

SOME FAST ALGORITHMS FOR CURVES IN SURFACES

MARC LACKENBY

ABSTRACT. We present some algorithms that provide useful topological information about curves in surfaces. One of the main algorithms computes the geometric intersection number of two properly embedded 1-manifolds C_1 and C_2 in a compact orientable surface S . The surface S is presented via a triangulation or a handle structure, and the 1-manifolds are given in normal form via their normal coordinates. The running time is bounded above by a polynomial function of the number of triangles in the triangulation (or the number of handles in the handle structure), and the logarithm of the weight of C_1 and C_2 . This algorithm represents an improvement over previous work, since its running time depends polynomially on the size of the triangulation of S and it can deal with closed surfaces, unlike many earlier algorithms. Another algorithm, with similar bounds on its running time, can determine whether C_1 and C_2 are isotopic. We also present a closely related algorithm that can be used to place a standard 1-manifold into normal form.

1. INTRODUCTION

Surfaces occupy a central position in low-dimensional topology and geometry. They are mostly studied by means of properly embedded 1-manifolds (that is, simple closed curves and arcs, that are pairwise disjoint, and with the endpoints of the arcs in the boundary of the surface). A foundational result is that two properly embedded essential 1-manifolds C_1 and C_2 in a compact surface S can be isotoped into an essentially unique position where $|C_1 \cap C_2|$ is minimal. The *geometric intersection number* $i(C_1, C_2)$ is then $|C_1 \cap C_2|$. This minimal position for C_1 and C_2 is easy to compute in principle, because when C_1 and C_2 are not in minimal position, then there is a bigon or half-bigon (as shown in Figure 1) in their complement, and this can be used to reduce the number of points of intersection. However, when C_1 and C_2 intersect many times, the computation of their geometric intersection number is not at all straightforward. Nevertheless, the following result shows that it can be computed extremely efficiently.

Theorem 1.1 (GEOMETRIC INTERSECTION NUMBER). *Let \mathcal{T} be a triangulation for a compact orientable surface S . Let C_1 and C_2 be normal 1-manifolds properly embedded in S . Then there is an algorithm that provides the geometric intersection number of C_1 and C_2 . The running time of this algorithm is bounded above by a polynomial function of $|\mathcal{T}|$, $\log w(C_1)$ and $\log w(C_2)$.*

In this result, and throughout this paper, we require the surface S to be compact. It may be closed or have non-empty boundary. In this result, we also require S to be orientable, but in some later theorems, we can drop this hypothesis. Here, $|\mathcal{T}|$

The author was partially supported by EPSRC grant EP/Y004256/1. For the purpose of open access, the author has applied a CC BY public copyright licence to any author accepted manuscript arising from this submission.

is the number of triangles of \mathcal{T} . Recall that a 1-manifold C in S is *normal* if it intersects each triangle in properly embedded arcs, each of which has endpoints in the interior of distinct edges of the triangle. The number $w(C)$ is the *weight* of a 1-manifold C , which is its number of points of intersection with the 1-skeleton of \mathcal{T} . Note that in the case where \mathcal{T} is a 1-vertex triangulation for a closed surface S , then $|\mathcal{T}|$ is equal to $2 - 2\chi(S)$ and hence the upper bound on the running time depends polynomially on $\chi(S)$. In the case where S has non-empty boundary and \mathcal{T} has one vertex on each boundary component and no vertices in the interior of S , then the upper bound on the running time again depends polynomially on $\chi(S)$.

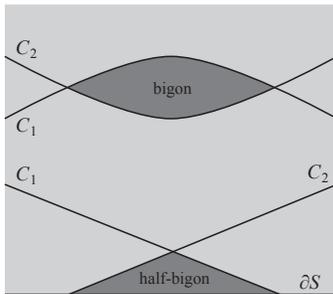


FIGURE 1. A bigon and a half-bigon for two properly embedded 1-manifolds C_1 and C_2

The computation of geometric intersection number has been studied in detail by several authors, including Bell and Webb [3], Dynnikov [7], Schaefer, Sedgwick and Štefankovič [14], Despré and Lazarus [5], Dubois [6] and Chang and de Mesmay [4]. In Section 1.1, we will compare those works with ours, highlighting what is new in Theorem 1.1.

In a similar vein, we have the following result.

Theorem 1.2 (ISOTOPY OF 1-MANIFOLDS). *Let \mathcal{T} be a triangulation for a compact orientable surface S . Let C_1 and C_2 be normal essential 1-manifolds properly embedded in S . Then there is an algorithm that determines whether C_1 and C_2 are isotopic. The running time of this algorithm is bounded above by a polynomial function of $|\mathcal{T}|$, $\log w(C_1)$ and $\log w(C_2)$.*

A similar result was proved by Schleimer [15, Theorem A.7], who provided a new solution to the conjugacy problem for $\pi_1(S)$. Two simple closed curves represent conjugate elements of $\pi_1(S)$ if and only if they are isotopic. Schleimer's algorithm takes, as its input, 'compressed' words in $\pi_1(S)$ and hence runs in polynomial time as a function of $\log w(C_1)$ and $\log w(C_2)$ with respect to a suitable triangulation \mathcal{T} . But how the running time in Schleimer's algorithm depends on the triangulation \mathcal{T} and, more generally on $\chi(S)$, is not explicitly addressed in [15]. In [13, Theorem 2(a)], another algorithm was given to determine whether two simple closed curves are isotopic, but it requires S to have non-empty boundary and the dependence of its running time on $\chi(S)$ is unclear.

Our next main result gives an algorithm that places a 1-manifold into normal form. For simplicity, we initially consider the case of surfaces with a given triangulation. Later, we will alternatively endow our surfaces with a handle structure.

A 1-manifold C properly embedded in a surface S with a triangulation \mathcal{T} is *standard* if it misses the vertices of \mathcal{T} , is transverse to the edges of \mathcal{T} and intersects each triangle in a collection of properly embedded arcs. A *simplifying disc* for C is a disc D lying in a triangle Δ of the triangulation such that

- (1) ∂D is the union of two arcs α and β that intersect at their endpoints;
- (2) $\alpha = D \cap C$;
- (3) $\beta = D \cap \partial\Delta$ is a subset of an edge of Δ .

See Figure 2. The *simplification number* $\text{simp}(C)$ is the number of simplifying discs for C .

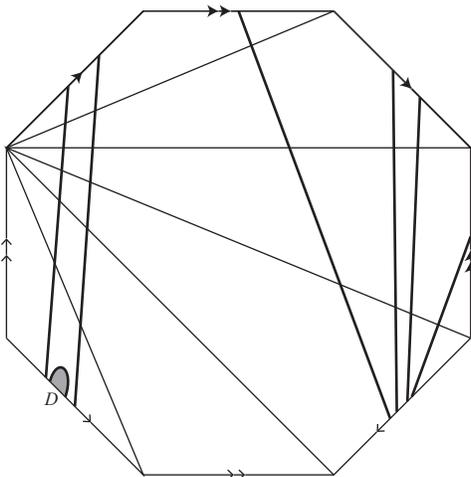


FIGURE 2. A standard curve in a triangulated surface, with a simplifying disc D

Theorem 1.3 (NORMALISE STANDARD 1-MANIFOLD IN A TRIANGULATION). *There is an algorithm that takes, as its input, the following data:*

- (1) a triangulation \mathcal{T} for a compact surface S ;
- (2) a standard 1-manifold C properly embedded in S ;

and provides, as its output, a normal 1-manifold \bar{C} that is obtained from C by an ambient isotopy and possibly removing some components that are parallel to arcs in ∂S or that bound discs in S . The running time for the algorithm is bounded above by a polynomial function of the number of triangles in \mathcal{T} , the simplification number of C and the logarithm of the weight of C .

In order for the statement of the theorem to be complete, we need to explain how a standard 1-manifold in a surface is entered into the algorithm. To do this, we first explain how to describe a system of disjoint properly embedded arcs C in a triangle Δ , where these arcs are disjoint from the vertices of Δ . Suppose that the edges of Δ are oriented in some way. The arcs intersect the edges some numbers of times, n_1 , n_2 and n_3 . Since each edge is oriented, we may identify these points of intersection with the i th edge with the integers $[1, n_i] \cap \mathbb{Z}$. We say that two arcs α_0 and α_1 of C are *parallel* if there is an embedding of $[0, 1] \times [0, 1]$ in Δ such that

- (1) $[0, 1] \times \{0\} = \alpha_0$,

- (2) $[0, 1] \times \{1\} = \alpha_1$,
- (3) $([0, 1] \times [0, 1]) \cap \partial\Delta = \{0, 1\} \times [0, 1]$,
- (4) this square is disjoint from the vertices of Δ ,
- (5) the intersection between $[0, 1] \times [0, 1]$ and each component of C is either empty or an arc of the form $[0, 1] \times \{t\}$.

This forms an equivalence relation on the components of C and we refer to the equivalence classes as *parallelism classes*. The endpoints of the arcs in a parallelism class consist of two intervals $[a, b] \cap \mathbb{Z}$ and $[c, d] \cap \mathbb{Z}$ in the intervals $[1, n_i] \cap \mathbb{Z}$ and $[1, n_j] \cap \mathbb{Z}$ corresponding to the relevant edges, and the arcs themselves determine an order-preserving or order-reversing bijection $[a, b] \cap \mathbb{Z} \rightarrow [c, d] \cap \mathbb{Z}$, which is termed a *pairing*. The collection of all these pairings is termed a *pairing system* for Δ . See Figure 3. It is clear that there is a simple procedure for determining whether a collection of pairings determines a pairing system. In fact, a collection of pairings gives a pairing system if and only if the following all hold:

- (1) the ranges and domains of all the pairings form a partition of $[1, n_1] \cap \mathbb{Z}$, $[1, n_2] \cap \mathbb{Z}$ and $[1, n_3] \cap \mathbb{Z}$;
- (2) a pairing is order-reversing if and only if the orientations of the edges containing the range and domain are consistent around $\partial\Delta$;
- (3) the ranges and domains of no two pairings interleave around $\partial\Delta$.

Thus, a standard 1-manifold C in a triangulated surface is given by the following data:

- (1) an orientation on each edge;
- (2) a non-negative integer $n(e)$ for each edge e ; this is the number of points of intersection between C and e ;
- (3) a pairing system for each face, where the number of endpoints of the arcs on each edge e is $n(e)$.

The above definition of a pairing system in a triangle readily generalises to the case of polygons, in an obvious way.

1.1. Comparison with previous work. As mentioned earlier, there has been a lot of previous work on the computation of geometric intersection number of curves in surfaces, which we now briefly summarise.

In [13], Schaefer, Sedgwick and Štefankovič used the theory of straight line programs to provide very efficient solutions to many computational problems about 1-manifolds properly embedded in surfaces. They built on this theory in [14] to provide an algorithm to compute the geometric intersection number of two properly embedded 1-manifolds. One significant difference between their work and ours is that they restricted to the case where S has non-empty boundary. It turns out when ∂S is non-empty, many algorithmic questions become somewhat easier. One way of understanding this is to note that the fundamental group of S is then free, and so closed curves then have a *unique* cyclically reduced representative in $\pi_1(S)$. Our algorithms, on the other hand, permit S to be closed. Another significant distinction between Theorem 1.1 and the work in [14] is that our algorithm has running time that depends only polynomially on $|\mathcal{T}|$.

An alternative approach to computing geometric intersection numbers was given by Bell and Webb in [3]. This relied on prior work of Bell [2], which provided a rapid method for modifying triangulations of surfaces. Given two 1-manifolds C_1 and C_2 properly embedded in S expressed as normal curves with respect to an ideal

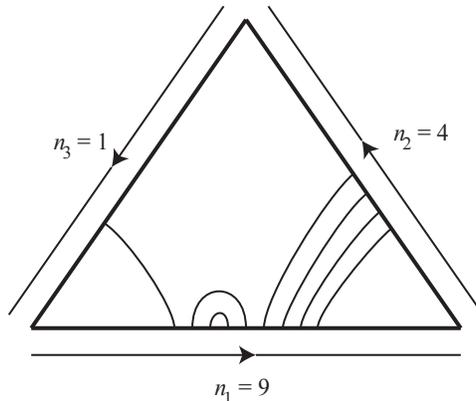


FIGURE 3. Shown is a system of arcs in triangle. This system is specified by picking orientations on the edges, as shown, and setting $n_1 = 9$, $n_2 = 4$ and $n_3 = 1$. The arcs are given by three pairings. For example, one set of arcs is given by the pairing $[2, 3] \cap \mathbb{Z} \rightarrow [4, 5] \cap \mathbb{Z}$, sending x to $7 - x$, and where the domain and range are both viewed as subsets of $[1, n_1] \cap \mathbb{Z}$.

triangulation \mathcal{T} , Bell and Webb were able to change \mathcal{T} to another ideal triangulation \mathcal{T}' that intersects C_1 in a very simple form. They could also keep track of how C_2 lies with respect to \mathcal{T}' , and then ensure that C_2 and C_1 have no complementary bigon. The geometric intersection number of C_1 and C_2 could then be directly read off. Like in [14], Bell and Webb require S to have non-empty boundary, since they make substantial use of the fact that an essential simple closed curve has a unique normal representative with respect to an ideal triangulation. Also as in [14], they make no analysis of how the running time depends on S . Nevertheless, the approach taken by Bell and Webb is possibly closest to the one we use, since we also use modifications of our given triangulation.

A third approach to geometric intersection numbers was by Dynnikov [7], who introduced a novel method of representing embedded curves on surfaces. He also used a novel method of expressing mapping classes of surfaces via matrices, so that composition of mapping classes can be computed via matrix multiplication. However, again, Dynnikov required the surface S to have non-empty boundary.

Despré and Lazarus [5] analysed a related but different problem. They considered 1-manifolds C in the surface S with $C \cap \partial S = \partial C$, but which might not be embedded. Such a curve has a well-defined self-intersection number, which is the minimal number of self-intersections among representatives with only double points, up to homotopy that keeps ∂C in ∂S . They give an efficient method for computing this self-intersection number. Despré and Lazarus had to focus on a connected 1-manifold. On the other hand, Chang and de Mesmay [4] and Dubois [6] considered the same problem, but for 1-manifolds with multiple components. By applying this to a 1-manifold that is the union of two embedded 1-manifolds C_1 and C_2 , this gives another method for computing their geometric intersection number. The algorithms of Despré and Lazarus, Dubois and Chang and de Mesmay have polynomial dependence of the genus of the surface and permit the surface to

be closed, but also depend polynomially on the initial number of self-intersection points of the given 1-manifold C . Hence, when applied to the union of two 1-manifolds C_1 and C_2 , the running time does not depend polynomially on $\log w(C_1)$ and $\log w(C_2)$.

The methods that we present in this paper rely (like [14, 3, 7]) very heavily on the assumption that the 1-manifolds are embedded. Nevertheless, there is some hope that our methods may be adapted and possibly combined with those of [5, 6, 4] to efficiently compute the geometric intersection number of two 1-manifolds C_1 and C_2 , which may be non-embedded but which have a relatively small number of self-intersections.

1.2. Structure of the paper. In Section 2, we introduce handle structures for surfaces and normal and standard 1-manifolds within them. We also state versions of Theorems 1.1 and 1.3 for handle structures instead of triangulations. In Section 3, we recall the algorithm of Agol, Hass and Thurston [1] and explain some of its uses in the study of normal and standard 1-manifolds. In Section 4, we examine the handle structure obtained by cutting along a standard 1-manifold, and explain how it can be computed efficiently. In Section 5, we prove Theorem 1.3 and its analogue for handle structures. This section is the heart of the paper, as its main ideas recur at several points later. In Section 6, we show how to place two 1-manifolds C and P into minimal position, via an isotopy of C , in the important case where P is simplicial in the triangulation. In Section 7, this is used, along with a theorem of the author and Yazdi [11], to complete the proof of Theorems 1.1 and 1.2. In Section 8, we consider a more general situation, where the surface S contains a ‘pattern’ which is a union of disjoint properly embedded 1-manifolds and graphs (with restrictions on the vertex degrees). We show how to arrange for a properly embedded 1-manifold to intersect a pattern ‘minimally’. The reason for studying patterns is that, in the analysis of hierarchies for 3-manifolds, the boundary of each 3-manifold in the hierarchy naturally inherits a pattern [9, 12].

