopy problem where, in addition, we require the isotopy to fix some vertices. Formally, for G_1 and G_2 two embeddings of a graph $G = (V, E)$ on the interior of a surface S and a set $V_f \subseteq V$ such that the embeddings of V_f are the same in both graphs, we want to test whether there exists an isotopy h_t between G_1 and G_2 such that $h_t|_{V_f}$ is the identity for all $t \in [0, 1]$. We call this the *fixed vertices graph isotopy problem*.

We will deal with this variant by applying our algorithm in the setting of surfaces with punctures. Surfaces with punctures are surfaces where we removed a finite number of points. We will denote such a surface by (S, P) , where P is the set of punctures on the surface without punctures S . Although such punctured surfaces are not compact, they share many properties with usual surfaces. In particular, they can be endowed with a hyperbolic metric if their Euler characteristic is negative, where we define it by $\chi(S, P) = 2 - 2g - b - |P|$.

Theorem 5 ([18, Theorem 1.2]) *Consider any surface, perhaps with punctures or boundary. If its Euler characteristic is negative, then there exists a complete, finite-area Riemannian metric of constant curvature -1 on the surface, with the property that the boundary of the surface is totally geodesic.*

Hyperbolic surfaces with punctures share many properties with usual hyperbolic surfaces; we refer the reader to [18] for details. We will mainly use the following proposition, which we give without proof. A cycle γ is *homotopic into a neighborhood of a puncture* if, for every neighborhood N of that puncture, there is a cycle homotopic to γ that lies entirely within N .

Proposition 6 ([18, Proposition 1.3]) *Let S be a hyperbolic surface. If α is a cycle in S that is not homotopic into a neighborhood of a puncture, then α is homotopic to a unique geodesic cycle γ .*

In a nutshell, the strategy for solving the fixed vertices graph isotopy problem is to put punctures on the fixed vertices of G_1 (or, equivalently, G_2). We essentially use the same algorithm as in the compact case, with some minor modifications listed below. To keep things simple, we only deal with the case where the resulting surface with punctures is hyperbolic; the remaining cases (the once or twice punctured sphere and the once punctured disk) are simpler and can be dealt with the same idea coupled with the techniques of Sect. 7.

Isotopies of Stables Families of Cycles The analog of Theorem 3 still holds in the case of surfaces with punctures:

Theorem 6 *Let (S, P) be an orientable surface with punctures and let $\Gamma_1 = (\gamma_{1,1}, \dots, \gamma_{1,n})$ and $\Gamma_2 = (\gamma_{2,1}, \dots, \gamma_{2,n})$ be two stable families of cycles on (S, P) in general position such that:*

1. *There exists an oriented homeomorphism h of (S, P) mapping each cycle $\gamma_{1,j}$ of Γ_1 to the corresponding cycle $\gamma_{2,j}$ of Γ_2 not necessarily pointwise, but preserving the orientations of the cycles, and*
2. *Each cycle of Γ_1 is homotopic to the corresponding cycle of Γ_2 .*

Then there is an isotopy of (S, P) that maps each cycle of Γ_1 to the corresponding cycle of Γ_2 , not necessarily pointwise, but preserving the orientations of the cycles.

Proof The proof is essentially the same as for Theorem 3 (in the hyperbolic case). Indeed, all the properties related to hyperbolicity are maintained, with a small caveat

stated in Proposition 6: For the cycles that are homotopic into the neighborhood of a puncture, there is no corresponding geodesic, as it would be infinitely small. But when we do local shortenings as in Proposition 1, the cycles homotopic to the neighborhood of a puncture are pushed into an ε -neighborhood of this puncture, and this allows us to define an *annular corridor* around it, where, using the same proof as Proposition 2, the cycles do not cross. Since each cycle γ of Γ_1 is homotopic to $h(\gamma)$, and h preserves the orientations of the cycles, each puncture with an annular corridor is mapped to itself by h . In this annular neighborhood, the orderings of Sect. 3.3 are the same because the oriented homeomorphism h maps the puncture to itself and can thus be extended to the annular corridor. The rest of the proof follows identically. $\square$

Preprocessing On a surface with punctures, we extend slightly our definition of graph embeddings to allow the graphs to hit punctures. Formally, a graph $G = (V, E)$ is embedded on (S, P) in general position if it is embedded on S and $G \cap P \subseteq V$. By a slight abuse of language, the *faces* of such an embedding will denote the faces of the usual embedding we obtain by removing the punctures under the vertices. The following lemma is straightforward.

