

Untangling Two Systems of Noncrossing Curves*

Jiří Matoušek^{1,2}, Eric Sedgwick³, Martin Tancer^{1,4}, and Uli Wagner⁵

¹ Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic

² Institute of Theoretical Computer Science, ETH Zürich, 8092 Zürich, Switzerland

³ School of CTI, DePaul University, 243 S. Wabash Ave, Chicago, IL 60604, USA

⁴ Institutionen för matematik, KTH, 100 44 Stockholm, Sweden

⁵ IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria

Abstract. We consider two systems $(\alpha_1, \dots, \alpha_m)$ and $(\beta_1, \dots, \beta_n)$ of curves drawn on a compact two-dimensional surface \mathcal{M} with boundary. Each α_i and each β_j is either an arc meeting the boundary of \mathcal{M} at its two endpoints, or a closed curve. The α_i are pairwise disjoint except for possibly sharing endpoints, and similarly for the β_j . We want to “untangle” the β_j from the α_i by a self-homeomorphism of \mathcal{M} ; more precisely, we seek an homeomorphism $\varphi: \mathcal{M} \rightarrow \mathcal{M}$ fixing the boundary of \mathcal{M} pointwise such that the total number of crossings of the α_i with the $\varphi(\beta_j)$ is as small as possible. This problem is motivated by an application in the algorithmic theory of embeddings and 3-manifolds.

We prove that if \mathcal{M} is planar, i.e., a sphere with $h \geq 0$ boundary components (“holes”), then $O(mn)$ crossings can be achieved (independently of h), which is asymptotically tight, as an easy lower bound shows. In general, for an arbitrary (orientable or nonorientable) surface \mathcal{M} with h holes and of (orientable or nonorientable) genus $g \geq 0$, we obtain an $O((m+n)^4)$ upper bound, again independent of h and g .

Keywords: Curves on 2-manifolds, simultaneous planar drawings, Lickorish’s theorem.

1 Introduction

Let \mathcal{M} be a surface, by which we mean a two-dimensional compact manifold with (possibly empty) boundary $\partial\mathcal{M}$.

By the classification theorem for surfaces, if \mathcal{M} is orientable, then \mathcal{M} is homeomorphic to a sphere with $h \geq 0$ holes and $g \geq 0$ attached handles; the number g is also called the *orientable genus* of \mathcal{M} . If \mathcal{M} is nonorientable, then it is homeomorphic to a sphere with $h \geq 0$ holes and with $g \geq 0$ *cross-caps*; in this case, the integer g is known as the *nonorientable genus* of \mathcal{M} . In the sequel, the word

* Research supported by the ERC Advanced Grant No. 267165. Moreover, the research of J.M. was also partially supported by Grant GRADR Eurogiga GIG/11/E023, the research of M.T. was supported by a Göran Gustafsson postdoctoral fellowship, and the research of U.W. was supported by the Swiss National Science Foundation (Grant SNSF-PP00P2-138948)

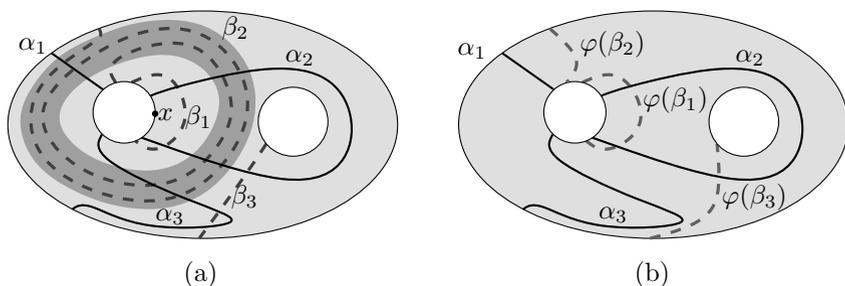


Fig. 1. Systems A and B of curves on a surface \mathcal{M} , with $g = 0$ and $h = 3$ (a), and a re-drawing of B via a ∂ -automorphism φ reducing the number of intersections (b).

“genus” will mean orientable genus for orientable surfaces and nonorientable genus for nonorientable surfaces.

We will consider curves in \mathcal{M} that are *properly embedded*, i.e., every curve is either a simple arc meeting the boundary $\partial\mathcal{M}$ exactly at its two endpoints, or a simple closed curve avoiding $\partial\mathcal{M}$. An *almost-disjoint system of curves* in \mathcal{M} is a collection $A = (\alpha_1, \dots, \alpha_m)$ of curves that are pairwise disjoint except for possibly sharing endpoints.