2. HANDLE STRUCTURES AND NORMAL 1-MANIFOLDS

2.1. Handle structures. In Theorems 1.1 and 1.3, the surface S was given by a triangulation. This is because triangulations are a very typical way of presenting a surface. However, it is frequently useful to deal with handle structures instead. (See, for example, Definition 4.1 where a useful construction involving handle structures is presented, and which has no direct analogue for triangulations.) As we will see, there are versions of Theorems 1.1 and 1.3 for handle structures. In this section, we introduce some of the terminology for handle structures.

Definition 2.1. A *handle structure* for a surface S is a decomposition of S into three sets $\mathcal{H}^0, \mathcal{H}^1, \mathcal{H}^2$, subject to the following conditions:

- (1) \mathcal{H}^0 is a union of disjoint closed discs;
- (2) each component of \mathcal{H}^1 is a disc of the form $[0, 1] \times [0, 1]$ with $([0, 1] \times [0, 1]) \cap \mathcal{H}^0 = \{0, 1\} \times [0, 1]$;
- (3) each component of \mathcal{H}^2 is a disc D with $\partial D = D \cap (\mathcal{H}^0 \cup \mathcal{H}^1)$.

Each component of \mathcal{H}^i is called an *i-handle*. See, for example, Figure 4. For each 1-handle $[0, 1] \times [0, 1]$, the arc $[0, 1] \times \{1/2\}$ is its *core* and $\{1/2\} \times [0, 1]$ is its *co-core*.

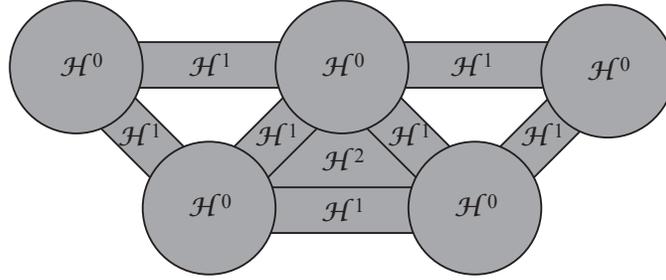


FIGURE 4. A handle structure on a 3-holed sphere

Definition 2.2. Let \mathcal{H} be a handle structure for a surface S . Its *size* $||\mathcal{H}||$ is the number of handles.

Remark 2.3. It is quite common to work instead with the number of 0-handles of \mathcal{H} , denoted $|\mathcal{H}|$. This is only a good measure of the complexity of \mathcal{H} when the 0-handles are constrained to fall into one of finitely many pre-arranged types. We do not wish to make that hypothesis in this paper, and we therefore work instead with $||\mathcal{H}||$.

Definition 2.4. Let \mathcal{H} be a handle structure for a surface S . Its *underlying cell structure* is the cell structure on S where each handle is a 2-cell. Each 1-cell arises either as a component of intersection between two handles or between a handle and ∂S . Each 0-cell arises either as a component of intersection between three handles or between two handles and ∂S .

For a cell structure \mathcal{X} and $i \in \{0, 1, 2\}$, \mathcal{X}^i denotes the set of i -cells.

Lemma 2.5. *The number of 1-cells in the underlying cell structure is $6|\mathcal{H}^1| \leq 6||\mathcal{H}||$.*

Proof. Each 1-cell lies in the boundary of a 0-handle or a 1-handle. The number of 1-cells in the boundary of a 1-handle but not the boundary of a 0-handle is $2|\mathcal{H}^1|$. The number of 1-cells in the boundary of a 0-handle is $4|\mathcal{H}^1|$. \square

A cell structure on a surface S is specified by giving a list of cells, together with their attaching maps onto the lower-dimensional cells. These attaching maps are specified as follows. Each 1-cell is oriented and its initial and terminal 0-cells are specified. The boundary of each 2-cell is specified as a sequence of 1-cells, together with information about whether each 1-cell is traversed consistently or inconsistently with the orientation.

A handle structure \mathcal{H} on a surface S is specified by giving its underlying cell structure \mathcal{X} , together with the index of each handle. The amount of data required is then bounded above by a polynomial function of $||\mathcal{H}||$ for the following reason. The attaching maps of the 1-cells require $2|\mathcal{X}^1| + ||\mathcal{H}||$ pieces of information, each of which is either the indexing number of a 0-cell or a spacing digit specifying a break between one 2-cell and the next. We may give each indexing number using at most $\log(||\mathcal{H}|| + 1)$ digits. The attaching maps of the 2-cells require at most $2|\mathcal{X}^1|$ pieces of information, each of which is the number of a 1-cell as well as its orientation. The number of 1-cells of \mathcal{X} is at most $6||\mathcal{H}||$, by Lemma 2.5. So the

amount of data required to specify \mathcal{H} is at most $O(\|\mathcal{H}\| \log(\|\mathcal{H}\| + 1))$ bits and at least $\|\mathcal{H}\|$ bits.

2.2. Standard and normal 1-manifolds.

Definition 2.6. Let S be a surface with a handle structure \mathcal{H} . Then a 1-manifold C properly embedded in S is *standard* if

- (1) it intersects each 0-handle in a collection of properly embedded arcs;
- (2) it intersects each 1-handle in a collection of properly embedded arcs that respect the handle's product structure and run parallel to the core of the handle;
- (3) it is disjoint from the 2-handles.

Definition 2.7. The *weight* $w(C)$ of a standard 1-manifold C is the number of components of $C \cap \mathcal{H}^1$.

Remark 2.8. For simple closed curves C , $w(C)$ is a reasonable measure of their complexity. However, when C has arc components, some of these may have weight zero. Therefore, $w(C) + |\partial C|$ is a better measure of the complexity of C .

Definition 2.9. Let \mathcal{H} be a handle structure for a compact surface S . Let \mathcal{X} be the associated cell structure of S , with a 2-cell for each handle. Then a *pairing system* for a standard 1-manifold C is the following:

- (1) it misses $H_0 \cap \mathcal{H}^1 \cap \mathcal{H}^2$ and $H_0 \cap \mathcal{H}^1 \cap \partial S$;
- (2) a choice of orientation on each 1-cell in \mathcal{X} ;
- (3) a non-negative integer for each 1-cell that gives the number of points of intersection between C and the 1-cell;
- (4) a pairing system on each 2-cell of \mathcal{H}^0 and \mathcal{H}^1 , that respects the product structure on each component of \mathcal{H}^1 , and that agrees on the boundary of the 2-cell with the data given in (2).

Definition 2.10. Let \mathcal{H} be a handle structure for a compact surface S , and let H_0 be a 0-handle of \mathcal{H} . An arc properly embedded in H_0 is *normal* if it satisfies the following conditions:

- (1) it does not have endpoints in the same component of $H_0 \cap \mathcal{H}^1$;
- (2) it does not have endpoints in the same component of $H_0 \cap \partial S$;
- (3) it does not have one endpoint in $H_0 \cap \mathcal{H}^1$ and the other endpoint in an adjacent component of $H_0 \cap \partial S$.

Definition 2.11. Let \mathcal{H} be a handle structure for a compact surface S . A 1-manifold C properly embedded in S is *normal* if it is standard and, in addition, each component of $C \cap \mathcal{H}^0$ is normal.

Definition 2.12. Let S be a surface with a handle structure \mathcal{H} . Let C be a standard 1-manifold properly embedded in S . Then a *simplifying disc* for C is a disc D lying in a 0-handle H_0 such that ∂D is the union of two arcs α and β that intersect at their endpoints, satisfying the following:

- (1) $\alpha = D \cap C$;
- (2) $\beta = D \cap \partial H_0$;
- (3) β intersects \mathcal{H}^1 in at most one arc;
- (4) β intersects ∂S in at most one arc.

If D intersects ∂S , it is a *boundary-simplifying disc*. Otherwise it is an *interior-simplifying disc*.

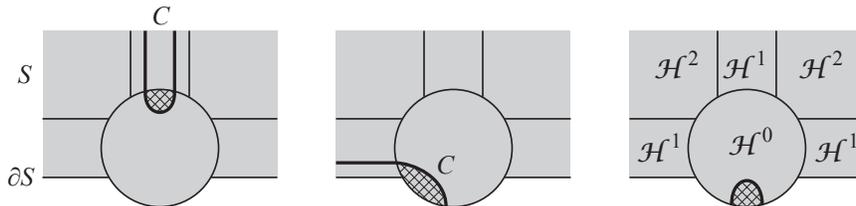


FIGURE 5. The three types of a simplifying disc for a 1-manifold in a surface. On the left is an interior-simplifying disc. In the middle and on the right are boundary-simplifying discs

Lemma 2.13. *Let \mathcal{H} be a handle structure for a compact surface S , and let C be a standard 1-manifold properly embedded in S . Then C is normal if and only if it has no simplifying disc.*

Proof. If C has a simplifying disc D , then $D \cap C$ is a non-normal arc. Conversely, if C has a non-normal arc α in some 0-handle H_0 , then some component D of $H_0 \setminus \alpha$ is a disc D that intersects \mathcal{H}^1 in at most one arc and ∂S in at most one arc. This disc may contain arcs of intersection with C , but then an outermost such arc separates off a simplifying disc for C . \square

Definition 2.14. The *simplification number* $\text{simp}(C)$ of C is the sum of the number of boundary-simplifying discs and twice the number of interior-simplifying discs.

In the proof of Lemma 5.4, we explain why boundary-simplifying discs and interior-simplifying discs are given different weightings in the above definition.

The following result is an analogue of Theorem 1.3 for handle structures, and is one of the main theorems of this paper.

Theorem 2.15 (NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE). *There is an algorithm that takes as its input the following:*

- (1) a handle structure \mathcal{H} for a compact surface S ;
- (2) a standard 1-manifold C properly embedded in S ;

and outputs a normal 1-manifold \bar{C} that is obtained from C by an isotopy plus possibly discarding some curves that bound discs and boundary-parallel arcs. The running time for the algorithm is bounded above by a polynomial function of $\|\mathcal{H}\|$, $\text{simp}(C)$ and $\log(w(C) + |\partial C|)$.

2.3. The size of a pairing system. Pairing systems are a very natural way of specifying standard 1-manifolds in surfaces. In this section, we investigate the amount of data required to specify a pairing system.

We will be given a handle structure \mathcal{H} for a surface S . We will also be given a standard 1-manifold C , as a pairing system. Recall from above that this means that each handle is viewed as a polygon, and then a pairing system is given in each polygon. We therefore introduce the following definition.

Definition 2.16. Let P be a polygon, and let C be a union of disjoint properly embedded arcs in P , disjoint from the vertices of P . Then a *simplifying disc* for C is a disc D embedded in P satisfying the following:

- (1) ∂D is the union of two arcs α and β that intersect at their endpoints;
- (2) $\alpha = D \cap C$;
- (3) $\beta = D \cap \partial P$ is a subset of an edge of P .

The simplification number $\text{simp}(C, P)$ is the number of simplifying discs.

When we view a 0-handle H_0 of \mathcal{H} as a polygon, then any simplifying disc for $C \cap H_0$ in the above sense is a simplifying disc for C in the sense of Definition 2.14. Hence, $\text{simp}(C) \geq \text{simp}(C, \mathcal{H}^0)$, where the latter is the sum of $\text{simp}(C, H_0)$ over each 0-handle H_0 .

The following lemma bounds the number of parallelism classes of arcs in a polygon in terms of their simplification number.

Lemma 2.17. *Let P be a polygon with k sides. Let C be a non-empty union of disjoint properly embedded arcs. Then the number of parallelism classes among the components of C is at most $2(k + \text{simp}(C, P)) - 3$.*

Proof. We start by subdividing the edges of the polygon P , by placing an extra vertex in each simplifying disc for C . The resulting polygon P' has $v = k + \text{simp}(C, P)$ vertices. Let \bar{C} be the arcs obtained by taking one representative of each arc type of C . Note that \bar{C} has no simplifying discs in P' . Then P' is cut into discs by \bar{C} . For each such disc, define its *complexity* to be the number of arcs of \bar{C} in its boundary, plus twice the number of vertices of P' that it contains, minus 2. Then the sum of the complexities of these discs is equal to $2|\bar{C}| + 2v - 2|P' \setminus \bar{C}| = 2v - 2$. A disc with negative complexity would have a single arc of \bar{C} in its boundary and no vertices of P' and hence would be a simplifying disc for \bar{C} , which does not exist. A disc with zero complexity must have two arcs of \bar{C} in its boundary and no vertices, and therefore these two arcs are parallel. However, by construction, no two arcs of \bar{C} are parallel. Hence, the number of components of $P' \setminus \bar{C}$ is at most $2v - 2$. So, the number of components of \bar{C} is at most $2v - 3$. \square

The following lemmas are immediate consequences.

Lemma 2.18. *Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Then the number of pairings in its associated pairing system is bounded above by $9|\mathcal{H}^1| + 2\text{simp}(C) \leq 9\|\mathcal{H}\| + 2\text{simp}(C)$.*

Proof. Note that C intersects only the 0-handles and 1-handles of \mathcal{H} . In each 1-handle, there is at most one parallelism class of arcs in C . So we focus on the 0-handles. Each 0-handle H_0 can be viewed as a polygon with $2|\mathcal{H}^1 \cap H_0|$ sides. Each 1-handle has two components of intersection with the 0-handles. Hence, the sum, over all 0-handles, of the number of sides is $4|\mathcal{H}^1|$. So, by Lemma 2.17, the number of arc types of C in the 0-handles is at most $8|\mathcal{H}^1| + 2\text{simp}(C, \mathcal{H}^0)$. As observed above, $\text{simp}(C, \mathcal{H}^0) \leq \text{simp}(C)$. The lemma follows immediately. \square

Definition 2.19. Let \mathcal{H} be a handle structure for a compact surface S . Let C_0 and C_1 be disjoint connected standard 1-manifolds properly embedded in S . Then they are *normally parallel* if there is an embedding of $C \times [0, 1]$ in S for some 1-manifold C , such that $C \times \{0\} = C_0$, $C \times \{1\} = C_1$, and for each $t \in [0, 1]$, $C \times \{t\}$ is standard and properly embedded.

Lemma 2.20. *Let \mathcal{H} be a handle structure of a compact surface S . Let C be a standard 1-manifold properly embedded in S . Suppose that no two components of C are normally parallel. Then $|C|$ is bounded above by $10\|\mathcal{H}\| + 2\text{simp}(C)$.*

Proof. It is a general fact that $|S \setminus C| \geq |C| - b_1(S; \mathbb{Z}_2)$. The reason is that if we cut S along C one component at a time, then we either increase the number of complementary regions or we find an element of $H^1(S; \mathbb{Z}_2)$ that is linearly independent from previously found elements. Furthermore, $b_1(S; \mathbb{Z}_2)$ is bounded above by $|\mathcal{H}^1| \leq \|\mathcal{H}\|$. As argued in the proof of Lemma 2.18, the number of arc types of $C \cap \mathcal{H}^0$ is bounded above $8|\mathcal{H}^1| + 2\text{simp}(C)$. Hence, the number of 0-handles of $S \setminus C$ that do not lie between parallel normal arcs of C is at most $|\mathcal{H}^0| + 8|\mathcal{H}^1| + 2\text{simp}(C)$. Consider any component of $S \setminus C$ that does not contain at least one such 0-handle. It cannot be an annulus or disc, since it would then co-bound normally parallel components of C . Hence, it must be a Möbius band, which therefore contributes to $b_1(S; \mathbb{Z}_2)$. Therefore, $|S \setminus C|$ is at most $8\|\mathcal{H}\| + 2\text{simp}(C) + b_1(S; \mathbb{Z}_2)$. We deduce that $|C|$ is bounded above by $10\|\mathcal{H}\| + 2\text{simp}(C)$. \square

We now consider how much information is required to specify a pairing system giving a 1-manifold C . The method for giving a pairing system is explained in Section 1. Each 1-cell must be oriented. The number of 1-cells is $6|\mathcal{H}^1| \leq 6\|\mathcal{H}\|$. The number of intersection points between each 1-cell and C must be given. The total number of intersection points is $2w(C) + |\partial C|$, since each component of $C \cap \mathcal{H}^1$ contributes 2 to this sum, and each component of ∂C contributes 1. Hence, the maximal number of intersection points with any 1-cell lies between $(2w(C) + |\partial C|)/6\|\mathcal{H}\|$ and $2w(C) + |\partial C|$. So, to give these numbers of intersection points requires of the order of $\|\mathcal{H}\| \log(w(C) + |\partial C|)$ bits of data. Finally, there is a pairing for each arc type, and the number of arc types lies between $\max\{1, \text{simp}(C)\}$ and $9\|\mathcal{H}\| + 2\text{simp}(C)$, by Lemma 2.18. Hence, the total amount of data is bounded above and below by linear functions of $(\|\mathcal{H}\| + \text{simp}(C)) \log(w(C) + |\partial C|)$.