Lemma 10 *Let $G = (V, E)$ be embedded as G_1 and G_2 on a surface S , such that both embeddings are identical for $V_f \subseteq V$. Then G_1 and G_2 are isotopic with V_f fixed if and only if G_1 and G_2 (seen as embeddings on (S, V_f)) are isotopic in (S, V_f) .*

Proof The same isotopy can be used in both settings. $\square$

Thus, solving the fixed vertices graph isotopy problem amounts to solving the graph isotopy problem on surfaces with punctures, where the punctures are placed on the fixed vertices.

The preprocessing step is slightly altered. We obviously do not want to remove from the graphs degree-one vertices that are supposed to stay fixed during the isotopy, so we use the following proposition instead.

Proposition 7 *Without loss of generality, we may assume that (1) G_1 has a single face, or none of its faces is a disk, (2) G_1 has no vertices of degree zero, and (3) The only vertices of G_1 of degree one are on punctures.*

The proof is the same, except that we do not touch the fixed vertices of degree one, i.e., vertices of degree one on punctures.

Construction of the Stable Family Unlike the compact case, we do not distinguish two cases in the construction of the stable families Γ_1 and Γ_2 . It follows closely the one described in Proposition 4, but we need to ensure that there is at most one puncture in the disks we obtain.

As in the proof of Proposition 4, we build the stable family Γ_1 corresponding to G_1 and then map it to Γ_2 by the homeomorphism that maps G_1 to G_2 . To build Γ_1 , we consider each connected component of G_1 in turn. Let (V', E') be such a connected

component and let $V'_f = V' \cap V_f$ be the set of vertices on the punctures. If $V'_f = \emptyset$, we apply the same procedure as in the proof of Proposition 4. Otherwise, we compute a spanning forest (V', E'') of (V', E') , where each tree of the forest contains exactly one vertex of V'_f . As in the proof of Proposition 4, we add to the stable family the cycles that are the boundaries of a tubular neighborhood of (V', E') , but adding a “crossover” at each edge in $E' \setminus E''$.

To see that this results in a stable family Γ_1 , note that (as before) the faces of the family are of two types:

- *inner disks*, which contain some pieces of the graph G ; such faces either contain exactly one puncture in their interior (in case they are built as indicated in the previous paragraph), or contain no puncture but have degree at least four (as in the proof of Proposition 4);
- *outer disks*, which have the topology of the faces of the graph G_1 , and thus, by the preprocessing step, cannot be disks, unless there is exactly one such disk, which must therefore be a cut graph of the surface S without boundary (and after disregarding the punctures). Such a disk has degree at least four if the surface has genus at least one; if the surface S is a sphere, the hyperbolicity hypothesis implies that there are at least three punctures, making at least two crossovers and thus an outer face of degree at least four.

Fixing the Map Automorphism We now need to prove the counterpart of Proposition 5. Let (V', E') be a connected component of G_1 . If V' contains no puncture, then we proceed as in the proof of Proposition 5: After testing the homotopy class of a given fundamental cycle, we can certify that the map automorphism is fixed (or, in other words, that some ambient isotopy maps the stable cycles corresponding to (V', E') *pointwise*). We now consider the case where V' contains at least one puncture.

Since each puncture is mapped to itself using the map automorphism, we know that each inner disk of (V', E') is mapped to itself. We use an argument similar to the proof of Proposition 5, but the counterpart of Lemma 9 is simpler in our case.

We choose an arbitrary *fundamental path* in (V', E') , which connects two punctures (possibly identical) using exactly one edge *not* in the spanning forest (V', E'') . (Such an edge exists, since, after the preprocessing step, it cannot be that (V', E') is a tree with a single puncture.) Let v be the vertex of the arrangement of Γ_1 corresponding to this fundamental path. We test whether this fundamental path is homotopic *with fixed endpoints* to its counterpart in G_2 *on the punctured surface, except that the punctures on the endpoints of the path are removed*.