In this paper we consider the following problem: We are given two almost-disjoint systems $A = (\alpha_1, \dots, \alpha_m)$ and $B = (\beta_1, \dots, \beta_n)$ of curves in \mathcal{M} , where the curves of B intersect those of A possibly very many times, as in Fig. 1(a). We would like to “redraw” the curves of B in such a way that they intersect those of A as little as possible.

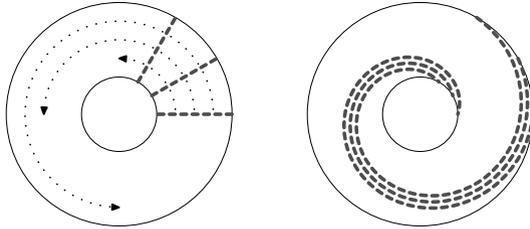
We consider re-drawings only in a restricted sense, namely, induced by ∂ -automorphisms of \mathcal{M} , where a ∂ -automorphism is a homeomorphism $\varphi: \mathcal{M} \rightarrow \mathcal{M}$ that fixes the boundary $\partial\mathcal{M}$ pointwise. Thus, given the α_i and the β_j , we are looking for a ∂ -automorphism φ such that the number of intersections (crossings) between $\alpha_1, \dots, \alpha_m$ and $\varphi(\beta_1), \dots, \varphi(\beta_n)$ is as small as possible (where sharing endpoints does not count). We call this minimum number of crossings achievable through any choice of φ the *entanglement number* of the two systems A and B .

In the orientable case, let $f_{g,h}(m, n)$ denote the maximum entanglement number of any two systems $A = (\alpha_1, \dots, \alpha_m)$ and $B = (\beta_1, \dots, \beta_n)$ of almost-disjoint curves on an orientable surface of genus g with h holes. Analogously, we define $\hat{f}_{g,h}(m, n)$ as the maximum entanglement number of any two systems A and B of m and n curves, respectively, on a nonorientable surface of genus g with h holes. It is easy to see that f and \hat{f} are nondecreasing in m and n , which we will often use in the sequel.

To give the reader some intuition about the problem, let us illustrate which re-drawings are possible with a ∂ -automorphism and which are not. In the example of Fig. 1, it is clear that the two crossings of β_3 with α_3 can be avoided by sliding β_3 aside.¹ It is perhaps less obvious that the crossings of β_2 can also be

¹ This corresponds to an *isotopy* of the surface that fixes the boundary pointwise.

eliminated: To picture a suitable ∂ -automorphism, one can think of an annular region in the interior of \mathcal{M} , shaded darkly in Fig. 1 (a), that surrounds the left hole and β_1 and contains most of the spiral formed by β_2 . Then we cut \mathcal{M} along the outer boundary of that annular region, twist the region two times (so that the spiral is unwound), and then we glue the outer boundary back. Here is an example of a single twist of an annulus; straight-line curves on the left are transformed to spirals on the right .



On the other hand, it is impossible to eliminate the crossings of β_1 or β_3 with α_2 by a ∂ -automorphism. For example, we cannot re-route β_1 to go around the right hole and thus avoid α_2 , since this re-drawing is not induced by any ∂ -automorphism φ : indeed, β_1 separates the point x on the boundary of left hole from the right hole, whereas α_2 does not separate them; therefore, the curve α_2 has to intersect $\varphi(\beta_1)$ at least twice, once when it leaves the component containing x and once when it returns to this component.

A rather special case of our problem, with $m = n = 1$ and only closed curves, was already considered by Lickorish [Lic62], who showed that the intersection of a pair of simple closed curves can be simplified via Dehn twists (and thus a ∂ -automorphism) so that they meet at most twice (also see Stillwell [Sti80]). The case with $m = 1$, n arbitrary, only closed curves, and \mathcal{M} possibly nonorientable was proposed in 2010 as a Mathoverflow question [Huy10] by T. Huynh. In an answer A. Putman proposes an approach via the “change of coordinates principle” (see, e.g., [FM11, Sec. 1.3]), which relies on the classification of 2-dimensional surfaces—we will also use it at some points in our argument.

The results. A natural idea for bounding $f_{g,h}(m, n)$ and $\hat{f}_{g,h}(m, n)$ is to proceed by induction, employing the change of coordinates principle mentioned above. This does indeed lead to finite bounds, but the various induction schemes we have tried always led to bounds at least exponential in one of m, n . Independently of our work, Geelen, Huynh, and Richter [GHR13] also recently proved bounds of this kind; see the discussion below. Partially influenced by the results on exponentially many intersections in representations of string graphs and similar objects (see [KM91,SSŠ03]), we first suspected that an exponential behavior might be unavoidable. Then, however, we found, using a very different approach, that polynomial bounds actually do hold.