2.4. A vector representation of normal curves and surfaces. In this subsection, we present an alternative way of specifying a standard 1-manifold C in a surface S with a handle structure \mathcal{H} . Let \mathcal{X} be the cell structure for S that is associated with \mathcal{H} . We subdivide \mathcal{X} by adding vertices to the interior of the 1-cells, one in the boundary of each simplifying disc, disjoint from C . After this has been done, each arc of intersection between C and each 2-cell joins distinct edges of the 2-cell. Thus, we may encode C as follows. For each 2-cell and for each pair of distinct edges of that 2-cell, we simply count the number of arcs of C running between these two edges. This list of non-negative integers is the *vector* (C) . Thus, we can specify C simply by specifying the subdivision of the 1-cells of the cell structure of \mathcal{H} and by giving the vector (C) .

Note that the amount of information required to specify a standard 1-manifold C in this way is at most a polynomial function of $\text{simp}(C)$, $\|\mathcal{H}\|$, and $\log(w(C) + |\partial C|)$, for the following reason. Before the subdivision of the cell structure, the total number of 1-cells was $6|\mathcal{H}^1|$ by Lemma 2.5. The number of subdivisions that we perform is at most $\text{simp}(C)$, and each such subdivision increases the number of 1-cells by 1. So afterwards, this number is at most $6|\mathcal{H}^1| + \text{simp}(C)$. Thus, the number of 1-cells in the boundary of the 2-cells is at most $12|\mathcal{H}^1| + 2\text{simp}(C)$. So the number of arc types in the 2-cells is at most $(12|\mathcal{H}^1| + 2\text{simp}(C))^2$. For each arc type, we specify the number of arcs of this type, and the total number of arcs

is at most $2w(C) + |\partial C|/2$, because each arc of C alternates between 0-handles and 1-handles of \mathcal{H} , and the number of times it runs over the 1-handles is its weight. Note that the addition of $|\partial C|/2$ is required to deal with arc components of C , each of which has one more component of intersection with the 0-handles than with the 1-handles. Hence, the number of digits required to specify each co-ordinate of (C) is at most $\log(2w(C) + |\partial C|/2)$.

3. THE ALGORITHM OF AGOL-HASS-THURSTON

A crucial tool that is used in this paper is the algorithm of Agol, Hass and Thurston [1]. We describe it in this section. The original use of the algorithm was to determine the number of components of a normal 1-manifold in a triangulated surface, and, in one higher dimension, the number of components of a normal surface in a triangulated 3-manifold. However, it is a very general algorithm with applications in many different situations in low-dimensional topology. (See, for example, [10] where it is used in several ways.)

A *pairing* is an order-preserving or order-reversing bijection $[a, b] \cap \mathbb{Z} \rightarrow [c, d] \cap \mathbb{Z}$ for integers $a \leq b$ and $c \leq d$. For an integer x and pairing g , we only write $g(x)$ if x is in the domain of g . Similarly, we only write $g^{-1}(x)$ if x is in the domain of g^{-1} . The input to the algorithm is:

- (1) non-negative integers N and r ;
- (2) a collection of r pairings $g_i: [a_i, b_i] \cap \mathbb{Z} \rightarrow [c_i, d_i] \cap \mathbb{Z}$, where $1 \leq a_i \leq b_i \leq N$ and $1 \leq c_i \leq d_i \leq N$.

Two integers x and y are said to be *in the same orbit* of these pairings if there is a sequence i_1, \dots, i_k of integers between 1 and r and integers $\varepsilon_1, \dots, \varepsilon_k \in \{-1, 1\}$ such that $g_{i_k}^{\varepsilon_k} \dots g_{i_1}^{\varepsilon_1}(x) = y$. This forms an equivalent relation, and the equivalence classes are called *orbits*.

The following is the basic algorithm that Agol, Hass and Thurston provided. They also gave an enhanced version that we will describe later.

Theorem 3.1 (BASIC AHT). *There is an algorithm that takes, as its input, non-negative integers N and r and a collection of pairings as above, and provides as its output the number of orbits. The running time is bounded above by a polynomial function of r and $\log N$.*

As an immediate consequence of this and Lemma 2.18, we obtain the following, which is a minor extension of a result of Agol, Hass and Thurston [1, Corollary 13] from normal 1-manifolds to standard 1-manifolds. The crucial observation is that a standard 1-manifold comes equipped with a pairing system, and orbits of this system correspond to components of the 1-manifold.

Theorem 3.2 (NUMBER OF COMPONENTS OF STANDARD 1-MANIFOLD). *There is an algorithm that takes, as its input, a triangulation \mathcal{T} or handle structure \mathcal{H} for a compact surface S and a standard 1-manifold C , and provides as its output the number of components of C . The running time is bounded above by a polynomial function of $\text{simp}(C)$, $\log(w(C) + |\partial C|)$ and $|\mathcal{T}|$ or $\|\mathcal{H}\|$.*

Proof. The standard 1-manifold C is specified using a pairing system. For each edge e of \mathcal{T} or each 1-cell e of the cell structure for \mathcal{H} , the quantity $|e \cap C|$ is given as part of the pairing system. Let N equal $\sum_e |e \cap C|$. In the case when we are given a triangulation, N is just the weight of C . In the case of a handle structure,

N is $|\partial C| + 2w(C)$. If we pick a total ordering on the edges, then, using the given orientation on the edges, we can identify $[1, N] \cap \mathbb{Z}$ with the points of intersection between C and the edges e . Thus we obtain pairings between sub-intervals of $[1, N] \cap \mathbb{Z}$. The number r of these pairings is at most $9\|\mathcal{H}\| + 2\text{simp}(C)$ in the case where we have a handle structure \mathcal{H} and is similarly bounded when we have a triangulation \mathcal{T} . Then BASIC AHT counts the number of orbits, which is the number of components of C . Its running time is bounded above by a polynomial function of $\text{simp}(C)$, $\log(w(C) + |\partial C|)$ and $|\mathcal{T}|$ or $\|\mathcal{H}\|$. \square

There is an enhanced version of the Agol-Hass-Thurston algorithm, which is as follows. Its input is the non-negative integers N and r , the pairings g_i and also a function $\phi: [1, N] \cap \mathbb{Z} \rightarrow \mathbb{Z}^m$ for some positive integer m . The ϕ -value of an orbit is then the sum of $\phi(x)$, over all elements x in the orbit. One measure of complexity of ϕ is $M = \max\{|\phi(x)|_1 + 1 : x \in [1, N] \cap \mathbb{Z}\}$, where $|\cdot|_1$ is the ℓ^1 norm. Another measure is the number s of integers x in $[2, N]$ for which $\phi(x) \neq \phi(x-1)$.

Theorem 3.3 (AHT). *There is an algorithm that takes as its input non-negative integers N and r , a collection of r pairings as above, and a function $\phi: [1, N] \cap \mathbb{Z} \rightarrow \mathbb{Z}^m$ as above, and provides as its output a list of orbits together with their ϕ -values. The running time is bounded above by a polynomial function of r , s , m , $\log N$ and $\log M$, where s and M are as defined above.*

This has many interesting applications. We present one of them now, which strengthens Theorem 3.2.

Theorem 3.4 (COMPONENTS OF STANDARD 1-MANIFOLD). *There is an algorithm that takes, as its input, a triangulation \mathcal{T} or handle structure \mathcal{H} for a compact surface S and a standard 1-manifold C , and provides as its output a list of the components of C , given as vectors. The running time is bounded above by a polynomial function of $\text{simp}(C)$, $\log(w(C) + |\partial C|)$ and $|\mathcal{T}|$ or $\|\mathcal{H}\|$.*

Proof. The 1-manifold C can be given as a pairing system or as a vector (C) as described in Section 2.4. If we are given it as a pairing system, we can convert it to a vector (C) . Similarly, if we are given (C) , we can convert it to a pairing system. We will use both such representations of C .

Say that there are m arc types of C in the 2-cells. For each arc type, we will define a function $\phi_i: [1, N] \cap \mathbb{Z} \rightarrow \mathbb{Z}$. These will combine to form the function $\phi: [1, N] \cap \mathbb{Z} \rightarrow \mathbb{Z}^m$. In other words, the composition of ϕ with projection onto the i th co-ordinate of \mathbb{Z}^m is ϕ_i . For the given arc type, this determines a pairing $[a, b] \cap \mathbb{Z} \rightarrow [c, d] \cap \mathbb{Z}$. Define ϕ_i to be 1 on the domain of this pairing and 0 elsewhere. Then if C' is a component of C , then C' corresponds to a subset of $[1, N]$ and if we apply ϕ to this subset, then the resulting integer is the co-ordinate of (C') corresponding to this arc type. Note that number s of values of x for which $\phi(x) \neq \phi(x-1)$ is at most twice the number of arc types.

Hence, when we apply AHT to this pairing system with this function ϕ , then the output is a list of the components of C , and the ϕ -value of each component C' gives its vector (C') . \square

4. CUTTING ALONG A STANDARD 1-MANIFOLD

One of the advantages of using handle structures rather than triangulations is that when we cut along a standard 1-manifold, the resulting manifold inherits a handle structure, as follows.

Definition 4.1. Let S be a compact surface with a handle structure \mathcal{H} . Let C be a standard 1-manifold properly embedded in S . Then the *inherited handle structure* on $S \setminus C$ has an i -handle for each component of $\mathcal{H}^i \setminus C$. In the case where S is instead given as a triangulation \mathcal{T} , we let \mathcal{H} be the handle structure dual to \mathcal{T} and then the inherited handle structure on $S \setminus C$ is defined as above.

The inherited handle structure on $S \setminus C$ is frequently not practical to use because it can have a large number of handles. Specifically, its number of 1-handles is at least $w(C)$. However, when $w(C)$ is large, many of the handles of $S \setminus C$ have the following form.

Definition 4.2. A *parallelity handle* for $S \setminus C$ is a handle in the inherited handle structure that lies between two adjacent normally parallel arcs of C in a handle H of \mathcal{H} . Thus, its boundary is the concatenation of four arcs. Two of these are components of $C \cap H$, and the other two are arcs in $\partial S \cap H$ or $H \cap \mathcal{H}^1 \cap \mathcal{H}^0$. The parallelity handle is an I -bundle over an interval, with the (∂I) -bundle being its intersection with C .

Definition 4.3. The *parallelity bundle* for $S \setminus C$ is the union of its parallelity handles. It is an I -bundle over a 1-manifold, with the (∂I) -bundle being its intersection with C .

In general, for an I -bundle B , its *horizontal boundary* $\partial_h B$ is the (∂I) -bundle. Its *vertical boundary* $\partial_v B$ is $\partial B \setminus \partial_h B$.

Note that two distinct components C_1 and C_2 of C are normally parallel precisely when some component of $S \setminus C$ is a component B of the parallelity bundle with $\partial_h B = C_1 \cup C_2$.

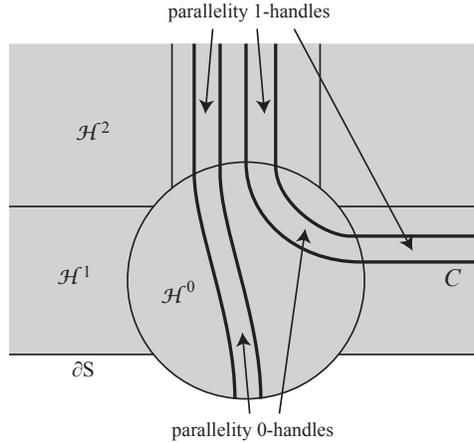


FIGURE 6. Parallelity handles of $S \setminus C$

Definition 4.4. Let S be a compact surface with a handle structure \mathcal{H} . Let C be a standard 1-manifold properly embedded in S . Then the *reduced handle structure* \mathcal{H}' is defined as follows. We start with the handle structure on $S \setminus \setminus C$ that it inherits. Let \mathcal{B} be its parallelity bundle. Then \mathcal{H}' is obtained by removing all the handles in \mathcal{B} and replacing each component of \mathcal{B} that misses ∂S and is an I -bundle over an interval by a single 1-handle. Each component of $S \setminus \setminus C$ that is a union of parallelity handles is removed.

Remark 4.5. As explained in the definition, components of $S \setminus \setminus C$ that are a union of parallelity handles are removed. Thus, the reduced handle structure might be not a handle structure for all of $S \setminus \setminus C$, but only certain components of $S \setminus \setminus C$. Moreover each component of \mathcal{B} that intersects ∂S in a single arc is removed, but this does not change the homeomorphism type of the surface.

Note that by Lemma 2.18, the number of handles of $S \setminus \setminus C$ that are not parallelity handles is at most $O(\|\mathcal{H}\| + \text{simp}(C))$. The remaining handles of the reduced handle structure are components B of the \mathcal{B} that are disjoint from ∂S and that are an I -bundle over an interval. For each such component B , $\partial_v B$ is incident to arcs of $C \cap \mathcal{H}^0$ that are outermost in their parallelism classes. Hence the number of such components of \mathcal{B} is also at most $O(\|\mathcal{H}\| + \text{simp}(C))$. Therefore, the number of handles in the reduced handle structure is at most $O(\|\mathcal{H}\| + \text{simp}(C))$.

Theorem 4.6 (CUT ALONG STANDARD 1-MANIFOLD). *There is an algorithm that takes, as its input, a handle structure \mathcal{H} for a compact surface S and a standard 1-manifold C properly embedded in S , and outputs the reduced handle structure on the relevant components of $S \setminus \setminus C$. The running time is at most a polynomial function of $\|\mathcal{H}\|$, $\text{simp}(C)$, and $\log(w(C) + |\partial C|)$.*

Proof. A closely related algorithm is given in [10, Theorems 9.2 and 9.3] and so we just sketch it here.

Let \mathcal{B} be the parallelity bundle of $S \setminus \setminus C$. The first stage in the procedure is to determine the handle structure on $(S \setminus \setminus C) \setminus \setminus \mathcal{B}$, as follows. For each handle H of \mathcal{H} , it is divided into handles $H \setminus \setminus C$. Those lying in \mathcal{B} are the ones lying between parallel arcs of $H \cap C$. Thus, we can form $(H \setminus \setminus C) \setminus \setminus \mathcal{B}$ by replacing each arc type of $H \cap C$ by a single arc, and then cutting along the resulting arcs. In this way, we can form the handle structure on $(S \setminus \setminus C) \setminus \setminus \mathcal{B}$, and also identify its components of intersection with $\partial_v \mathcal{B}$.

Say that there are k arc types of C in the 0-handles and 1-handles of \mathcal{H} . For $1 \leq i \leq k$, say there are N_i arcs of that type. Then there are $\max\{0, N_i - 1\}$ parallelity handles of $S \setminus \setminus C$ between these arcs. Let $N = 2 \sum \max\{0, N_i - 1\}$, which is the total number of components of $\partial_h H'$, as H' runs over all parallelity handles. This will be one of the inputs into the AHT algorithm.