If these paths are not homotopic, clearly G_1 and G_2 are not isotopic, since any ambient isotopy between G_1 and G_2 (fixing a subset of the vertices) yields a homotopy of the fundamental path with fixed vertices.

On the other hand, assume that the images of the fundamental path in G_1 and G_2 are homotopic with fixed endpoints. Assuming that the cycles are homotopic in Γ_1 and Γ_2 , we obtain by Theorem 6 that some isotopy i maps each cycle in Γ_1 to the corresponding cycle in Γ_2 , not necessarily pointwise, but preserving the orientations of the cycles. As in the proof of Proposition 5, i can be viewed as an orientation-preserving map automorphism of Γ_1 . We recall that Γ'_1 denotes the arrangement of

cycles of Γ_1 corresponding to a connected component of G . To conclude, it suffices to prove that, for every connected component of G , this map automorphism restricts to the identity on Γ'_1 , and by connectedness of Γ'_1 it actually suffices to prove that it maps v to itself. But if i were to map v into another vertex, it would map a fundamental path into a different one.

There just remains to prove that two different fundamental paths are *not* homotopic with fixed endpoints on this surface. Indeed, we can perturb them slightly so as to make them disjoint (except at their endpoints), so if they are homotopic, they bound a degree-two disk. If this disk contains no other component of G_1 , then G_1 has a face that is a degree-two disk, which is excluded by the preprocessing step and the hyperbolicity hypothesis. If it contains another component of G_1 , this component is contractible and without fixed vertices, which is again excluded by the preprocessing.

Therefore i is the identity, which concludes the proof of the counterpart of Proposition 5.

End of Proof The components of $S \setminus \Gamma_1 = S \setminus \Gamma_2$ are now punctured disks instead of disks, but we made sure that they only contain a *single* puncture. Hence, since Alexander's Lemma also applies to once-punctured disks [18, Sect. 2.2.1], the end of the proof remains the same.

Conclusion In summary, some minor modifications of our algorithms can also handle the case where some vertices of the graph have to be fixed during the entire isotopy. In some cases, we need to use homotopy tests for paths instead of cycles (which is a computationally simpler problem, since it reduces to testing contractibility of cycles). In particular, here are the counterparts of Theorems 1 and 2:

Theorem 7 *Let S be an orientable surface, possibly with boundary. Let H be a fixed graph cellularly embedded on S . Let G_1 and G_2 be two graph embeddings of the same graph G on S , each in general position with respect to H . Furthermore, let V_f be a set of vertices of G . Given the combinatorial map of the arrangement of G_1 with H (resp., G_2 with H), of complexity k_1 (resp., k_2), we can determine whether G_1 and G_2 are isotopic, fixing the vertices in V_f , in $O(k_1 + k_2)$ time.*

Theorem 8 *Let P be a set of p points in the plane, and let G_1 and G_2 be two piecewise-linear graph embeddings of the same graph G in $\mathbb{R}^2 \setminus P$, of complexities (number of segments) k_1 and k_2 respectively. Furthermore, let V_f be a set of q vertices of G . We can determine whether G_1 and G_2 are isotopic in $\mathbb{R}^2 \setminus P$, fixing the vertices in V_f , in time $O(n^{3/2} \log n)$ time, where n is the total size of the input. In more detail, the running time is, for any $\varepsilon > 0$,*

$$O((k_1 + p + q) \log(k_1 + p + q) + (k_2 + p + q) \log(k_2 + p + q) + \min\{(k_1 + k_2)(p + q), (p + q)^{1+\varepsilon} + (k_1 + k_2)\sqrt{p + q} \log(p + q)\}).$$

Moreover, the above proof implies that Corollaries 3 and 4 also extend to the case where some vertices of the graph embeddings have to be fixed: any isotopy between

Fig. 9 Gadget reducing the case of fixed vertices to the general case: A fixed vertex v of degree d is replaced with d vertices and edges forming a cycle γ_v around a new boundary component of the surface

G_1 and G_2 extends to an ambient isotopy, and that ambient isotopy can be chosen piecewise-linear if G_1 and G_2 are piecewise-linear.