For planar \mathcal{M} , i.e., $g = 0$, we obtain an asymptotically tight bound:

Theorem 1. *For planar \mathcal{M} , we have $f_{0,h}(m, n) = O(mn)$, independent of h .*

Here and in the sequel, the constants implicit in the O -notation are absolute, independent of g and h .

A simple example providing a lower bound of $2mn$ is obtained, e.g., by replicating α_2 in Fig. 1 m times and β_1 n times.

In general, we obtain the following bounds:

- Theorem 2.** (i) *For the orientable case, $f_{g,h}(m, n) = O((m + n)^4)$.*
 (ii) *For the nonorientable case, $\hat{f}_{g,h}(m, n) = O((m + n)^4)$.*

Both parts of Theorems 2 are derived from the planar case, Theorem 1. In the orientable case, we use the following results on genus reduction. For a more convenient notation, let us set $L = \max(m, n)$.

Proposition 1 (Orientable genus reductions).

- (i) *For $g > L$, we have $f_{g,h}(m, n) \leq f_{L,g+h-L}(m, n)$.*
 (ii) *$f_{g,h}(m, n) \leq f_{0,h+1}(cg(m + g), cg(n + g))$ for a suitable constant $c > 0$.*

Theorem 2 (i), the orientable case, follows immediately from Proposition 1 and the planar bound.

In the nonorientable case, Theorem 2 (ii) is derived in two steps. First, analogous to Proposition 1 (i), we have the following reduction:

Proposition 2 (Nonorientable genus reduction). *For $g > 4L + 2$, we have $\hat{f}_{g,h}(m, n) \leq \hat{f}_{g',h'}(m, n)$, where $g' = 4L + 2 - (g \bmod 2)$ and $h' = h + \lceil g/2 \rceil - 2L - 1$.*

The second step is a reduction to the orientable case.

Proposition 3 (Orientability reduction). *There is a constant c such that $\hat{f}_{g,h}(m, n) \leq f_{\lfloor (g-1)/2 \rfloor, h+1+(g \bmod 2)}(c(g + m), c(g + n))$.*

Table 1 summarizes the proof of Theorem 2.

Motivation. We were led to the question concerning untangling curves on surfaces while working on a project on 3-manifolds and embeddings. Specifically, we are interested in an algorithm for the following problem: given a 3-manifold M with boundary, does M embed in the 3-sphere? A special case of this problem, with the boundary of M a torus, was solved in [JS03]. The problem is motivated, in turn, by the question of algorithmically testing the embeddability of a 2-dimensional simplicial complex in \mathbb{R}^3 ; see [MTW11].

In our current approach, which has not yet been completely worked out, we need just a finite bound on $f_{g,h}(m, n)$. However, we consider the problem investigated in this paper interesting in itself and contributing to a better understanding of combinatorial properties of curves on surfaces.

As mentioned above, the question studied in the present paper has also been investigated independently by Geelen, Huynh, and Richter [GHR13], with a rather different and very strong motivation stemming from the theory of graph minors, namely the question of obtaining explicit upper bounds for the graph minor algorithms of Robertson and Seymour. Phrased in the language of the present paper, Geelen et al. [GHR13, Theorem 2.1] show that $f_{g,h}(m, n)$ and $\hat{f}_{g,h}(m, n)$ are both bounded by $n3^m$.

Table 1. A summary of the proof

1. For a planar surface, temporarily remove the holes not incident to any α_i or β_j , and contract the remaining “active” holes, augment the resulting planar graphs to make them 3-connected. Make a simultaneous plane drawing of the resulting planar graphs G_1 and G_2 with every edge of G_1 intersecting every edge of G_2 at most $O(1)$ times. Decontract the active holes and put the remaining holes back into appropriate faces (Theorem 1; Section 2).
2. If the genus is larger than $c(m+n)$, find handles or cross-caps avoided by the α_i and β_j , temporarily remove them, untangle the α_i and β_j , and put the handles or cross-caps back (Propositions 1 (i) and 2; the proofs are omitted from this extended abstract).
3. If the surface is nonorientable, make it orientable by cutting along a suitable curve that intersects the α_i and β_j at most $O(m+n)$ times, untangle the resulting pieces of the α_i and β_j , and glue back (Proposition 3).
4. Make the surface planar by cutting along a suitable system of curves (canonical system of loops), untangle the resulting pieces of the α_i and β_j , and glue back (Proposition 1 (ii)).

2 Planar Surfaces

In this section we prove Theorem 1. In the proof we use the following basic fact (see, e.g., [MT01]).