Let m be the number of boundary components of $\partial_v \mathcal{B}$ disjoint from ∂S . This will be another input to AHT. Note that these components of $\partial(\partial_v \mathcal{B})$ lie in $(S \setminus \setminus C) \setminus \setminus \mathcal{B}$, and hence m is bounded above by a linear function of $\|\mathcal{H}\|$ and $\text{simp}(C)$. We will define a function $\phi: [1, N] \rightarrow \mathbb{Z}^m$. For each integer between 1 and m , we identify it with a specific boundary component of $\partial_v \mathcal{B} \setminus \setminus \partial S$. Consider a 1-handle H_1 of \mathcal{H} attached onto a 0-handle H_0 . Suppose that $H_1 \cap C$ consists of more than 1 arc. Then between these arcs lie handles of \mathcal{B} . Each such handle H'_1 has two components of $\partial_h H'_1$, each of which is identified with an integer i between 1 and N . The handle H'_1 is incident to a handle of $H_0 \setminus \setminus C$ that may or may not be parallelity. If it is not

a parallelity handle, then each component of $\partial_h H'_1$ meets a component of $\partial(\partial_v \mathcal{B})$, the j th component say. Then we define $\phi(i)$ to be 1 at the j co-ordinate and 0 at the remaining co-ordinates. For the remaining handles H'_1 that are incident to parallelity 0-handles in H_0 , this incidence defines pairings, each going from a subset of $[1, N]$ to another subset of $[1, N]$.

When we run AHT, the output is a collection of orbits. Each orbit corresponds to a component of $\partial_h \mathcal{B}$. For each such component α , we examine its ϕ -value to determine the location of $\partial\alpha \setminus \partial S$. The number of components of $\partial\alpha \setminus \partial S$ is 0, 1 or 2. If it is 2, then the component of \mathcal{B} containing α is an I -bundle over an interval and is disjoint from ∂S , and hence forms a 1-handle in the reduced handle structure. On the other hand, if this number is less than 2, then the component of \mathcal{B} containing α is removed when forming the reduced handle structure.

We can therefore compute which components of \mathcal{B} form 1-handles of the reduced handle structure and how they are attached onto $(S \setminus C) \setminus \mathcal{B}$. \square

Remark 4.7. It is worth stating that the reduced handle structure \mathcal{H}' on $S \setminus C$ is provided not just as an abstract handle structure but also with information about how it embeds in \mathcal{H} . More precisely, for each handle of \mathcal{H}' that did not come from a component of the parallelity bundle \mathcal{B} , we have an identification between it and the handle of $S \setminus C$ that it came from. We can also determine the location of the handles arising from components of \mathcal{B} , in the following sense.

Theorem 4.8 (LOCATION OF PARALLELITY BUNDLE COMPONENT). *There is an algorithm that takes, as its input, the following:*

- (1) a handle structure \mathcal{H} for a compact surface S ;
- (2) a standard 1-manifold C properly embedded in S ;
- (3) a point p in the intersection between $S - C$ and a 1-cell of the associated cell structure on S ;

and provides the following:

- (1) the vector(s) in S for the component(s) of $\partial_h B$, where B is the component of the parallelity bundle for $S \setminus C$ containing p (or the zero vector if p does not lie in the parallelity bundle);
- (2) for each such component α of $\partial_h B$, the location of $\partial\alpha - \partial S$.

The running time is at most a polynomial function of $|\mathcal{H}|$, $\text{simp}(C)$, and $\log(w(C) + |\partial C|)$.

Proof. In the proof of Theorem 4.6, N was the number of horizontal boundary components of the parallelity handles for $S \setminus C$. Each such handle has an associated arc type of C . Say that there are L arc types. Then instead of defining ϕ to have image in \mathbb{Z}^m , we define it to have image in \mathbb{Z}^{m+L+1} . The first m co-ordinates are as in the proof of Theorem 4.6. The next L co-ordinates correspond to the L arc types and record the arc type of a parallelity handle. The final co-ordinate is defined to be 1 if the handle contains p , and zero otherwise. The output of AHT is a list of orbits, each of which corresponds to a component of $\partial_h \mathcal{B}$. For each such component α of $\partial_h \mathcal{B}$, we can determine whether it lies in the same component of \mathcal{B} as p by reading off the final component of the ϕ -value of α . We can read off the vector of α by examining the previous L co-ordinates. \square

It will be useful to consider the following close relative of the parallelity bundle.

Definition 4.9. A *semi-parallelity handle* for $S \setminus C$ is a handle H of its inherited handle structure such that ∂H is a concatenation of an arc in C , an arc in $\mathcal{H}^0 \cap \mathcal{H}^1$, an arc in ∂S and an arc in $\mathcal{H}^0 \cap \mathcal{H}^1$. It is an I -bundle, where the (∂I) -bundle is its intersection with $C \cup \partial S$. The *semi-parallelity bundle* is the union of the semi-parallelity handles. It too is an I -bundle, where the (∂I) -bundle is its intersection with $C \cup \partial S$.

Note that it is straightforward to compute the location of the semi-parallelity bundle, because the number of semi-parallelity handles is at most twice the number of components of $\partial S \cap \mathcal{H}^1$, and this is at most $2|\mathcal{H}^1| \leq 2\|\mathcal{H}\|$.

5. FAST NORMALISATION OF STANDARD 1-MANIFOLDS

In this section, we prove Theorem 2.15. The proof of Theorem 1.3 is entirely analogous and is omitted.

Theorem 2.15 (NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE). *There is an algorithm that takes as its input the following:*

- (1) a handle structure \mathcal{H} for a compact surface S ;
- (2) a standard 1-manifold C properly embedded in S ;

and outputs a normal 1-manifold \bar{C} that is obtained from C by an isotopy plus possibly discarding some curves that bound discs and boundary-parallel arcs. The running time for the algorithm is bounded above by a polynomial function of $\|\mathcal{H}\|$, $\text{simp}(C)$ and $\log(w(C) + |\partial C|)$.

We are given a handle structure \mathcal{H} for a compact surface S and a standard 1-manifold C properly embedded in S .

We first note that we may assume that no two components of C are normally parallel, via the following procedure. We may compute a list of components of C by COMPONENTS OF STANDARD 1-MANIFOLD (Theorem 3.4). We may then replace C by C_- , which has a single copy of each normal parallelism type of the components of C . Suppose that we can isotope C_- to a normal 1-manifold \bar{C}_- . Then we can place C into normal form by replacing each component of \bar{C}_- by a suitable number of copies. So, we will now assume that no two components of C are normally parallel. This is useful because in this situation, the reduced handle structure is a handle structure for a surface homeomorphic to $S \setminus C$, except possibly with some Möbius band components removed. (See Remark 4.5.)

Suppose that C is not normal. Then it has a simplifying disc D , by Lemma 2.13. If D is disjoint from \mathcal{H}^1 , then $D \cap C$ is a boundary-parallel arc, and we can remove this arc and all normally parallel copies of it. Suppose therefore that $D \cap C$ intersects \mathcal{H}^1 . By definition of a simplifying disc, this intersection is an arc. We may then isotope C across D and across the incident 1-handle. The resulting 1-manifold C' has smaller weight than C . It might not in fact be standard, because it might have a component of intersection with a 0-handle that is a simple closed curve. But in that case, we can remove this component. So, we may assume that C' is standard. Hence, by induction on the weight of C , we can normalise it. This is the usual normalisation process.

However, this does not obviously lead to an efficient algorithm, for two reasons. Firstly, the weight of C' is only $w(C) - 1$ or $w(C) - 2$. Hence, there may need to be of the order of $w(C)$ such isotopies before C is normalised, whereas we require

the algorithm to terminate in polynomial time as a function of $\log(w(C) + |\partial C|)$. Secondly, we have seen in Section 2.3 that the amount of information required to describe a standard curve C is of the order $(\text{simp}(C) + \|\mathcal{H}\|) \log(w(C) + |\partial C|)$. Thus, to prevent this quantity increasing without bound, we would need to control $\text{simp}(C)$.

Instead of using $w(C)$ as our measure of progress through the algorithm that we will present, we use $\text{simp}(C)$. However $\text{simp}(C')$ might actually be greater than $\text{simp}(C)$. An example is shown in Figure 7.

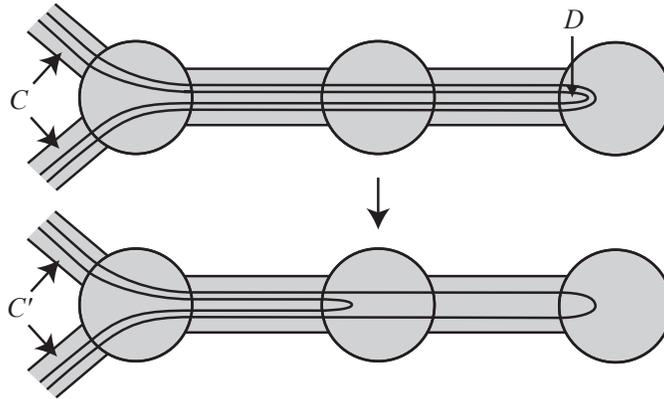


FIGURE 7. Shown is a standard curve C with a simplifying disc D . When an isotopy is performed across D and the incident 1-handle, the effect is to increase the simplification number

In general, this happens as follows. For simplicity, we initially assume that S is closed. Let H_0 be the 0-handle of \mathcal{H} containing D . Running parallel to $D \cap C$, there might be another arc of $H_0 \cap C$. This would be part of the boundary of a new simplifying disc for C' . But also the new arc of $\mathcal{H}^0 \cap C'$ created by the isotopy across D and the incident 1-handle may also bound a new simplifying disc D' . There are two possibilities for D' : either the interior of D' intersects the disc region between C and C' or it is disjoint from it. These two cases are shown in Figure 8.

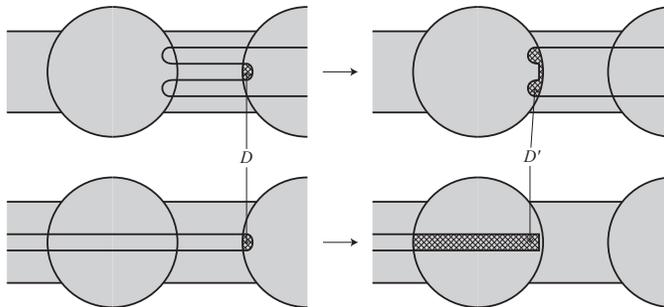


FIGURE 8. Shown on the left are simplifying discs D for the standard curve C . When the associated isotopy of C is performed, new simplifying discs D' are created.

In the former case, then D' was actually formed from two simplifying discs of C , both of which are removed in the isotopy. Thus, in this case, it is possible to deduce that $\text{simp}(C') < \text{simp}(C)$. So suppose that the interior of D' is disjoint from the region between C and C' . Then D' lay in a parallelity handle of $S \setminus C$. Thus, we are led to the following definition.

Definition 5.1. Let \mathcal{H} be a handle structure for a closed surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C . Then the *isotopy region* associated with D is the union of the following subsets of $S \setminus C$:

- (1) the simplifying disc D ;
- (2) the component B of the parallelity bundle for $S \setminus C$ that is incident to D ;
- (3) a small regular neighbourhood of $\partial_v B$ in $S \setminus C$.

(See Figure 9.) We say that the *horizontal boundary* of R is $\partial R \cap C \setminus \partial D$.

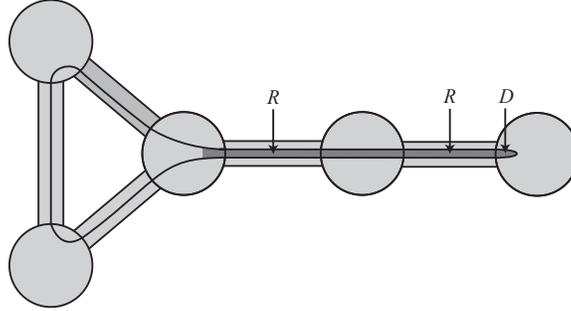


FIGURE 9. The isotopy region R associated with the simplifying disc D is shaded

The version of the above definition when S has boundary is a bit more complicated. It involves the following type of 0-handle of $S \setminus C$.

Definition 5.2. Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . A 0-handle of $S \setminus C$ is an *interpolating 0-handle* if its boundary is a concatenation of an arc of $C \cap \mathcal{H}^0$, an arc in $\mathcal{H}^0 \cap \mathcal{H}^1$, an another arc of $C \cap \mathcal{H}^0$, another arc in $\mathcal{H}^0 \cap \mathcal{H}^1$ and an arc in ∂S . See Figure 10.

Definition 5.3. Let \mathcal{H} be a handle structure for a compact surface S with non-empty boundary. Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C . Then the *isotopy region* R associated with D is the union of the following subsets of $S \setminus C$:

- (1) the simplifying disc D ;
- (2) the component B_1 of the parallelity or semi-parallelity bundle that is incident to D ;
- (3) if D is an interior-simplifying disc and B_1 is incident to an interpolating 0-handle of $S \setminus C$, then this 0-handle is also part of R and so is the incident component B_2 of the semi-parallelity bundle;
- (4) a small regular neighbourhood of $\partial_v B_1$ in $S \setminus C$, as well as a small regular neighbourhood of $\partial_v B_2$ if B_2 is defined.

(See Figure 10.) We say that the *horizontal boundary* of R is $\partial R \cap (C \cup \partial S) \setminus \partial D$.

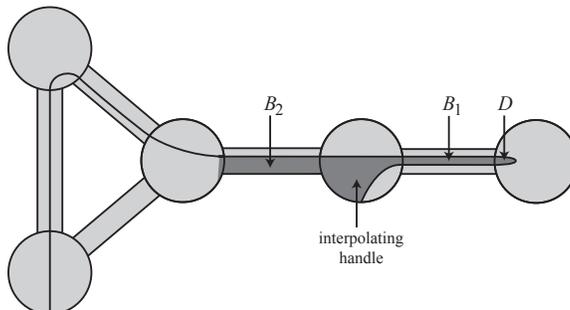


FIGURE 10. The isotopy region associated with a simplifying disc D in the case where $\partial S \neq \emptyset$.

The reason why an interpolating handle might be included in the isotopy region is as follows. Suppose that C consists of two parallel copies of the 1-manifold shown in Figure 10. If we were to isotope one of these copies across D and B_1 but no further, then the resulting 1-manifold C' would have higher simplification number than C . Although the interior-simplifying disc D would have been removed, an interior-simplifying disc and a boundary-simplifying disc would have been created.

It is also worth pointing out that here the situation for triangulations is somewhat simpler than for handle structures. In the case where S comes with a triangulation, then Definition 5.1 is used rather than Definition 5.3. In particular, there is no analogue of an interpolation handle.

The following key lemma states that the simplification number does not go up when an isotopy is performed across an isotopy region.

Lemma 5.4. *Let \mathcal{H} be a handle structure for a compact surface S , and let C be a standard properly embedded 1-manifold. Let C' be the 1-manifold obtained by isotoping C across the isotopy region associated with a simplifying disc D . Then $\text{simp}(C') \leq \text{simp}(C)$ and $w(C') < w(C)$. Furthermore, if $\text{simp}(C') = \text{simp}(C)$, then there is an arc of $\mathcal{H}^0 \cap C'$ that is normally parallel to $D \cap C$, and all the components of $C \cap \mathcal{H}^0$ between these two arcs are normally parallel to $D \cap C$.*

Proof. We define a function

$$f: \{\text{simplifying discs for } C'\} \rightarrow \mathbb{P}\{\text{simplifying discs for } C\}$$

as follows. For a simplifying disc D' for C' , we set $f(D')$ to be the set of simplifying discs for C that lie in D' . We will establish the following properties:

- (1) $f(D') \neq \emptyset$ for every simplifying disc D' for C' ;
- (2) $f(D'_1) \cap f(D'_2) = \emptyset$ for distinct simplifying discs D'_1 and D'_2 for C' ;
- (3) if D' is an interior-simplifying disc for C' , then $f(D')$ consists of interior-simplifying discs;
- (4) if $f(D') = \{D_1\}$ for a single disc D_1 that is interior-simplifying exactly when D' is interior-simplifying, then $C \cap \text{int}(D')$ consists of a (possibly empty) collection of parallel arcs, all normally parallel to $D' \cap C'$.