Actually, if it were known that two piecewise-linear graph embeddings are isotopic if and only if they are piecewise-linearly isotopic (with fixed vertices), our result of Sect. 3 would extend rather directly to the case of fixed surfaces, by the construction shown in Fig. 9, where each vertex of G that has to be fixed is replaced with a cycle of vertices and edges enclosing a boundary component of the surface: Any piecewise-linear isotopy with fixed vertices in the original graph corresponds to a piecewise-linear isotopy (without fixed vertices) in the new graph and conversely, so we can apply our results on graph isotopies without fixed vertices. However, this fact seems harder to prove for arbitrary continuous isotopies.

7 Exceptional Surfaces

In this section, we prove that Theorem 3 also holds for surface of nonnegative Euler characteristic. We note that in these cases we obtain a stronger theorem than in the general case, as the ambient isotopy we obtain maps the families of cycles pointwise. We split the proof in two parts depending on whether the surface is a plane with boundaries in Sect. 7.1 or a torus in Sect. 7.2. As a foreword, we note that the results of Lemmas 4, Corollary 2, and Lemma 5 and the first result of Proposition 1 (no cycle in a stable family is null-homotopic) also hold in the nonnegative Euler characteristic case.

7.1 Sphere, Disk, and Annulus

Since the Euler characteristic of a surface is $\chi(S) = 2 - 2g - b$, the only cases where the surface is a plane with boundaries and has nonnegative Euler characteristic are the sphere, the disk, and the annulus. In these cases, the proof of Theorem 3 is a simple corollary of Lemmas 4 and 5. We even obtain a slightly stronger statement, because we can take the isotopy to map *pointwise* Γ_1 to Γ_2 :

Theorem 9 *Let S be a sphere, a disk or an annulus and let $\Gamma_1 = (\gamma_{1,1}, \dots, \gamma_{1,n})$ and $\Gamma_2 = (\gamma_{2,1}, \dots, \gamma_{2,n})$ be two stable families of cycles on S in general position such that:*

1. *There exists an oriented homeomorphism h of S mapping each cycle $\gamma_{1,j}$ of Γ_1 to the corresponding cycle $\gamma_{2,j}$ of Γ_2 not necessarily pointwise, but preserving the orientations of the cycles, and*
2. *Each cycle of Γ_1 is homotopic to the corresponding cycle of Γ_2 .*

Then there is an ambient isotopy of S mapping each cycle of Γ_1 to the corresponding cycle of Γ_2 pointwise.

Proof If the surface S is a sphere or a disk, all the cycles in Γ_1 and Γ_2 are null-homotopic, which is impossible as noted above. Hence these families are empty and the theorem is trivial.

If the surface is an annulus, for $i = 1, 2$, we claim that there are no crossing points in Γ_i , i.e., all the cycles are simple and two distinct cycles do not intersect each other. Indeed, if there were a crossing point, the connected component of Γ_i containing it would form a planar graph such that every vertex has degree four. Hence, by Corollary 2, there would be at least a k -gon with $k \leq 3$, contradicting, with Lemma 5, the stability of the family Γ_i .

Thus, Γ_1 is a family of disjoint simple cycles homotopic to the boundaries of the annulus, and the same holds for Γ_2 . There is an isotopy of S mapping one family into the other if and only if they have the same ordering, as defined in the proof for the hyperbolic case. But this is exactly what the oriented homeomorphism between them ensures. This concludes the proof. $\square$

7.2 Torus

The proof in the case of the torus is slightly more involved. Let us introduce a few definitions before delving into it. We choose a Euclidean metric on the torus, which induces one on its universal cover $\mathbb{R}^2$. This allows to define *translations* on the torus, which are projections of the usual translations of $\mathbb{R}^2$. Geodesics on the torus lift to straight lines in the plane, and two geodesics are homotopic if and only if these lines have the same slope, as a slope $s = \frac{m}{n}$ determines a unique element (m, n) with $m \wedge n = 1$ of the fundamental group of the torus. When we mention the *slope* of a geodesic on the torus, we refer to the slope of one of its lifts in the universal cover. Note that as a translation is an isometry, it maps a geodesic to another geodesic.