Lemma 1. *If G is a maximal planar simple graph (a triangulation), then for every two planar drawings of G in S^2 there is an automorphism ψ of S^2 converting one of the drawings into the other (and preserving the labeling of the vertices and edges). Moreover, if an edge e is drawn by the same arc in both of the drawings, w.l.o.g. we may assume that ψ fixes it pointwise.*

Let us introduce the following piece of terminology. Let G be as in the lemma, and let D_G, D'_G be two planar drawings of G . We say that D_G, D'_G are *directly equivalent* if there is an orientation-preserving automorphism of S^2 mapping D_G to D'_G , and we call D_G, D'_G *mirror-equivalent* if there is an orientation-reversing automorphism of S^2 converting D_G into D'_G .

We will also rely on a result concerning simultaneous planar embeddings; see [BKR12]. Let V be a vertex set and let $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ be two planar graphs on V . A planar drawing D_{G_1} of G_1 and a planar drawing D_{G_2} of G_2 are said to form a *simultaneous embedding* of G_1 and G_2 if each vertex $v \in V$ is represented by the same point in the plane in both D_{G_1} and D_{G_2} .

We note that G_1 and G_2 may have common edges, but they are not required to be drawn in the same way in D_{G_1} and in D_{G_2} . If this requirement is added, one speaks of a *simultaneous embedding with fixed edges*. There are pairs of planar graphs known that do not admit any simultaneous embedding with fixed edges

(and consequently, no simultaneous straight-line embedding). An important step in our approach is very similar to the proof of the following result.

Theorem 3 (Erten and Kobourov [EK05]). *Every two planar graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ admit a simultaneous embedding in which every edge is drawn as a polygonal line with at most 3 bends.*

We will need the following result, which follows easily from the proof given in [EK05]. For the reader's convenience, instead of just pointing out the necessary modifications, we present a full proof.

Theorem 4. *Every two planar graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ admit a simultaneous, piecewise linear embedding in which every two edges e_1 of G_1 and e_2 of G_2 intersect at least once and at most C times, for a suitable constant C .²*

In addition, if both G_1 and G_2 are maximal planar graphs, let us fix a planar drawing D'_{G_1} of G_1 and a planar drawing D'_{G_2} of G_2 . The planar drawing of G_1 in the simultaneous embedding can be required to be either directly equivalent to D'_{G_1} , or mirror-equivalent to it, and similarly for the drawing of G_2 (each of the four combinations can be prescribed).

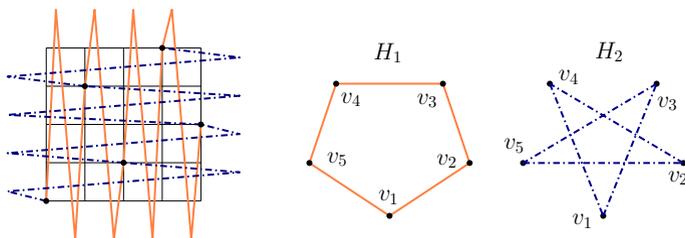
Proof. For the beginning, we assume that both graphs are Hamiltonian. Later on, we will drop this assumption.

Let v_1, v_2, \dots, v_n be the order of the vertices as they appear on (some) Hamiltonian cycle H_1 of G_1 . Since the vertex set V is common for G_1 and G_2 , there is a permutation $\pi \in S(n)$ such that $v_{\pi(1)}, \dots, v_{\pi(n)}$ is the order of the vertices as they appear on some Hamiltonian cycle H_2 of G_2 .

We draw the vertex v_i in the grid point $p_i = (i, \pi(i))$, $i = 1, 2, \dots, n$. Let S be the square $[1, n] \times [1, n]$. A *bispiked* curve is an x -monotone polygonal curve with two bends such that it starts inside S ; the first bend is above S , the second bend is below S and it finishes in S again.

The $n - 1$ edges $v_i v_{i+1}$, of H_1 , $i = 1, 2, \dots, n - 1$, are drawn as bispiked curves starting in p_i and finishing in p_{i+1} . In order to distinguish edges and their drawings, we denote these bispiked curves by $c(i, i + 1)$.

Similarly, we draw the edges $v_{\pi(i)} v_{\pi(i+1)}$ of H_2 , $i = 1, 2, \dots, n - 1$, as y -monotone analogs of bispiked curves, where the first bend is on the left of S and the second is on the right of S ; here is an example:



² An obvious bound from the proof is $C \leq 36$, since every edge in this embedding is drawn using at most 5 bends. By a more careful inspection, one can easily get $C \leq 25$, and a further improvement is probably possible.

We continue only with description of how to draw G_1 ; G_2 is drawn analogously with the grid rotated by 90 degrees.