These properties quickly imply the lemma, as follows. Via (1), each simplifying disc for C' gives rise to a non-empty collection of simplifying discs for C . As a result of (2), these sets are pairwise disjoint. Hence, the number of simplifying discs for C is at least the number of simplifying discs for C' . Moreover, by (3), each interior-simplifying disc for C' gives rise to at least one interior-simplifying disc for C . So we deduce that $\text{simp}(C') \leq \text{simp}(C)$. In the case of equality, each interior-simplifying disc for C' contains exactly one simplifying disc for C , which is interior-simplifying; and each boundary-simplifying disc for C' contains exactly one simplifying disc for C , which is boundary-simplifying; and each simplifying disc for C lies in a simplifying disc for C' . (It is here that we use the fact that interior-simplifying discs and boundary-simplifying discs for C have different contributions to $\text{simp}(C)$.) In particular, the simplifying disc D must lie in a simplifying disc D' for C' . Since $|f(D')| = 1$ and the hypotheses of (4) are satisfied, then (4) implies that $C \cap \text{int}(D')$ consists of a (possibly empty) collection of parallel arcs, all normally parallel to $D' \cap C'$. In fact, these arcs are non-empty in this case, because otherwise $D' = D$, but D is removed by the isotopy. Thus, we deduce that, in the case where $\text{simp}(C') = \text{simp}(C)$, there is an arc of $\mathcal{H}^0 \cap C'$ that is normally parallel to $D \cap C$, and all the components of $C \cap \mathcal{H}^0$ between these two arcs are normally parallel to $D \cap C$, as required.

We now prove (1) - (4). Properties (2) and (3) are trivial. For (2), note that distinct simplifying discs D'_1 and D'_2 for C' are disjoint, and therefore the simplifying discs for C that lie within them are disjoint. For (3), note that an interior-simplifying disc D' is disjoint from ∂S and hence so is every disc that it contains.

For (4), observe that if the arcs $C \cap \text{int}(D')$ are non-empty and not all parallel in D' , then at least two separate off sub-discs of D' that are disjoint from $\partial D' \cap C'$. These discs are then simplifying discs, implying that $|f(D')| > 1$. So if $f(D') = \{D_1\}$, then the arcs $C \cap \text{int}(D')$ are all parallel in D' . Moreover when D_1 and D' are both interior-simplifying or both boundary-simplifying, then the arcs $C \cap \text{int}(D')$ are all normally parallel to $D' \cap C'$, as required.

For a simplifying disc D' for C' that is also a simplifying disc for C , property (1) is immediate. Therefore, what concerns us are simplifying discs for C' that are not simplifying discs for C . Let D' be one such disc. Let α' be the arc $\partial D' \cap C'$. We consider two possibilities: either α' is a subset of C or it is not.

Suppose first that α' is a subset of C . Then the reason that D' is not a simplifying disc for C must have been because C intersected the interior of D' and the arcs of intersection between C and the interior of D' were removed in the isotopy. In that case, at least one of these arcs bounded a simplifying disc for C that was removed in the isotopy. This establishes (1) in this case.

Now consider the case where α' is not a subset of C . It therefore contains the arc $C' - C$. Thus, the isotopy region R between C and C' has non-empty intersection with α' . Let R_0 be the component of $R \cap \mathcal{H}^0$ incident to α' . This either lies in D' or has interior disjoint from D' .

Suppose first that the interior of R_0 is disjoint from D' . (See the second case in Figure 8.) Then $R_0 \cup D'$ would have formed a parallelity, semi-parallelity or interpolating 0-handle of $S \setminus C$. Moreover, in the case where $R_0 \cup D'$ is an interpolating handle, R_0 would have been incident to the parallelity bundle for $S \setminus C$. Hence, when the interior of R_0 is disjoint from D' , $R_0 \cup D'$ would have been part of

the isotopy region, contradicting the assumption that D' is not part of the isotopy region.

Now suppose that R_0 lies in D' . (An example is shown in the first case in Figure 8.) In the case where D' is an interior-simplifying disc, R_0 cuts D' into at least two discs, and hence D' contains at least two interior-simplifying discs for C , both of which are removed by the isotopy. In the case where D' is a boundary-simplifying disc, $D' \setminus R_0$ has at least one component disjoint from ∂S , and therefore D' contains at least one interior-simplifying disc for C . Thus, in both cases, $f(D')$ is non-empty. \square

Let D be a simplifying disc for C . Let C' be the result of performing the isotopy across the isotopy region associated with D . If $\text{simp}(C') = \text{simp}(C)$, then one might hope to iterate this procedure, to remove the arc of $\mathcal{H}^0 \cap C'$ that runs parallel to $D \cap C$, and so on. In fact, this is a good strategy, but the situation is complicated by the following possibility. The arc of $C' \cap \mathcal{H}^0$ incident to the isotopy region might also be parallel to $D \cap C$. So, when we perform the isotopy across the isotopy region associated with D , we might not decrease the number of parallel copies of $D \cap C$. (See Figure 11.) We formalise this as follows.

Definition 5.5. Let \mathcal{H} be a handle structure for a compact surface S , and let C be a standard properly embedded 1-manifold. Let D be a simplifying disc for C and let R be the associated isotopy region. Let C' be the 1-manifold obtained by the isotopy across R . Then D and R are *self-returning* if the arc of $C' \cap \mathcal{H}^0$ incident to R is normally parallel to $D \cap C$ and all the components of $C \cap \mathcal{H}^0$ between these two arcs are normally parallel to $D \cap C$.

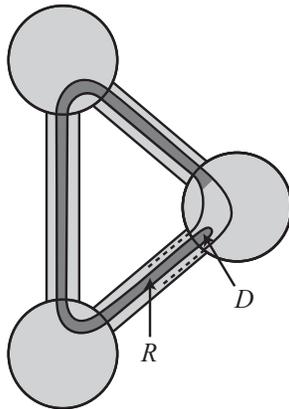


FIGURE 11. The simplifying disc D and its isotopy region R are self-returning

We will initially focus on the case where D is not self-returning.

We know that if $\text{simp}(C') = \text{simp}(C)$, then there is an arc of $\mathcal{H}^0 \cap C'$ that runs parallel to $D \cap C$. If D is not self-returning, then this is actually an arc of $\mathcal{H}^0 \cap C$. It is natural to attempt to perform the associated isotopy at the same time as the one across D . We are therefore led to the following definition.

Definition 5.6. Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C . Let R be the associated isotopy region. Then an *enlargement* of R consists of a disc D_+ containing R and where $C \cap D_+$ consists of arcs all normally parallel to $R \cap C$, one of which lies in ∂D_+ . An enlargement is *maximal* if it does not lie in a bigger enlargement, up to isotopy preserving C and the handles of \mathcal{H} . The *isotopy across D_+* replaces the arcs $D_+ \cap C$ with arcs parallel to $\partial R \setminus (C \cup \partial S)$.

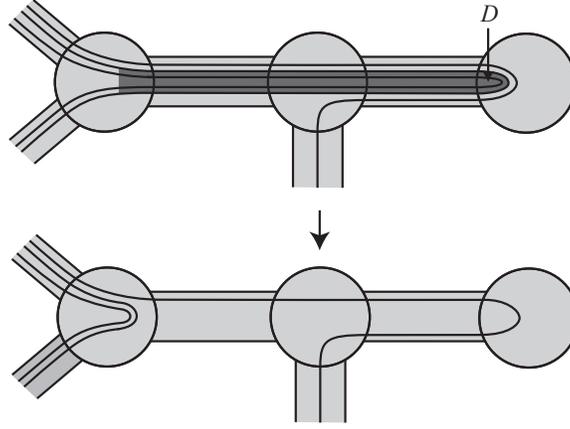


FIGURE 12. The maximal enlargement of the isotopy region is shaded. The associated isotopy is shown.

Clearly, the isotopy across D_+ can be expressed as a composition of isotopies across isotopy regions. Hence, the following result is immediate.

Lemma 5.7. *Let \mathcal{H} be a handle structure for a compact surface S , and let C be a properly embedded standard 1-manifold. Let C' be the 1-manifold that is obtained from C by an isotopy across a maximal enlargement of an isotopy region. Then $\text{simp}(C') \leq \text{simp}(C)$.*

As is evident from Figure 12, the isotopy across D_+ might not reduce the simplification number. However, the following lemma asserts that one need only perform such isotopies a controlled number of times and then the simplification number must strictly decrease.

Lemma 5.8. *Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C . Suppose that D is not self-returning. Then after performing at most $18\|\mathcal{H}\| + 4\text{simp}(C)$ isotopies across maximal enlargements of isotopy regions containing D , the resulting 1-manifold C'' satisfies $\text{simp}(C'') < \text{simp}(C)$.*

Proof. Let R be the isotopy region for D and let D_1 be the maximal enlargement of R . Suppose first that every arc of $C \cap \mathcal{H}^0$ parallel to $D \cap C$ is contained in D_1 . Then, when we perform the isotopy across D_1 , we remove all the normal arcs parallel to $D \cap C$. We are assuming that D is not self-returning, and so this isotopy does not introduce any normal arcs parallel to $D \cap C$. Hence, this isotopy has reduced the simplification number, by Lemma 5.4.

So, we may assume that there is a normal arc of $C \cap \mathcal{H}^0$ parallel to $D \cap C$ that is not contained in D_1 . Consider the outermost one in \mathcal{H}^0 , in other words, the one closest to D . Then we can extend this arc maximally, so that the resulting arc β runs parallel to the boundary of D_1 . But at some point this arc must diverge from D_1 . At this point, ∂D_1 contains an arc of $C \cap \mathcal{H}^0$ that is outermost in its parallelism class. We enlarge the disc D_1 to form a disc D_2 that includes β and all arcs normally parallel to it. Now consider the next arc of $C \cap \mathcal{H}^0$ that is parallel to $D \cap C$ but not contained in D_2 . It runs parallel to ∂D_2 until it diverges from it. Enlarge D_2 to form a disc D_3 that includes this arc and all arcs that are normally parallel to it. Continue in this way. The number of times that this can be iterated is at most $18\|\mathcal{H}\| + 4\text{simp}(C)$. This because each step gives rise to an arc of $C \cap \mathcal{H}^0$ that is outermost in its parallelism class. The number of such arcs is at most $18\|\mathcal{H}\| + 4\text{simp}(C)$, by Lemma 2.18. Hence after at most this number of isotopies, we end with a 1-manifold having smaller simplification number than C . \square

The following algorithm establishes that we can implement the isotopy across the maximal enlargement of an isotopy region.

Lemma 5.9 (ISOTOPY ACROSS MAXIMAL ENLARGEMENT). *There is an algorithm that takes as its input the following:*

- (1) a handle structure \mathcal{H} for a compact surface S ;
- (2) a standard 1-manifold C properly embedded in S ;
- (3) an arc of $C \cap \mathcal{H}^0$ transversely oriented into a simplifying disc D ;

and outputs the standard 1-manifold C' that is obtained by performing the isotopy across the maximal enlargement of the isotopy region associated with D , and possibly removing some boundary-parallel arc components. It also outputs the number of arcs of intersection between C and this maximal enlargement. The running time is bounded above by a polynomial function of $\|\mathcal{H}\|$, $\text{simp}(C)$, and $\log(w(C) + |\partial C|)$.

Proof. Let R be the isotopy region associated with D and let α be $\partial R \cap C$. By the algorithm LOCATION OF PARALLELITY BUNDLE COMPONENT given in Theorem 4.8, we can determine the location on C of the point or points $\partial\alpha - \partial S$. If $\partial\alpha - \partial S = \emptyset$, then α is a boundary-parallel arc component of C , in which case we simply remove α and all parallel components. So suppose that $\partial\alpha - \partial S \neq \emptyset$. We need to be able to determine the collection of arcs in C that run parallel to α . The point or points of $\partial\alpha - \partial S$ lie (after a small isotopy) in a component β of $\mathcal{H}^0 \cap \mathcal{H}^1$. We plan to cut α along its intersection with a suitable sub-arc of β , so that the resulting components contain all the sub-arcs of C parallel to α . We might not be able to cut C along the entirety of $C \cap \beta$ because in doing so, we might cut up α or one of its parallel arcs. We therefore need to determine the points of $\alpha \cap \beta$ that are closest to $\partial\alpha$. More precisely, we pick a point x of $\partial\alpha - \partial S$ and we consider the component of β' of $\beta \setminus x$ that does not initially go into the isotopy region. We wish to determine the point of $(\beta' - x) \cap \alpha$ that is closest to x or to determine that there is no such point. This can be achieved using AHT, as follows.

We start by cutting C along all the points of $C \cap \beta'$. We need to determine whether this decomposes α . We can do this by computing the components of $\alpha \setminus \beta'$, identifying the components that contain $\partial\alpha$ and seeing whether there is just one such component (in which case α was not decomposed) or two (in which case α was decomposed). If we did not decompose α , then we know that there are no other points of $\beta' \cap \alpha$ other than x . On other hand, suppose that we did decompose

α . In this case, we shrink β' to a sub-arc containing x that has half as many points of intersection with C (rounded to a nearest integer). We cut C along β' and again determine using AHT whether this decomposes α . If it does, then we shrink β' again to half its size. On the other hand, if the decomposition along β' did not decompose α , then we enlarge β' to $3/2$ of its former size. Repeating in this way, we can eventually home in on the point of $(\beta' - x) \cap \alpha$ closest to x .

In the case where there are no other points of $\beta \cap \alpha$ other than $\partial\alpha$, we simply cut C along $C \cap \beta$, use AHT to determine its components and then count how many are parallel to α . In the case where there are points of $\beta \cap \alpha$ other than $\partial\alpha$, we let d be the smallest number of points of $C \cap \beta$ strictly between $\partial\alpha$ and another point of intersection with $\beta \cap \alpha$. Let β'' be the sub-arc of β containing $\partial\alpha$ and that extends $\lfloor d/2 \rfloor$ points of $C \cap \beta$ in each direction from $\partial\alpha$. We cut C along $C \cap \beta''$ and then use AHT to determine the number of components parallel to α .

In this way, we determine the number of parallel copies of α in the maximal enlargement of the isotopy region. We remove these from C and replace them by the same number of short arcs in the 0-handle incident to β . The result is C' . \square

We now consider the case where D is self-returning. We would like to argue that the associated isotopy region R can be combined with D to form an embedded annulus or Möbius band. However, there is a situation where this might not be possible: when the horizontal boundary of R contains an arc that is parallel to $D \cap C$. (See Figure 13.) We deal with this possibility in the following lemma.

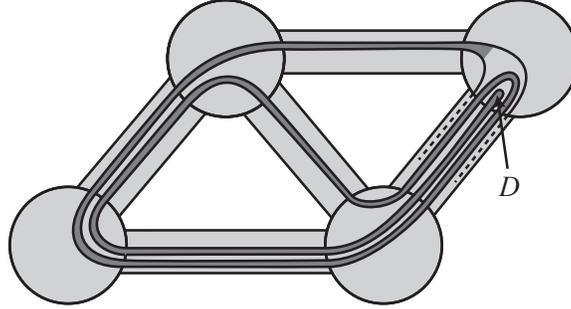


FIGURE 13. The isotopy region R for D is shown shaded. It is self-returning. Its horizontal boundary contains arcs of intersection with \mathcal{H}^0 that are parallel to $D \cap C$.

Lemma 5.10. *Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C , and let R be the associated isotopy region. Suppose that the horizontal boundary of R contains a standard arc in \mathcal{H}^0 that is parallel to $D \cap C$. Then after performing at most $36\|\mathcal{H}\| + 8\text{simp}(C)$ isotopies across maximal enlargements of isotopy regions containing D , the resulting 1-manifold C'' satisfies $\text{simp}(C'') < \text{simp}(C)$.*

Proof. Let α be the sub-arc of C that is the union of $D \cap C$ and the horizontal boundary of R . By Lemma 5.8, we may assume that R is self-returning. As in the proof of Lemma 5.8, we define a nested collection of discs embedded in S . The first is the maximal enlargement D_1 of R . If there is an arc of $C \cap \mathcal{H}^0$ normally parallel to $D \cap C$ that is not contained in D_1 , then we consider one that

is outermost in \mathcal{H}^0 and extend it as far as possible running parallel to ∂D_1 , to form an arc β . Note that β cannot run parallel to the entire length of α , since in that case it would be part of D_1 . We define the next disc D_2 to contain D_1 and to include β and all normally parallel copies. We continue in this way until there are no more arcs of $C \cap \mathcal{H}^0$ normally parallel to $D \cap C$. This forms a collection of discs $D_1 \subset D_2 \subset \dots \subset D_r$ where $r \leq 18\|\mathcal{H}\| + 4\text{simp}(C)$. We then perform all these isotopies, giving a 1-manifold C' . Since R is self-returning, a component of $C' \cap \mathcal{H}^0$ is normally parallel to $C \cap D$. An outermost such arc in \mathcal{H}^0 bounds a simplifying disc D' . This contains D . We claim that D' is not self-returning. Assuming this claim, Lemma 5.8 then implies that after $18\|\mathcal{H}\| + 4\text{simp}(C)$ further isotopies across maximal enlargements of isotopy regions, we obtain a 1-manifold C'' satisfying $\text{simp}(C'') < \text{simp}(C') \leq \text{simp}(C)$.