Theorem 10 *Let S be a torus and let $\Gamma_1 = (\gamma_{1,1}, \dots, \gamma_{1,n})$ and $\Gamma_2 = (\gamma_{2,1}, \dots, \gamma_{2,n})$ be two stable families of cycles on S in general position such that:*

1. *There exists an oriented homeomorphism h of S mapping each cycle $\gamma_{1,j}$ of Γ_1 to the corresponding cycle $\gamma_{2,j}$ of Γ_2 not necessarily pointwise, but preserving the orientations of the cycles, and*
2. *Each cycle of Γ_1 is homotopic to the corresponding cycle of Γ_2 .*

Then there is an ambient isotopy of S mapping each cycle of Γ_1 to the corresponding cycle of Γ_2 pointwise.

Proof For all the cycles γ in Γ_1 or Γ_2 , we start by applying de Graaf and Schrijver [11, Proposition 13] as in the proof of Proposition 1: Up to applying an isotopy of S ,

we can assume that γ is contained in the ε -neighborhood of one of its corresponding geodesics (or of a point, if γ is contractible). As in the proof of Proposition 1, we infer that γ is not contractible. Since in a torus, geodesic cycles are either simple or multiple concatenations of the same simple cycle, their ε -neighborhoods are annuli. Hence, every cycle in Γ_1 and Γ_2 can be assumed to lie in an annulus.

If one of these cycles $\gamma \in \Gamma_i$ is non-simple, it forms a graph embedded on an annulus such that every vertex has degree four. By Corollary 2, one of the faces of this graph is a disk with degree lower than four, which with Lemma 5 contradicts the stability of Γ_i . Thus all the cycles in Γ_1 and Γ_2 are simple. By the same argument, for $i = 1, 2$, two homotopic cycles in Γ_i do not cross each other.

The isotopy can then be found as follows. We split the proof in two cases.

Case 1: If all the cycles in Γ_1 are homotopic or inverse homotopic,⁹ we just pick an arbitrary one, say $\gamma_{1,1}$, and apply the pointwise isotopy mapping it to $\gamma_{2,1}$, which exists because they are homotopic.¹⁰ Cutting the surface along $\gamma_{1,1} = \gamma_{2,1}$ gives an annulus. In this annulus, since all the other cycles in Γ_1 and Γ_2 are disjoint from these, the existence of an isotopy between them follows from the case of the annulus in Sect. 7.1. After gluing back the boundaries together, this gives the desired isotopy of the torus.

Case 2: If there are at least two homotopy classes (modulo inversion) in Γ_1 , we pick two representatives, say $\gamma_{1,1}$ and $\gamma_{1,2}$. As two couple of lines with the same slopes pairwise can be moved one to the other with a translation, by doing a translation of the torus, we can assume that $\gamma_{1,1}$ and $\gamma_{2,1}$ lie in the neighborhood of the same geodesic, as well as $\gamma_{1,2}$ and $\gamma_{2,2}$, and furthermore that the crossing points between $\gamma_{1,1}$ and $\gamma_{1,2}$ lie in a ε -neighborhood of the corresponding crossing points between $\gamma_{2,1}$ and $\gamma_{2,2}$.¹¹

We are now in the same situation as in Sect. 3.2: Since both couples of cycles $\gamma_{1,1}$ and $\gamma_{2,1}$ and $\gamma_{1,2}$ and $\gamma_{2,2}$ lie in the ε -neighborhood of the same geodesic, if we take as stable families $\Gamma'_1 = \gamma_{1,1} \cup \gamma_{1,2}$ and $\Gamma'_2 = \gamma_{2,1} \cup \gamma_{2,2}$, we can similarly define corridors, as well as edge and vertex polygons. Then Proposition 2 holds with exactly the same proof. Since there is only one cycle of Γ'_1 in each corridor, Lemma 6 also holds trivially. Hence by applying the same techniques, we can conclude that there exists an ambient isotopy mapping Γ'_1 to Γ'_2 . Note that here, since the crossing points of the cycles in Γ'_1 have been matched, the isotopy we obtain is also pointwise.