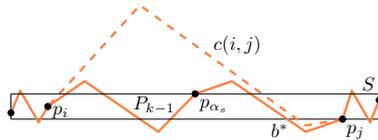
Let D'_{G_1} be a planar drawing of G_1 . Every edge from E_1 that is not contained in H_1 is drawn either inside D'_{H_1} or outside. Thus, we split $E_1 \setminus E(H_1)$ into two sets E'_1 and E''_1 .

Let P_0 be the polygonal path obtained by concatenation of the curves $c(1, 2), c(2, 3), \dots, c(n - 1, n)$. Now our task is to draw the edges of $E'_1 \cup \{v_1v_n\}$ as bispiked curves, all above P_0 , and then the edges of E''_1 below P_0 .

We start with E'_1 and we draw edges from it one by one, in a suitably chosen order, while keeping the following properties.

- (P1) Every edge v_iv_j , where $i < j$, is drawn as a bispiked curve $c(i, j)$ starting in p_i and ending in p_j .
- (P2) The x -coordinate of the second bend of $c(i, j)$ belongs to the interval $[j - 1, j]$.
- (P3) The polygonal curve P_k that we see from above after drawing the k th edge is obtained as a concatenation of some curves $c(1, i_1), c(i_1, i_2), \dots, c(i_\ell, n)$.

Here is an illustration; the square S is deformed for the purposes of the drawing:



Initially, before drawing the first edge, the properties are obviously satisfied.

Let us assume that we have already drawn $k - 1$ edges of E'_1 , and let us focus on drawing the k th edge. Let $e = v_iv_j \in E'_1$ be an edge that is not yet drawn and such that all edges below e are already drawn, where “below e ” means all edges $v_{i'}v_{j'} \in E'_1$ with $i \leq i' < j' \leq j$, $(i, j) \neq (i', j')$. (This choice ensures that we will draw all edges of E'_1 .)

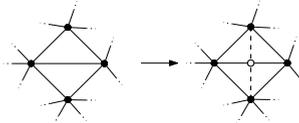
Since D'_{G_1} is a planar drawing, we know that there is no edge $v_{i'}v_{j'} \in E'_1$ with $i < i' < j < j'$ or $i' < i < j' < j$, and so the points p_i and p_j have to belong to P_{k-1} . The subpath P' of P_{k-1} between p_i and p_j is the concatenation of curves $c(i, \alpha_1), c(\alpha_1, \alpha_2), \dots, c(\alpha_s, j)$ as in the inductive assumptions. In particular, the x -coordinate of the second bend b^* of $c(\alpha_s, j)$ belongs to the interval $[j - 1, j]$. We draw $c(i, j)$ as follows: The second bend of $c(i, j)$ is slightly above b^* but still below the square S . The first bend of S is sufficiently high above S (with the x -coordinate somewhere between i and $j - 1$) so that the resulting bispiked curve $c(i, j)$ does not intersect P_{k-1} . The properties (P1) and (P2) are obviously satisfied by the construction. For (P3), the path P_k is obtained from P_{k-1} by replacing P' with $c(i, j)$.

After drawing the edges of E'_1 , we draw v_1v_n in the same way. Then we draw the edges of E''_1 in a similar manner as those of E'_1 , this time as bispiked curves below P_0 . This finishes the construction for Hamiltonian graphs.

Now we describe how to adjust this construction for non-Hamiltonian graphs, in the spirit of [EK05].

First we add edges to G_1 and G_2 so that they become planar triangulations. This step does not affect the construction at all, except that we remove these edges in the final drawing.

Next, we subdivide some of the edges of G_i with *dummy* vertices. Moreover, we attach two new *extra* edges to each dummy vertex, as in the following illustration:



By choosing the subdivided edges suitably, one can obtain a 4-connected, and thus Hamiltonian, graph; see [EK05, Proof of Theorem 2] for details (this idea previously comes from [KW02]). An important property of this construction is that each edge of G_i is subdivided at most once.

In this way, we obtain new Hamiltonian graphs G'_1 and G'_2 , for which we want to construct a simultaneous drawing as in the first part of the proof. A little catch is that G'_1 and G'_2 do not have same vertex sets, but this is easy to fix. Let d_i be the number of dummy vertices of G'_i , $i = 1, 2$, and say that $d_1 \geq d_2$. We pair the d_2 dummy vertices of G'_2 with some of the dummy vertices of G'_1 . Then we iteratively add $d_1 - d_2$ new triangles to G'_2 , attaching each of them to an edge of a Hamiltonian cycle. This operation keeps Hamiltonicity and introduces $d_1 - d_2$ new vertices, which can be matched with the remaining $d_1 - d_2$ dummy vertices in G'_1 .