Suppose that the claim is not true. Consider the isotopy region R' associated with D' . This consists of D' , followed by a component of the parallelity bundle or semi-parallelity bundle of $S \setminus C'$, possibly followed by an interpolation 0-handle, possibly followed by a component of the semi-parallelity bundle. The first 1-handle of R' contains the first 1-handle of R . The next 0-handle of R' contains the next 0-handle of R , and so on. Since we are assuming that R' is self-returning, we deduce that there is a one-one correspondence between the handles of R and the handles of R' , where the correspondence is that the former handle lies in the latter. In particular, each handle of R' contains exactly one handle of R . But we are assuming that the horizontal boundary of R contains a standard arc that is parallel to $D \cap C$. Hence, R enters the space between $D \cap C$ and $D' \cap C'$. This implies that the first handle of R' contains more than one handle of R . This is a contradiction. \square

We are now left with the situation where D is self-returning, but where the horizontal boundary of the isotopy region R does not run parallel to $D \cap C$ at any point. The following lemma partially deals with this case.

Lemma 5.11. *Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C , and let R be the associated isotopy region. Suppose R is self-returning and that no arc of $\partial_h R$ is normally parallel to $D \cap C$. Let D_+ be the maximal enlargement of R . Suppose that $|D_+ \cap C|$ is strictly less than the number of arcs of $C \cap \mathcal{H}_0$ normally parallel to $D \cap C$, including $D \cap C$. Let C' be the 1-manifold obtained from C by the isotopy across this maximal enlargement. Then $C' \cap \mathcal{H}^0$ has a component that is parallel to $D \cap C$ and that bounds a simplifying disc for C' which is not self-returning.*

Proof. Since there are more arcs of $C \cap \mathcal{H}_0$ normally parallel to $D \cap C$ than there are arcs of $C \cap D_+$, and since the horizontal boundary of R does not run parallel to $D \cap C$ at any point, we deduce that some arc of $C \cap \mathcal{H}^0$ parallel to $D \cap C$ is not moved in the isotopy. An outermost such arc bounds a simplifying disc D' for C' . The associated isotopy region R' cannot be self-returning, as otherwise $\partial_h R'$ would have been included in D_+ . \square

We still have not fully dealt with the case of self-returning isotopy regions. For such regions containing an interior-simplifying disc, the following structure will be useful.

Definition 5.12. Let S be a compact surface with a handle structure \mathcal{H} . Let C be a standard 1-manifold properly embedded in S . A *Reeb region* is a subsurface A of S with the following properties:

- (1) A is a thin regular neighbourhood of a standard simple closed curve, and hence is an annulus or Möbius band;
- (2) the intersection between A and C consists of a union of essential arcs in A , all of which are normally parallel;
- (3) each such arc is a concatenation of three arcs $\alpha_-, \alpha_0, \alpha_+$;
- (4) α_- starts at a point of $\partial A \cap \mathcal{H}^0 \cap \mathcal{H}^1$, goes into \mathcal{H}^0 and then runs normally parallel to ∂A in some direction, possibly winding several times around that component but going at least once around that component, and then ending on the same component of $\mathcal{H}^0 \cap \mathcal{H}^1$ that is started on;
- (5) α_0 is non-normal arc properly embedded in \mathcal{H}^0 with endpoints in the same component of $\mathcal{H}^0 \cap \mathcal{H}^1$;
- (6) α_+ starts at a point of ∂A in the same component of $\mathcal{H}^0 \cap \mathcal{H}^1$ as $\partial\alpha_-$, then winds around ∂A the same number of times as α_- and then ends on the same component of $\mathcal{H}^0 \cap \mathcal{H}^1$ that is started on.

A *Reeb isotopy* is supported in a Reeb region as above. It replaces $A \cap C$ with normally parallel non-normal arcs in \mathcal{H}^0 . A Reeb region A is *maximal* if, for any other Reeb region A' containing it, there is an isotopy of S preserving C and each handle of \mathcal{H} taking A' to A .

In the case of a boundary-simplifying disc, we have the following structure.

Definition 5.13. Let S be a compact surface with a handle structure \mathcal{H} . Let C be a standard 1-manifold properly embedded in S . A *semi-Reeb region* is an annular subsurface A of S with the following properties:

- (1) A is a thin regular neighbourhood of a component of ∂S ;
- (2) the intersection between A and C consists of a union of essential arcs in A , all of which are normally parallel;
- (3) each such arcs is a concatenation of two arcs α_0, α_+ ;
- (4) α_+ starts at a point of $(\partial A - \partial S) \cap \mathcal{H}^0 \cap \mathcal{H}^1$, goes into \mathcal{H}^0 and then runs normally parallel to $\partial A - \partial S$ in some direction, possibly winding several times around that component but going at least once around that component, and then ending on the same component of $\mathcal{H}^0 \cap \mathcal{H}^1$ that it started on;
- (5) α_0 is non-normal arc properly embedded in \mathcal{H}^0 , ending on ∂S .

A *semi-Reeb isotopy* is supported in a semi-Reeb region as above. It replaces $A \cap C$ with normally parallel arcs in \mathcal{H}^0 . A semi-Reeb region A is *maximal* if, for any other semi-Reeb region A' containing it, there is an isotopy of S preserving C and each handle of \mathcal{H} taking A' to A .

Lemma 5.14. *Let S be a compact surface with a handle structure \mathcal{H} . Let C be a standard 1-manifold properly embedded in S , and let C' be the 1-manifold that results from a Reeb isotopy of C . Then $\text{simp}(C') \leq \text{simp}(C)$.*

Proof. We use the terminology of Definitions 5.12 and 5.13. Suppose that each arc α_+ winds n times around ∂A . Then the Reeb or semi-Reeb isotopy is the composition of n isotopies across maximal enlargements of isotopy regions. Hence, by Lemma 5.7, this does not increase the simplification number. \square

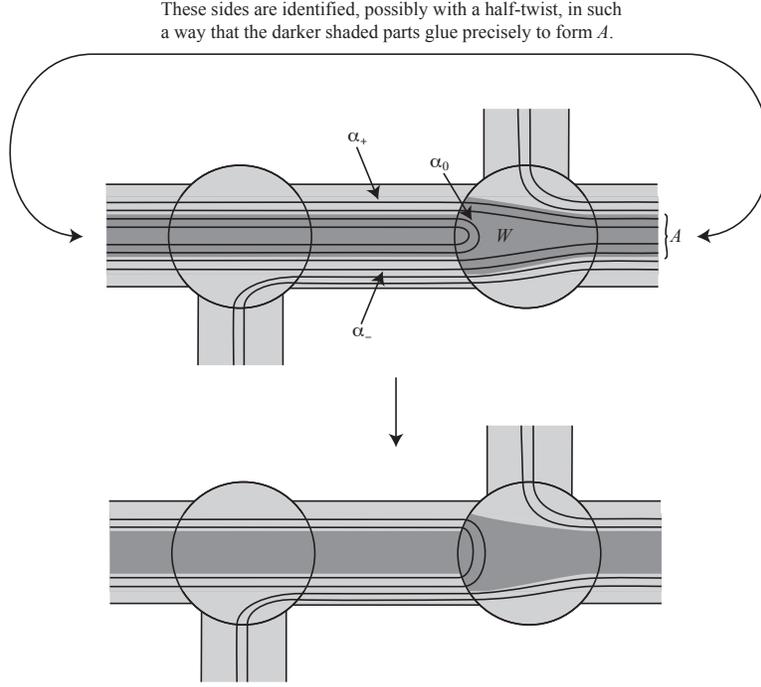


FIGURE 14. A maximal Reeb region A is shaded, and the associated Reeb isotopy is shown. The left and right sides of these figures are identified, but possibly via a reflection, in which case A would be a Möbius band

Clearly, a Reeb or semi-Reeb isotopy across a maximal Reeb or semi-Reeb region A does not necessarily reduce the simplification number since there is still a simplifying disc in A for the resulting 1-manifold.

Lemma 5.15. *Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C , and let R be the associated isotopy region. Suppose R is self-returning and that no arc of $\partial_h R$ is normally parallel to $D \cap C$. Let D_+ be the maximal enlargement of R . Suppose that $|D_+ \cap C|$ is not strictly less than the number of arcs of $C \cap \mathcal{H}_0$ normally parallel to $D \cap C$, including $D \cap C$. Then D_+ lies in a Reeb or semi-Reeb region for C .*

Proof. Let C' be the 1-manifold obtained by isotoping C across D_+ . The arc $\partial D_+ \cap C$ is incident to a non-normal arc of $C' \cap \mathcal{H}^0$. Let E be the component of $\mathcal{H}^0 \setminus C$ that intersects \mathcal{H}^1 in a single component and intersects ∂S in at most one component. Then $D_+ \cup E$ is a Reeb annulus for C . \square

Lemma 5.16. *Let \mathcal{H} be a handle structure for a compact surface S . Let C be a standard 1-manifold properly embedded in S . Let D be a simplifying disc for C , and suppose that D lies in a maximal Reeb or semi-Reeb region A . Let C' be the 1-manifold obtained by the Reeb or semi-Reeb isotopy, and let D' be the resulting simplifying disc for C' . Suppose that D' is self-returning. Then D' satisfies the hypotheses of Lemma 5.11.*

Proof. Let R' be the isotopy region for D' . Since R' is self-returning, it is a regular neighbourhood of the isotopy region for D . Hence, no arc of $\partial_h R' \cap \mathcal{H}^0$ is normally parallel to $D' \cap C'$, which is one of the hypotheses of Lemma 5.11. Furthermore, suppose that the number of arcs of intersection between C' and the maximal enlargement of R' is not strictly less than the number of arcs of $C' \cap \mathcal{H}^0$ normally parallel to $D' \cap C'$. Then we could have expanded A to be a bigger Reeb or semi-Reeb region for C , contradicting its maximality. Thus, D' does satisfy the hypotheses of Lemma 5.11. \square

Thus, we are now in a position to give the algorithm NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE.

Proof of Theorem 2.15. We perform the following operations in a loop. First, we search for the existence of a maximal Reeb or semi-Reeb region. We will explain how to do this below. If there is such a Reeb or semi-Reeb region, we perform this Reeb or semi-Reeb isotopy, creating a 1-manifold C' . Denote the simplifying disc for C' in the Reeb or semi-Reeb region by D' . If there was no Reeb or semi-Reeb region, then we let $C' = C$ and we pick any simplifying disc D' for C' .

We then perform the isotopy across the maximal enlargement of the isotopy region associated with D' . If this reduces the simplification number, then we return to the start of the loop. On the other hand, if the simplification number is unchanged, then we know that the new 1-manifold C'' has an arc of $C'' \cap \mathcal{H}^0$ that runs parallel to $D' \cap C'$. An outermost such arc bounds a simplifying disc. We repeat the procedure in this paragraph with this disc. We continue until the simplification number decreases, at which point we return to the start of the loop. This continues until there are no more simplifying discs.

We now explain why this procedure works. If there is no Reeb or semi-Reeb region, then, by Lemma 5.15, each simplifying disc D' for C' satisfies at least one of the following:

- (1) D' is not self-returning;
- (2) the horizontal boundary of its isotopy region contains a standard arc that runs parallel to $D' \cap C'$;
- (3) the horizontal boundary of the isotopy region R' does not contain a standard arc that runs parallel to $D' \cap C'$, and the number of arcs parallel to $D' \cap C'$ (including $D' \cap C'$) is greater than the number of arcs of C' in the maximal enlargement of R' .

In the first case, Lemma 5.8 gives that after performing at most $18\|\mathcal{H}\| + 4\text{simp}(C)$ isotopies across maximal enlargements, the simplification number goes down. In the second case, a similar conclusion holds by Lemma 5.10. In the third case, Lemma 5.11 implies that the 1-manifold obtained by isotoping C' across the maximal enlargement of R' has a simplifying disc that is not self-returning. Hence, after at most $18\|\mathcal{H}\| + 4\text{simp}(C)$ further isotopies, the simplification number goes down. On the other hand, if there is a Reeb or semi-Reeb region for C , then there is a maximal one A , and we let C' be the 1-manifold obtained by the Reeb or semi-Reeb isotopy and we let D' be the simplifying disc for C' contained in A . Lemma 5.16 and Lemma 5.11 imply that either D' is not self-returning or the 1-manifold obtained by the isotopy across the maximal enlargement of D' has such a simplifying disc. Hence, after at most $18\|\mathcal{H}\| + 4\text{simp}(C)$ further isotopies, the simplification number decreases.

We now describe how to find a maximal Reeb region if one exists. We can readily find all the interior-simplifying discs for the 1-manifold C , by considering each parallelism class of arcs of $C \cap \mathcal{H}^0$. We consider each interior-simplifying disc in turn. So let us focus on one such disc D and its isotopy region R . We first determine whether R is self-returning as follows. By applying CUT ALONG STANDARD 1-MANIFOLD (Theorem 4.6) to $S \setminus C$ and Remark 4.7, we can determine the component W of $\mathcal{H}^0 \setminus C$ containing $\partial_v R \setminus D$. (See Figure 14.) Then R is self-returning if and only if the following all hold:

- (1) W lies in the interior of S ;
- (2) the intersection between ∂W and $\mathcal{H}^0 \cap C$ consists of three arcs β_- , β_0 and β_+ ;
- (3) one of these arcs β_0 is normally parallel to $D \cap C$ and all arcs of $C \cap \mathcal{H}^0$ between β_0 and $D \cap C$ are normally parallel;
- (4) the remaining two arcs β_- and β_+ of $\mathcal{H}^0 \cap C \cap \partial W$ are normally parallel to each other.

These conditions are easily checked. Suppose that R is self-returning. Then we can easily check whether the horizontal boundary of R contains no arc of $\mathcal{H}^0 \cap C$ that is normally parallel to $D \cap C$ using LOCATION OF PARALLELITY BUNDLE COMPONENT (Theorem 4.8). Using ISOTOPY ACROSS MAXIMAL ENLARGEMENT (Lemma 5.9), we can determine the number p of arcs of intersection between C and the maximal enlargement of R . For R to be part of a Reeb region, this number needs to be equal to the number of components of $C \cap \mathcal{H}^0$ normally parallel to $C \cap D$, including $C \cap D$ itself.

Thus, using the above procedure, we can determine whether R is part of a Reeb region. Suppose it is. We then find the maximal Reeb region containing R , as follows. Let β be the component of $\mathcal{H}^0 \cap \mathcal{H}^1$ incident to D . As in the proof of Lemma 5.9, we can determine whether R intersects $\beta - D$. If it does not, then we cut C along β and count the number of components of $C - \beta$ normally parallel to each component of $\partial_h R$. Let q_1 and q_2 be these number of arcs. Let q be the largest integer multiple of p that is at most $\min\{q_1, q_2\}$. Then the union of this many copies of $\partial_h R$, together with all arcs of $C \cap \mathcal{H}^0$ parallel to $C \cap D$, forms the arcs of intersection between C and the maximal Reeb region A . Removing all these arcs and replacing them with this many copies of $C \cap D$ achieves the Reeb isotopy.

There is a similar procedure for finding a maximal semi-Reeb region if one exists and for performing the associated semi-Reeb isotopy. \square

Remark 5.17. The isotopies given in the above proof can be realised as a union of isotopies across simplifying discs and their incident 1-handles. However, of course, the number of such isotopies may be considerably higher than the number of moves in the above algorithm.

6. MINIMAL POSITION FOR CURVES IN SURFACES

In this section, we embark on the computation of the geometric intersection number of two properly embedded 1-manifolds P and C . In fact, we will actually isotope C into minimal position with respect to P , but in the special case where P has very restricted intersection with the triangulation or handle structure.