Finally, cutting along these cycles cuts the surface into one or more disks, and the isotopy between Γ_1 and Γ_2 is obtained by applying Alexander's Lemma separately on each of these disks. $\square$

⁹We say that γ_1 and γ_2 are inverse homotopic if γ_1 is homotopic to γ_2^{-1} .

¹⁰Simple and non-contractible cycles which are homotopic are also isotopic, as proved by Epstein [15, Theorem 2.1].

¹¹This is not necessarily the case a priori, since a given crossing point c between $\gamma_{1,1}$ and $\gamma_{2,1}$ can be matched to another crossing point than $h(c)$. Note that this is why the result in the torus case is stronger than in the general case, in which the crossing points can *not* necessarily be matched.

8 Conclusion

We have given an optimal algorithm to test whether two graph embeddings of the same graph on an orientable surface are isotopic, where the vertices are allowed to move. On the other hand, the case of non-orientable surfaces remains open. Our characterization heavily relies on testing the existence of an oriented homeomorphism; it is not clear how to adapt this test in the non-orientable case.

Regarding the case of the punctured plane, the most obvious open question is to improve the running time of the algorithm. Bespamyatnikh [3, Theorem 7] describes an $O(n^{4/3} \text{polylog}(n))$ -time algorithm for testing path homotopy; however, it is not clear that this algorithm extends to the homotopy test for cycles. One obstacle for this extension is that our cycles are not simple and may make turns always in the same direction: the algorithm by Bespamyatnikh [3, Sect. 5.5] considers maximal subpaths that always turn in the same direction, but in our case such maximal subpaths may be cycles without “starting” and “ending” points, for which the same approach does not seem to work.

Also, we only test the existence of a topological isotopy: The edges are allowed to bend during the deformation. It is easy to see that, in the presence of obstacles, the existence of a topological isotopy between two straight-line embeddings does not imply the existence of a straight-line isotopy, in contrast to the case without obstacle [4, 20]. Could it be that, in such a situation, there exists a straight-line isotopy after splitting each edge in two (or a constant number of) segments? Computing such an isotopy efficiently may be a not easy task, but related techniques [30] might apply.

Finally, in both the surface model and the punctured plane model, computing shortest graph embeddings within a given isotopy class would be very interesting, and would generalize known results for computing shortest paths within a given homotopy class [2, 8–10, 13, 23], even though we expect the problem to be much harder.

Acknowledgement We thank Francis Lazarus and Julien Rivaud for sending us an advance copy of their paper [28], which inspired us for the construction of the stable family in Sect. 4, and we thank the anonymous referees for their careful reading of the manuscript and their useful suggestions.

References

1. Bentley, J.L., Ottmann, Th.: Algorithms for reporting and counting geometric intersections. In: IEEE Transactions on Computing, vol. 28, pp. 643–647 (1979)
2. Bespamyatnikh, S.: Computing homotopic shortest paths in the plane. *J. Algorithms* **49**, 284–303 (2003)
3. Bespamyatnikh, S.: Encoding homotopy of paths in the plane. In: LATIN 2004: Theoretical Informatics, the 6th Latin American Symposium. Lecture Notes in Computer Science, vol. 2976, pp. 329–338. Springer, Berlin (2004)
4. Bing, R.H., Starbird, M.: Linear isotopies in E^2 . *Trans. Am. Math. Soc.* **237**, 205–222 (1978)
5. Buser, P.: Geometry and Spectra of Compact Riemann Surfaces. Progress in Mathematics, vol. 106. Birkhäuser, Basel (1992)
6. Cabello, S., Liu, Y., Mantler, A., Snoeyink, J.: Testing homotopy for paths in the plane. *Discrete Comput. Geom.* **31**, 61–81 (2004)
7. Colin de Verdière, É.: Raccourcissement de courbes et décomposition de surfaces. Ph.D. thesis, Université Paris 7 (2003). English translation available at <http://www.di.ens.fr/~colin/03these.html>
8. Colin de Verdière, É., Erickson, J.: Tightening nonsimple paths and cycles on surfaces. *SIAM J. Comput.* **39**, 3784–3813 (2010)