After drawing resulting graphs, we remove all extra dummy vertices and extra edges added while introducing dummy vertices. An original edge e that was subdivided by a dummy vertex is now drawn as a concatenation of two bispiked curves. Therefore, each edge is drawn with at most 5 bends.

Two edges with 5 bends each may in general have at most 36 intersections, but in our case, there can be at most 25 intersections, since the union of the two segments before and after a dummy vertex is both x -monotone and y -monotone.

Because of the bispiked drawing of all edges, it is also clear that every edge of G_1 crosses every edge of G_2 at least once.

Finally, the requirements on directly equivalent or mirror-equivalent drawings can easily be fulfilled by interchanging the role of top and bottom in the drawing of G_1 or left and right in the drawing of G_2 . Theorem 4 is proved. \square

Proof of Theorem 1. Let a planar surface \mathcal{M} and the curves $\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_n$ be given; we assume that \mathcal{M} is a subset of S^2 . From this we construct a set V of $O(m+n)$ vertices in S^2 and planar drawings D_{G_1} and D_{G_2} of two simple graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ in S^2 , as follows.

1. We put all endpoints of the α_i and of the β_j into V .

2. We choose a new vertex in the interior of each α_i and each β_j , or two distinct vertices if α_i or β_j is a loop with a single endpoint, or three vertices if α_i or β_j is a closed curve, and we add all of these vertices to V . These new vertices are all distinct and do not lie on any curves other than where they were placed.
3. If the boundary of a hole in \mathcal{M} already contains a vertex introduced so far, we add more vertices so that it contains at least 3 vertices of V . This finishes the construction of V .
4. To define the edge set $E_1 = E(G_1)$ and the planar drawing D_{G_1} , we take the portions of the curves $\alpha_1, \dots, \alpha_m$ between consecutive vertices of V as edges of E_1 . Similarly, we make the arcs of the boundaries of the holes into edges in E_1 ; these will be called the *hole edges*. By the choice of the vertex set V above, this yields a simple plane graph.
5. Then we add new edges to E_1 so that we obtain a drawing D_{G_1} in S^2 of a maximal planar simple graph G_1 (i.e., a triangulation) on the vertex set V . While choosing these edges, we make sure that all holes containing no vertices of G lie in faces of D_{G_1} adjacent to some of the α_i . New edges drawn in the interior of a hole are also called *hole edges*.
6. We construct $G_2 = (V, E_2)$ and D_{G_2} analogously, using the curves β_1, \dots, β_m . We make sure that all hole edges are common to G_1 and G_2 .

After this construction, each hole of \mathcal{M} contains either no vertex of V on its boundary or at least three vertices. In the former case, we speak of an *inner hole*, and in the latter case, of a *subdivided hole*. A face f of D_{G_1} or D_{G_2} is a *non-hole face* if it is not contained in a subdivided hole. An inner hole H has its *signature*, which is a pair (f_1, f_2) , where f_1 is the unique non-hole face of D_{G_1} containing H , and f_2 is the unique non-hole face of D_{G_2} containing H .³ By the construction, each f_1 appearing in a signature is adjacent to some α_i , and each f_2 is adjacent to some β_j .

In the following claim, we will consider different drawings D'_{G_1} and D'_{G_2} for G_1 and G_2 . By Lemma 1, the faces of D_{G_1} are in one-to-one correspondence with the faces of D'_{G_1} . For a face f_1 of D_{G_1} , we denote the corresponding face by f'_1 , and similarly for a face f_2 of D_{G_2} and f'_2 .

Claim 5. *The graphs G_1 and G_2 as above have planar drawings D'_{G_1} and D'_{G_2} , respectively, that form a simultaneous embedding in which each edge of G_1 crosses each edge of G_2 at most C times, for a suitable constant C ; moreover, D'_{G_1} is directly equivalent to D_{G_1} ; D'_{G_2} is directly equivalent to D_{G_2} ; all hole edges are drawn in the same way in D'_{G_1} and D'_{G_2} ; and whenever (f_1, f_2) is a signature of an inner hole, the interior of the intersection $f'_1 \cap f'_2$ is nonempty.*

We postpone the proof of Claim 5, and we first finish the proof of Theorem 1 assuming this claim.

³ Classifying inner holes according to the signature helps us to obtain a bound independent on the number of holes. Inner holes with same signature are all treated in the same way, independent of their number.

For each inner hole H with signature (f_1, f_2) , we introduce a closed disk B_H in the interior of $f'_1 \cap f'_2$. We require that these disks are pairwise disjoint. In sequel, we consider holes as subsets of S^2 homeomorphic to closed disks (in particular, a hole H intersects \mathcal{M} in ∂H).