Definition 6.1. Let C and P be 1-manifolds properly embedded in a compact surface S that intersect transversely. A *bigon* for C and P is a disc D embedded in

the interior of S with interior disjoint from $C \cup P$, and with boundary that is the concatenation of an arc in C and an arc in P . A *half-bigon* a disc D embedded in S with interior disjoint from $C \cup P$, and with boundary that is the concatenation of an arc in C , an arc in P and an arc in ∂S . The *bigon number* $\text{bigon}(C, P)$ is equal to the sum of the number of half-bigons and twice the number of bigons. We say that C and P are in *minimal position* if they have no bigon or half-bigon.

The following key fact is well known [8, Proposition 1.7 and Section 1.2.7].

Lemma 6.2. *Let C and P be 1-manifolds properly embedded in a compact surface S that intersect transversely. Then $|C \cap P|$ is minimised up to isotopy if and only if they have no bigon or half-bigon.*

Theorem 6.3 (MINIMAL POSITION FOR 1-MANIFOLDS). *There is an algorithm that takes, as its input, the following:*

- (1) a triangulation \mathcal{T} or handle structure \mathcal{H} of a compact surface S ;
- (2) a 1-manifold P properly embedded in S that either is a subset of the 1-skeleton of \mathcal{T} or that intersects \mathcal{H} in the following way: it is disjoint from 0-handles; it intersects each 1-handle in the empty set or in the co-core of the handle; it intersects each 2-handle in at most one properly embedded arc;
- (3) a standard 1-manifold C ;

and provides as its output a normal 1-manifold C' that is obtained from C by an isotopy, plus possibly removing some components that are boundary parallel arcs and curves that bound discs. This 1-manifold C' is in minimal position with respect to P and has weight bounded above by $w(C)(1 + 2|\mathcal{T}|)$ or $w(C)(1 + 2\|\mathcal{H}\|)$. The running time is bounded above by a polynomial function of $\text{simp}(C)$, $\log(w(C) + |\partial C|)$, and $|\mathcal{T}|$ or $\|\mathcal{H}\|$.

When P is arranged with respect to the handle structure \mathcal{H} as in (2) above, we say that it is *cellular with respect to the cell structure dual to \mathcal{H}* .

The proof will follow that in Section 5, but with bigons and half-bigons replacing the role of simplifying discs. In Section 5, the primary measure of complexity of the standard 1-manifold C was the simplification number $\text{simp}(C)$, but here we use $\text{bigon}(C, P)$. The reader may wonder why the running time for MINIMAL POSITION FOR 1-MANIFOLDS does not depend on $\text{bigon}(C, P)$, given that this quantity is our primary measure of complexity. In fact, the running time does depend on $\text{bigon}(C, P)$, but this can be bounded above in terms of $\|\mathcal{H}\|$ and $\text{simp}(C)$, as follows.

Lemma 6.4. *Let S be a compact surface with a handle structure \mathcal{H} . Let P be a properly embedded 1-manifold that is cellular with respect to the cell structure dual to \mathcal{H} , and let C be a standard 1-manifold. Then*

$$\text{bigon}(C, P) \leq 20\|\mathcal{H}\| + 4\text{simp}(C).$$

Proof. We obtain a handle structure \mathcal{H}' for $S \setminus P$ by removing the handles intersecting P . Hence, $\|\mathcal{H}'\| \leq \|\mathcal{H}\|$. The 1-manifold $C \setminus P$ is standard in \mathcal{H}' . Each simplifying disc for $C \setminus P$ is a simplifying disc for C , but it may have become boundary-simplifying when previously it was interior-simplifying. Hence, $\text{simp}(C \setminus P) \leq \text{simp}(C)$. The components of $C \setminus P$ fall into at most $10\|\mathcal{H}'\| + 2\text{simp}(C \setminus P)$ parallelism classes, by Lemma 2.20. In each parallelism class, only the outermost

two 1-manifolds can be part of a bigon or half-bigon. The lemma follows immediately. \square

Corollary 6.5. *Let S be a compact surface with a triangulation \mathcal{T} . Let P be a properly embedded 1-manifold that is a subset of the 1-skeleton, and let C be a normal 1-manifold. Then $\text{bigon}(C, P) \leq 154|\mathcal{T}|$.*

Proof. Dual to \mathcal{T} is a handle structure \mathcal{H} . Each handle of \mathcal{H} arises from a simplex of \mathcal{T} that does not lie wholly in ∂S . The number of edges of \mathcal{T} not in ∂S is at most $(3/2)|\mathcal{T}|$. The number of vertices of \mathcal{T} is at most $3|\mathcal{T}|$. Hence, the number of handles of \mathcal{H} is at most $(11/2)|\mathcal{T}|$. The 1-manifold C might not be normal with respect to \mathcal{H} , but it has only boundary-simplifying discs and at most two of these can lie in each 0-handle. So by Lemma 6.4, $\text{bigon}(C, P) \leq 20\|\mathcal{H}\| + 4\text{simp}(C) \leq 28\|\mathcal{H}\| \leq 154|\mathcal{T}|$. \square

In MINIMAL POSITION FOR 1-MANIFOLDS, our aim will be to isotope C to a 1-manifold C' with $\text{bigon}(C', P) = 0$. Just as in Figure 7, an isotopy of C across a bigon or half-bigon may result in a 1-manifold C' with $\text{bigon}(C', P) > \text{bigon}(C, P)$. In order to avoid this, we need to consider a more substantial move. We therefore need to introduce an analogue of the isotopy region for a simplifying disc, and to do so, we will introduce analogues of many of the definitions in Sections 2, 4 and 5. We will use the adjective ‘topological’ for many of these analogues.

Definition 6.6. A *topological parallelity region* for C with respect to P consists of a disc with interior disjoint from C and P , and with boundary equal to the concatenation of an arc in C , an arc in P or ∂S , an arc in C and another arc in P or ∂S . The *topological parallelity bundle* is the subset of $S \setminus C$ consisting of the union of the topological parallelity regions. It is an I -bundle \mathcal{B} over a 1-manifold, with $\partial_h \mathcal{B} = \mathcal{B} \cap C$. See Figure 15.

Definition 6.7. A *topological semi-parallelity region* for C with respect to P is a disc with interior disjoint from C and P , and with boundary equal to the concatenation of an arc in C , an arc in P , an arc in ∂S and another arc in P . The *topological semi-parallelity bundle* is the subset of $S \setminus C$ consisting of the union of the semi-parallelity regions. It is an I -bundle \mathcal{B} over a 1-manifold, with $\partial_h \mathcal{B} = \mathcal{B} \cap (C \cup \partial S)$.

Definition 6.8. A *topological interpolating region* for C and P is a disc with interior disjoint from C and P , and with boundary consisting of an arc in C , an arc in P , an arc in C , an arc in P and an arc in ∂S .

Definition 6.9. Let D be a bigon or half-bigon for C and P . Then the *associated topological isotopy region* R is the union of the following subsets of $S \setminus C$:

- (1) the bigon or half-bigon D ;
- (2) the component B_1 of the topological parallelity or topological semi-parallelity bundle incident to D ;
- (3) if D is a bigon and B_1 is incident to a topological interpolating region, then this topological interpolating region is also part of R and so is the incident component B_2 of the topological semi-parallelity bundle;
- (4) a small regular neighbourhood of $\partial_v B_1$ in $S \setminus C$, as well as a small regular neighbourhood of $\partial_v B_2$ if B_2 is defined.

We say that the *horizontal boundary* of R is $(\partial R \cap C) \setminus D$.

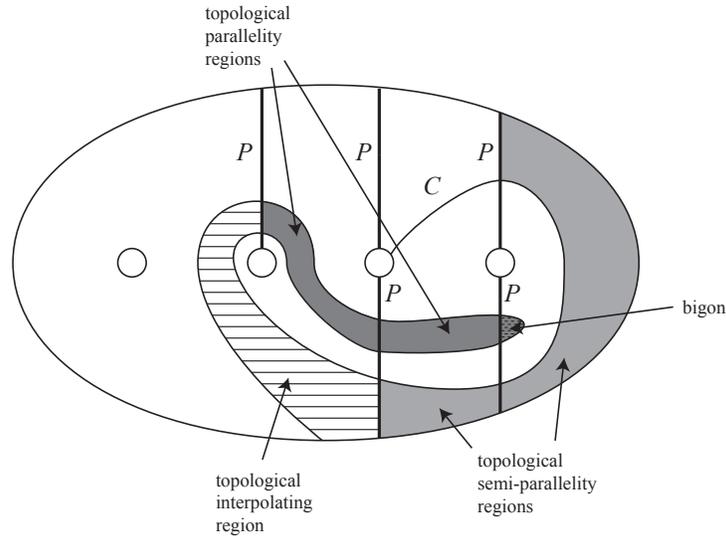


FIGURE 15. A surface S containing properly embedded 1-manifolds C and P . Shown are various complementary regions, which are topological parallelity regions, topological semi-parallelity regions, a topological interpolating region and a bigon.

Definition 6.10. We say that two arcs of $C \setminus P$ in $S \setminus P$ are *topologically parallel* if they cobound a topological parallelity region. More generally, if α_1 and α_2 are sub-arcs of C with $\partial\alpha_1 \subset P$ and $\partial\alpha_2 \subset P$, we say that these arcs are *topologically parallel* if there is a disc R in S , such that ∂R is the concatenation of α_1 , an arc in P or ∂S , α_2 and another arc in P or ∂S , and R is a union of topological parallelity regions.

The following analogue of Lemma 5.4 holds, with essentially the same proof, which is omitted.

Lemma 6.11. *Let C and P be 1-manifolds properly embedded in a compact surface S that intersect transversely. Let D be a bigon or half-bigon for C and P . Let C' be the 1-manifold obtained by isotoping C across the topological isotopy region associated with D . Then $\text{bigon}(C', P) \leq \text{bigon}(C, P)$ and $|C' \cap P| < |C \cap P|$. Furthermore, if $\text{bigon}(C', P) = \text{bigon}(C, P)$, then there is an arc of $C' \setminus P$ that is topologically parallel to $D \cap C$, and all the components of $C \setminus P$ between these two arcs are topologically parallel to $D \cap C$.*

The following is the version of Definition 5.6.

Definition 6.12. Let C and P be 1-manifolds properly embedded in S that intersect transversely. Let D be a bigon or half-bigon for C and P . Let R be the associated topological isotopy region. Then an *enlargement* of R consists of a disc D_+ containing R and where $C \cap D_+$ consists of arcs all parallel to $R \cap C$, one of which lies in ∂D_+ . An enlargement is *maximal* if it does lie in a bigger enlargement, up to an isotopy that preserves C and P . The *isotopy across D_+* replaces the arcs $D_+ \cap C$ with arcs parallel to $\partial R \setminus (C \cup \partial S)$.

We also have the analogue of ISOTOPY ACROSS MAXIMAL ENLARGEMENT (Lemma 5.9) and LOCATION OF PARALLELITY BUNDLE COMPONENT (Theorem 4.8).

Lemma 6.13 (ISOTOPY ACROSS MAXIMAL ENLARGEMENT OF BIGON/HALF-BIGON). *There is an algorithm that takes as its input the following:*

- (1) a triangulation \mathcal{T} or handle structure \mathcal{H} for a compact surface S ;
- (2) a properly embedded essential 1-manifold P that either is a subset of the 1-skeleton of \mathcal{T} or is cellular with respect to the cell complex dual to \mathcal{H} ;
- (3) a standard 1-manifold C properly embedded in S ;
- (4) an arc of $C \cap \mathcal{H}^0$ transversely oriented into a bigon or half-bigon D for C and P ;

and outputs the standard 1-manifold C' that is obtained by the isotopy across the maximal enlargement of the isotopy region associated with D and possibly removing some boundary-parallel arc components. It also provides the location of the sub-arc of P that the components of $C' \setminus C$ run along. The running time is bounded above by a polynomial function of $|\mathcal{T}|$ or $\|\mathcal{H}\|$, $\text{simp}(C)$, and $\log(w(C) + |\partial C|)$.

We also have the analogue of Definition 5.5.

Definition 6.14. Let D be a bigon or half-bigon for C and P , and let R be the associated topological isotopy region. Let C' be the 1-manifold obtained by the isotopy across R . Then D and R are *self-returning* if the arc of $C' \setminus P$ incident to R is topologically parallel to $D \cap C$, and all the components of $C' \setminus P$ between these two arcs are topologically parallel to $D \cap C$.

The following is a version of Lemma 5.8, with almost the same proof.

Lemma 6.15. *Let C and P be 1-manifolds properly embedded in a compact surface S that intersect transversely. Let D be a bigon or half-bigon for C and P . Suppose that D is not self-returning. Then after performing at most $-8\chi(S) + 3\text{bigon}(C, P) + |\partial P|$ isotopies across maximal enlargements of topological isotopy regions containing D , the resulting 1-manifold C'' satisfies $\text{bigon}(C'', P) < \text{bigon}(C, P)$.*

The only deviation from the proof of Lemma 5.8 is that instead of using Lemma 2.18, we use the following.

Lemma 6.16. *Let C and P be 1-manifolds properly embedded in the compact surface S that intersect transversely. Suppose that no component of C or P bounds a disc in the interior of S with interior disjoint from $C \cup P$, and no component is topologically parallel to an arc in ∂S , via a disc with interior disjoint from $C \cup P$. Then the number of topological parallelism classes of components of $C' \setminus P$ that are not disjoint from P is at most $-4\chi(S) + (3/2)\text{bigon}(C, P) + |\partial P|/2$.*

Proof. First note that we may remove components of C that are disjoint from P . Similarly, we may remove components of P that are disjoint from C . Define the *index* of a component R of $S \setminus (C \cup P)$ to be

$$I(R) = -\chi(R) + |R \cap C \cap P|/4 + |R \cap C \cap \partial S|/4 + |R \cap P \cap \partial S|/4.$$

Then it is an easy calculation that the sum, over all components R , of the index of R , is equal to $-\chi(S)$. We now consider the regions R with negative index. Necessarily, such a region R must be a disc intersecting $(C \cap P) \cup (C \cap \partial S) \cup (P \cap \partial S)$ at most

three times. By our hypotheses, it is therefore a bigon or half-bigon. A bigon has index $-1/2$ and a half-bigon has index $-1/4$. Hence,

$$\sum_{I(R)>0} I(R) = -\chi(S) + \text{bigon}(C, P)/4.$$

But each region R that is not a topological parallelity or topological semi-parallelity region and is not a bigon or half-bigon has positive index at least $1/4$. (Here we are using the assumption that every component of C intersects P and that every component of P intersects C . This is to rule out disc regions R consisting of an arc in ∂S , an arc of C or P , another arc in ∂S and another arc of C or P . It also rules out annuli and Möbius band components with boundary disjoint from $C \cap P$.) Therefore, the number of such regions is at most $-4\chi(S) + \text{bigon}(C, P)$. The number of regions with negative index is at most $\text{bigon}(C, P)$. Hence, in total, the number of regions that are not parallelity or semi-parallelity regions is at most $-4\chi(S) + 2\text{bigon}(C, P)$.

To bound the number of topological parallelism classes of components of $C \setminus P$, we replace each topological parallelism class with a single component. We view the resulting 1-manifold α as properly embedded in $S \setminus P$. The complementary regions of α correspond to the complementary regions of $C \cup P$ that are not topological parallelity regions. Each such region R with positive index has at most $8I(R)$ arcs of C in its boundary, where the extreme case is given by a topological interpolating region with index $1/4$ and with 2 arcs of C in its boundary. So, the number of arcs of α in the boundary of the regions with positive index is at most $8 \sum_{I(R)>0} I(R)$, which is at most $-8\chi(S) + 2\text{bigon}(C, P)$. The number of arcs of α in the boundary of the regions with negative index is at most $\text{bigon}(C, P)$. The number of arcs of α in the boundary of regions with zero index is at most the number of topological semi-parallelity regions, which is at most $|\partial P|$. So in total, the number of arcs of α in the boundary of these regions is at most $-8\chi(S) + 3\text{bigon}(C, P) + |\partial P|$. Each arc of α is counted twice, and hence we find the required upper bound on the number of topological parallelism classes of $C \setminus P$. \square

The following is an analogue of Lemma 5.10.