9. Colin de Verdière, É., Lazarus, F.: Optimal system of loops on an orientable surface. *Discrete Comput. Geom.* **33**, 507–534 (2005)
10. Colin de Verdière, É., Lazarus, F.: Optimal pants decompositions and shortest homotopic cycles on an orientable surface. *J. ACM* **54** (2007). Article 18
11. de Graaf, M., Schrijver, A.: Making curves minimally crossing by Reidemeister moves. *J. Comb. Theory, Ser. B* **70**, 134–156 (1997)
12. Dey, T.K., Guha, S.: Transforming curves on surfaces. *J. Comput. Syst. Sci.* **58**, 297–325 (1999)
13. Efrat, A., Kobourov, S.G., Lubiw, A.: Computing homotopic shortest paths efficiently. *Comput. Geom.* **35**, 162–172 (2006)
14. Eppstein, D.: Dynamic generators of topologically embedded graphs. In: *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 599–608 (2003)
15. Epstein, D.B.A.: Curves on 2-manifolds and isotopies. *Acta Math.* **115**, 83–107 (1966)
16. Erickson, J., Whittlesey, K.: Greedy optimal homotopy and homology generators. In: *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1038–1046 (2005)
17. Erickson, J., Whittlesey, K.: Transforming curves on surfaces redux. In: *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1646–1655 (2013)
18. Farb, B., Margalit, D.: *A Primer on Mapping Class Groups*. Princeton University Press, Princeton (2011)
19. Feustel, C.D.: Homotopic arcs are isotopic. *Proc. Am. Math. Soc.* **17**, 891–896 (1966)
20. Gotsman, C., Surazhsky, V.: Guaranteed intersection-free polygon morphing. *Comput. Graph.* **25**, 67–75 (2001)
21. Hamidi-Tehrani, H.: On complexity of the word problem in braid groups and mapping class groups. *Topol. Appl.* **105**, 237–259 (2000)
22. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002). Available at <http://www.math.cornell.edu/~hatcher/>
23. Hershberger, J., Snoeyink, J.: Computing minimum length paths of a given homotopy class. *Comput. Geom.* **4**, 63–98 (1994)
24. Hirsch, M.W.: *Differential Topology*. Graduate Texts in Mathematics, vol. 33. Springer, Berlin (1994). Corrected reprint of the 1976 original
25. Kettner, L.: Using generic programming for designing a data structure for polyhedral surfaces. *Comput. Geom.* **13**, 65–90 (1999)
26. Ko, K.H., Lee, S.J., Cheon, J.H., Han, J.W., Kang, J.-s., Park, C.: New public-key cryptosystem using braid groups. In: *Advances in Cryptology—CRYPTO 2000*. Lecture Notes in Computer Science, vol. 1880, pp. 166–183. Springer, Berlin (2000)
27. Ladegaillerie, Y.: Classes d'isotopie de plongements de 1-complexes dans les surfaces. *Topology* **23**, 303–311 (1984)
28. Lazarus, F., Rivaud, J.: On the homotopy test on surfaces. In: *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 440–449 (2012)
29. Lins, S.: Graph-encoded maps. *J. Comb. Theory, Ser. B* **32**, 171–181 (1982)
30. Lubiw, A., Petrick, M.: Morphing planar graph drawings with bent edges. *Electron. Notes Discrete Math.* **31**, 45–48 (2008)
31. Mosher, L.: Mapping class groups are automatic. *Ann. Math.* **142**, 303–384 (1995)
32. Ringel, G.: Teilungen der Ebene durch Geraden oder topologische Geraden. *Math. Z.* **64**, 79–102 (1956)
33. Stillwell, J.: *Classical Topology and Combinatorial Group Theory*, 2nd edn. Springer, New York (1993)