Claim 6. *There is an orientation-preserving automorphism φ_1 of S^2 transforming every inner hole H to B_H and D_{G_1} to D'_{G_1} .*

Proof. Using Lemma 1 again, there is an orientation-preserving automorphism ψ_1 transforming D_{G_1} into D'_{G_1} (since D_{G_1} and D'_{G_1} are directly equivalent).

Let f_1 be a face of D_{G_1} . The interior of f'_1 contains images $\psi_1(H)$ of all holes H with signature (f_1, \cdot) , and it also contains the disks B_H for these holes. Therefore, there is a boundary- and orientation-preserving automorphism of f'_1 that maps each $\psi_1(H)$ to B_H .

By composing these automorphisms on every f'_1 separately, we have an orientation-preserving automorphism ψ_2 fixing D'_{G_1} and transforming each $\psi_1(H)$ to B_H . The required automorphism is $\varphi_1 = \psi_2\psi_1$. □

Claim 7. *There is an orientation-preserving automorphism φ_2 of S^2 that fixes hole edges (of subdivided holes), fixes B_H for every inner hole H , and transforms $\varphi_1(D_{G_2})$ to D'_{G_2} .*

Proof. By Lemma 1 there is an orientation-preserving automorphism ψ_3 of S^2 that fixes hole edges and transforms $\varphi_1(D_{G_2})$ to D'_{G_2} .

If an inner hole H has a signature (\cdot, f_2) , then both $\psi_3(B_H)$ and B_H belong to the interior of f'_2 . Therefore, as in the proof of the previous claim, there is an orientation-preserving homeomorphism ψ_4 that fixes D'_{G_2} and transforms $\psi_3(B_H)$ to B_H . We can even require that $\psi_4\psi_3$ is identical on B_H . We set $\varphi_2 := \psi_4\psi_3$. □

To finish the proof of Theorem 1, we set $\varphi = \varphi_1^{-1}\varphi_2\varphi_1$. We need that φ fixes the holes (inner or subdivided) and that $\alpha_1, \dots, \alpha_m$ and $\varphi(\beta_1), \dots, \varphi(\beta_m)$ have $O(mn)$ intersections. It is routine to check all the properties:

If H is a hole (inner or subdivided), then φ_2 fixes $\partial\varphi_1(H)$. Therefore, φ also restricts to a ∂ -automorphism of \mathcal{M} .

The collections of curves $\alpha_1, \dots, \alpha_m$ and $\varphi(\beta_1), \dots, (\beta_m)$ have same intersection properties as the collections $\varphi_1(\alpha_1), \dots, \varphi_1(\alpha_m)$ and $\varphi_2(\varphi_1(\beta_1)), \dots, \varphi_2(\varphi_1(\beta_m))$. Since each α_i and each β_j was subdivided at most three times in the construction, by Claims 5, 6, and 7, these collections have at most $O(mn)$ intersections. The proof of the theorem is finished, except for Claim 5.

Proof of Claim 5. Given G_1 and G_2 , we form auxiliary planar graphs \tilde{G}_1 and \tilde{G}_2 on a vertex set \tilde{V} by contracting all hole edges and removing the resulting loops and multiple edges. We note that a loop cannot arise from an edge that was a part of some α_i or β_j .

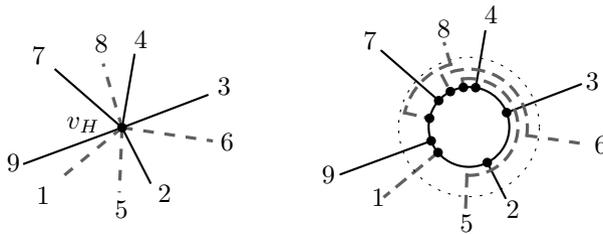
Then we consider planar drawings $D_{\tilde{G}_1}$ and $D_{\tilde{G}_2}$ forming a simultaneous embedding as in Theorem 4, with each edge of \tilde{G}_1 crossing each edge of \tilde{G}_2 at least once and most a constant number of times.

Let $v_H \in \tilde{V}$ be the vertex obtained by contracting the hole edges on the boundary of a hole H . Since the drawings $D_{\tilde{G}_1}$ and $D_{\tilde{G}_2}$ are piecewise linear, in a sufficiently small neighborhood of v_H the edges are drawn as radial segments.

We would like to replace v_H by a small circle and thus turn the drawings $D_{\tilde{G}_1}, D_{\tilde{G}_2}$ into the required drawings D'_{G_1}, D'_{G_2} . But a potential problem is that the edges in $D_{\tilde{G}_1}, D_{\tilde{G}_2}$ may enter v_H in a wrong cyclic order.

We claim that the edges in $D_{\tilde{G}_1}$ entering v_H have the same cyclic ordering around v_H as the corresponding edges around the hole H in the drawing D_{G_1} . Indeed, by contracting the hole edges in the drawing D_{G_1} , we obtain a planar drawing $D_{\tilde{G}_1}^*$ of \tilde{G}_1 in which the cyclic order around v_H is the same as the cyclic order around H in D_{G_1} . Since \tilde{G}_1 was obtained by edge contractions from a maximal planar graph, it is maximal as well (since an edge contraction cannot create a non-triangular face), and its drawing is unique up to an automorphism of S^2 (Lemma 1). Hence the cyclic ordering of edges around v_H in $D_{\tilde{G}_1}$ and in $D_{\tilde{G}_1}^*$ is either the same (if $D_{\tilde{G}_1}$ and $D_{\tilde{G}_1}^*$ are directly equivalent), or reverse (if $D_{\tilde{G}_1}$ and $D_{\tilde{G}_1}^*$ are mirror-equivalent). However, Theorem 4 allows us to choose the drawing $D_{\tilde{G}_1}$ so that it is directly equivalent to $D_{\tilde{G}_1}^*$, and then the cyclic orderings coincide. A similar consideration applies for the other graph G_2 .

The edges of $D_{\tilde{G}_1}$ may still be placed to wrong positions among the edges in $D_{\tilde{G}_2}$, but this can be rectified at the price of at most one extra crossing for every pair of edges entering v_H , as the following picture indicates (the numbering specifies the cyclic order of the edges around H in $D_{G_1} \cup D_{G_2}$):



It remains to draw the edges of G_1 and G_2 that became loops or multiple edges after the contraction of the hole edges. Loops can be drawn along the circumference of the hole, and multiple edges are drawn very close to the corresponding single edge.

In this way, every edge of G_1 still has at most a constant number of intersections with every edge of G_2 , and every two such edges intersect at least once unless at least one of them became a loop after the contraction. Consequently, whenever (f_1, f_2) is a signature of an inner hole, the corresponding faces f'_1 and f'_2 intersect. This finishes the proof. □

Acknowledgement. We would like to thank the authors of [GHR13] for making a draft of their paper available to us, and, in particular, T. Huynh for an e-mail correspondence.

References

- BKR12. Bläsius, T., Kobourov, S.G., Rutter, I.: Simultaneous embedding of planar graphs (2012) (preprint), <http://arxiv.org/abs/1204.5853>
- EK05. Erten, C., Kobourov, S.G.: Simultaneous embedding of planar graphs with few bends. *J. Graph Algorithms Appl.* 9(3), 347–364 (2005) (electronic)
- FM11. Farb, B., Margalit, D.: *A primer on mapping class groups*. Princeton University Press, Princeton (2011)
- GHR13. Geelen, J., Huynh, T., Richter, R.B.: Explicit bounds for graph minors (May 2013) (preprint), <http://arxiv.org/abs/1305.1451>
- Huy10. Huynh, T.: Removing intersections of curves in surfaces. *Mathoverflow* (2010), <http://mathoverflow.net/questions/33963/removing-intersections-of-curves-in-surfaces/33970#33970>
- JS03. Jaco, W., Sedgwick, E.: Decision problems in the space of Dehn fillings. *Topology* 42(4), 845–906 (2003)
- KM91. Kratochvíl, J., Matoušek, J.: String graphs requiring huge representations. *J. Combin. Theory Ser. B* 53(1), 1–4 (1991)
- KW02. Kaufmann, M., Wiese, R.: Embedding vertices at points: few bends suffice for planar graphs. *J. Graph Algorithms Appl.* 6(1), 115–129 (2002); (electronic) *Graph drawing and representations*, Prague (1999)
- Lic62. Lickorish, W.B.R.: A representation of orientable combinatorial 3-manifolds. *Ann. Math* (2) 76, 531–540 (1962)
- MT01. Mohar, B., Thomassen, C.: *Graphs on Surfaces*. Johns Hopkins University Press, Baltimore (2001)
- MTW11. Matoušek, J., Tancer, M., Wagner, U.: Hardness of embedding simplicial complexes in \mathbb{R}^d . *J. Eur. Math. Soc.* 13(2), 259–295 (2011)
- SSŠ03. Schaefer, M., Sedgwick, E., Štefankovič, D.: Recognizing string graphs in NP. *J. Comput. Syst. Sci.* 67(2), 365–380 (2003)
- Sti80. Stillwell, J.: *Classical Topology and Combinatorial Group Theory*. Springer, New York (1980)