Lemma 6.17. *Let C and P be 1-manifolds properly embedded in a compact surface S that intersect transversely. Let D be a bigon or half-bigon for C and P . Let R be the associated topological isotopy region. Suppose that ∂R contains an arc of $C \setminus P$ that is topological parallel to $D \cap C$ but not equal to $D \cap C$. Then after performing at most $-16\chi(S) + 6\text{bigon}(C, P) + 2|\partial P|$ isotopies across maximal enlargements of topological isotopy regions containing D , the resulting 1-manifold C'' satisfies $\text{bigon}(C'', P) < \text{bigon}(C, P)$.*

There is also an obvious analogue of Lemma 5.11, which we omit.

Finally, we introduce the analogue of a Reeb region in this context.

Definition 6.18. Let S be a compact surface containing properly embedded 1-manifolds P and C that intersect transversely. A *topological Reeb region* for C with respect to P is a subsurface A in the interior of S with the following properties:

- (1) A is an annulus or Möbius band;
- (2) $A \cap P$ is a collection of arcs embedded in A , where each such arc is a concatenation of an arc or point in ∂A , a properly embedded essential arc in A , and another arc or point in ∂A ;

- (3) A contains a single bigon D in its interior;
- (4) the intersection between A and C consists of a union of essential arcs in A , all of which are topologically parallel;
- (5) each such arc is a concatenation of three arcs $\alpha_-, \alpha_0, \alpha_+$;
- (6) α_- starts at a point of $\partial A \cap P$ in the same component of $A \cap P$ as $D \cap P$ and then runs topologically parallel to ∂A in some direction, possibly winding several times around that component, but running at least once along that component;
- (7) α_0 is topologically parallel to (or equal to) $D \cap C$;
- (8) α_+ starts at a point of $\partial A \cap P$ in the same component of $A \cap P$ as $D \cap P$, and winds around ∂A the same number of times as α_- and in the same direction.

A *topological Reeb isotopy* is supported in a topological Reeb region as above. It replaces $A \cap C$ with parallel arcs that are essential in A and with interiors disjoint from P . A topological Reeb region A is *maximal* if any other topological Reeb region containing it is isotopic to A , via an isotopy preserving P and C .

We also have the following analogue of a semi-Reeb region.

Definition 6.19. Let S be a compact surface, and let C and P be 1-manifolds properly embedded in S that intersect transversely. A *topological semi-Reeb region* for C with respect to P is an annular subsurface A of S with the following properties:

- (1) A is a regular neighbourhood of a component of ∂S ;
- (2) $A \cap P$ is a collection of arcs embedded in A , where each such arc is a concatenation of an arc or point in $\partial A - \partial S$ and a properly embedded essential arc in A ;
- (3) A contains a single half-bigon D and no bigon;
- (4) the intersection between A and C consists of a union of essential arcs in A , all of which are topologically parallel;
- (5) each such arcs is a concatenation of two arcs α_0, α_+ ;
- (6) α_+ starts at a point of $(\partial A - \partial S) \cap P$ in the same component of $A \cap P$ as $D \cap P$ and then runs topologically parallel to $\partial A - \partial S$ in some direction, possibly winding several times around that component, but running at least once along that component;
- (7) α_0 is topologically parallel to (or equal to) $\partial D \setminus (\partial S \cup P)$.

(See Figure 16.) A *semi-Reeb isotopy* is supported in a semi-Reeb region as above. It replaces $A \cap C$ with parallel arcs that are essential in A and with interiors disjoint from P . A semi-Reeb region A is *maximal* if any other semi-Reeb region containing it is isotopic to A , via an isotopy preserving P and C .

There are obvious analogues of Lemmas 5.15 and 5.16 for topological Reeb or semi-Reeb regions.

The following lemma controls the increase in the simplification number of the 1-manifold C as it is modified by the above isotopies.

Lemma 6.20. *Let S be a compact surface with a handle structure \mathcal{H} . Let P be a 1-manifold properly embedded in S that is cellular with respect to the cell structure dual to \mathcal{H} . Let C be a standard 1-manifold properly embedded in S . Let C' be obtained from C either by an isotopy across the maximal enlargement of an isotopy region associated with a bigon or half-bigon or by a Reeb or semi-Reeb isotopy*

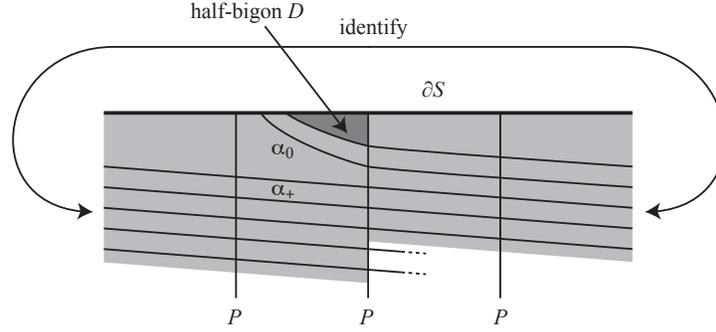


FIGURE 16. A topological semi-Reeb region

associated with a bigon or half-bigon. Then $\text{simp}(C') \leq \text{simp}(C) + 8\|\mathcal{H}\|$ and $w(C') \leq w(C) + 2\|\mathcal{H}\|(|C \cap P| - |C' \cap P|)$.

Proof. In the isotopy, sub-arcs of C are replaced by arcs normally parallel to a sub-arc of an arc of $P \setminus C$. Specifically, this arc of $P \setminus C$ is shifted out of the 2-handles and 1-handles so that it becomes an arc α running through the 0-handles and 1-handles. When this arc is shifted into the 0-handles and 1-handles, it runs over each 1-handle at most twice. So the number of components of $\alpha \cap \mathcal{H}^0$ is at most $2|\mathcal{H}^1| \leq 2\|\mathcal{H}\|$. Each component of $\alpha \cap \mathcal{H}^0$ may give rise to multiple components of $(C' \setminus C) \cap \mathcal{H}^0$, but only the outermost two of these can be part of a simplifying disc for C' . Thus the number of new simplifying discs that are created is at most $4\|\mathcal{H}\|$. Each simplifying disc contributes at most 2 to $\text{simp}(C')$ and so $\text{simp}(C') \leq \text{simp}(C) + 8\|\mathcal{H}\|$.

The number of arcs of $C' \setminus P$ that are inserted is at most $|P \cap C| - |P \cap C'|$, because each new arc that is inserted leads to a reduction in the number of intersections with P by at least 1. As observed above, each of these new arcs has weight at most $2\|\mathcal{H}\|$. Hence, $w(C') \leq w(C) + 2\|\mathcal{H}\|(|C \cap P| - |C' \cap P|)$. \square

Proof of Theorem 6.3. First we search for a maximal Reeb or semi-Reeb region. If there is one, then we perform the Reeb or semi-Reeb isotopy, creating a 1-manifold C' . Let D' be the bigon or half-bigon for C' in the Reeb or semi-Reeb region. If there was no Reeb or semi-Reeb region, then we pick any bigon or half-bigon D' and set $C' = C$. We then perform the isotopy of C' across the maximal enlargement of the isotopy region associated with D' . If this reduces the bigon number, then we return to the start of the loop. On the other hand, if the bigon number is unchanged, then we know, by Lemma 6.11, that the new 1-manifold C'' has an arc of $C'' \setminus P$ that runs parallel to $D' \cap C'$. An outermost such arc bounds a bigon or half-bigon. We repeat the above procedure with this disc. We continue until the bigon number decreases, at which point we return to the start of the loop. This continues until there are no more bigons or half-bigons.

The method of finding maximal Reeb and semi-Reeb regions and performing the isotopies is exactly analogous to that in the proof of Theorem 2.15. The argument that this procedure works is also similar, using Lemmas 6.15 and 6.17 and analogues of Lemmas 5.11, 5.15 and 5.16. Note that the upper bounds on the number of isotopies in these lemmas depend on $|\partial P|$. But when P is as in (2) in Theorem 6.3, $|\partial P|$ is at most $2\|\mathcal{H}\|$ or $3|\mathcal{T}|$.

The output from the above procedure is a standard curve C'' that is in minimal position with respect to P , but that is not necessarily normal.

By Lemma 6.20, each of the isotopies that we have used increases the simplification number by at most $4\|\mathcal{H}\|$. Hence, $\text{simp}(C'')$ is bounded above by a polynomial function of $\text{simp}(C)$ and $\|\mathcal{H}\|$.

We also need to control the weight of C'' . This 1-manifold is produced from C in steps, where each step is a Reeb isotopy, a semi-Reeb isotopy or an isotopy across the maximal enlargement of an isotopy region. Let C_i be the 1-manifold obtained after the i th step. We claim that the weight of C_i is at most $w(C) + 2(|C \cap P| - |C_i \cap P|)|\mathcal{T}|$ or $w(C) + 2(|C \cap P| - |C_i \cap P|)\|\mathcal{H}\|$. When $i = 0$, this is trivially true. When we have a handle structure \mathcal{H} , then in going from C_i to C_{i+1} , the weight increases by at most $2\|\mathcal{H}\|(|C_i \cap P| - |C_{i+1} \cap P|)$, by Lemma 6.20. So, we have

$$\begin{aligned} w(C_{i+1}) &\leq w(C_i) + 2\|\mathcal{H}\|(|C_i \cap P| - |C_{i+1} \cap P|) \\ &\leq w(C) + 2(|C \cap P| - |C_i \cap P|)\|\mathcal{H}\| + 2\|\mathcal{H}\|(|C_i \cap P| - |C_{i+1} \cap P|) \\ &= w(C) + 2(|C \cap P| - |C_{i+1} \cap P|)\|\mathcal{H}\|. \end{aligned}$$

This establishes the claim. Thus, $w(C'') \leq w(C) + 2|C \cap P|\|\mathcal{H}\| \leq w(C)(1 + 2\|\mathcal{H}\|)$. A similar statement is true when S has a triangulation.

We can therefore apply NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE (Theorem 2.15) or NORMALISE STANDARD 1-MANIFOLD IN A TRIANGULATION (Theorem 1.3) to normalise C'' in polynomial time. Note that by Corollary 6.22 below, the resulting 1-manifold C''' remains in minimal position with respect to P . Furthermore, $w(C''') \leq w(C'')$. \square

Lemma 6.21. *Let C be a standard 1-manifold properly embedded in a compact surface S with a handle structure \mathcal{H} . Let P be a properly embedded 1-manifold that is cellular with respect to the cell structure dual to \mathcal{H} . Suppose that C and P are in minimal position. Let C' be obtained from C by isotopy across a simplifying disc and the incident 1-handle. Then C' and P are in minimal position.*

Proof. The simplifying disc does not extend to a bigon or half-bigon of C and P , by the assumption that they are in minimal position. Hence, the isotopy across the simplifying disc is an isotopy supported in the complement of P . Hence, the resulting 1-manifold C' remains in minimal position with respect to P . \square

Corollary 6.22. *Let S , C and \mathcal{H} or \mathcal{T} be as in Theorem 2.15 or Theorem 1.3. Let P be a 1-manifold properly embedded in S that is either cellular with respect to the cell structure dual to \mathcal{H} or a subset of the 1-skeleton of \mathcal{T} . Suppose that C and P are in minimal position. Then the output C' from the algorithm NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE or NORMALISE STANDARD 1-MANIFOLD IN A TRIANGULATION is also in minimal position with respect to P .*

Proof. As discussed in Remark 5.17, the isotopies used in NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE or NORMALISE STANDARD 1-MANIFOLD IN A TRIANGULATION can be viewed as a composition of isotopies, each of which slides the 1-manifold across a simplifying disc and an incident 1-handle. By Lemma 6.21, each of these isotopies keeps the 1-manifold and P in minimal position. \square

7. GEOMETRIC INTERSECTION NUMBER AND THE ISOTOPY PROBLEM

In this section, we complete the proof of Theorems 1.1 and 1.2.

Theorem 1.1 (GEOMETRIC INTERSECTION NUMBER). *Let \mathcal{T} be a triangulation for a compact orientable surface S . Let C_1 and C_2 be normal 1-manifolds properly embedded in S . Then there is an algorithm that provides the geometric intersection number of C_1 and C_2 . The running time of this algorithm is bounded above by a polynomial function of $|\mathcal{T}|$, $\log w(C_1)$ and $\log w(C_2)$.*

Theorem 1.2 (ISOTOPY OF 1-MANIFOLDS). *Let \mathcal{T} be a triangulation for a compact orientable surface S . Let C_1 and C_2 be normal essential 1-manifolds properly embedded in S . Then there is an algorithm that determines whether C_1 and C_2 are isotopic. The running time of this algorithm is bounded above by a polynomial function of $|\mathcal{T}|$, $\log w(C_1)$ and $\log w(C_2)$.*

We follow the approach of Bell [2] and Bell and Webb [3], where they modify \mathcal{T} to a new triangulation \mathcal{T}' with the property that C_2 intersects each edge of \mathcal{T}' as few times as possible. They also keep track of how C_1 sits in \mathcal{T}' .

In order to make our results as general as possible, we also consider *ideal triangulations* of a compact surface S . Recall that this is a representation of $S - \partial S$ as a union of triangles with their edges identified in pairs and with their vertices removed.

Definition 7.1. Let \mathcal{T} be a triangulation or ideal triangulation of a compact surface S . Let e be an edge of \mathcal{T} that is adjacent to distinct triangles. Then the operation of removing e , forming a square, and then inserting the other diagonal of the square, creates a new triangulation or ideal triangulation that is obtained from \mathcal{T} by a *flip*. This is also called a *2-2 Pachner move*.

Definition 7.2. Let \mathcal{T} be a triangulation or ideal triangulation of a compact surface S . Let C be a normal simple closed curve that intersects each edge of \mathcal{T} at most three times. Then, for a natural number n , a *twist of order n* about C is the modification to \mathcal{T} that results from Dehn twisting n times about C .

The following result is essentially a restatement of [11, Theorems 7.4 and 7.5]. A similar result was proved by Bell and Webb [3, Corollary 3.5], but they did not control the running time in terms of the Euler characteristic of the surface.

Theorem 7.3 (SIMPLIFY NORMAL 1-MANIFOLD). *There is an algorithm that takes as its input the following:*

- (1) a 1-vertex triangulation \mathcal{T} of a closed orientable surface S or an ideal triangulation \mathcal{T} of a compact orientable surface S with non-empty boundary,
- (2) a properly embedded connected essential normal 1-manifold C ,

and outputs a sequence of 1-vertex triangulations or ideal triangulations

$$\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_k$$

and, for each i , a properly embedded 1-manifold C_i that is normal in the handle structure dual to \mathcal{T}_i , such that:

- (1) each \mathcal{T}_i is obtained from \mathcal{T}_{i-1} by a flip or a twist of order n_i , where n_i is bounded above by a polynomial function of $w(C)$ and $|\mathcal{T}|$;
- (2) C_i is isotopic to C_{i-1} ;
- (3) each component of C_k intersects each edge of \mathcal{T}_k at most one twice;
- (4) the number of moves k is bounded above by a polynomial function of $|\mathcal{T}|$ and $\log w(C)$.

The running time of the algorithm is bounded above a polynomial function of $|\mathcal{T}|$ and $\log w(C)$.

Remark 7.4. In Theorems 7.4 and 7.5 of [11], the surface S was required to be orientable, and so the same restriction is made here and hence in Theorems 1.1 and 1.2. The requirement that S was orientable was necessary for the proofs in [11], which relied on an analysis of the AHT algorithm applied to train tracks, and certain operations in the AHT algorithm can be ruled out under the assumption that S is orientable. Nevertheless, we expect that Theorem 7.3 remains true even when S is non-orientable, and hence we expect that the orientability hypothesis can be dropped from Theorems 1.1 and 1.2.

Proof of Theorem 1.1. We are given a triangulation \mathcal{T} for a compact orientable surface S , and normal co-ordinates for two properly embedded 1-manifolds C_1 and C_2 . We wish to compute the geometric intersection number $i(C_1, C_2)$.

We may assume that C_2 is connected, for the following reason. Using COMPONENTS OF STANDARD 1-MANIFOLD (Theorem 3.4), we may compute a list of components of C_2 . It is easy to see that $i(C_1, C_2)$ is equal to the sum of $i(C_1, C'_2)$ over all components C'_2 of C_2 . Hence, if we can compute $i(C_1, C'_2)$ for each such component, then we can compute $i(C_1, C_2)$.

We may assume that S is not a sphere, disc or annulus, since geometric intersection number is straightforward to compute in these cases. Hence, we may reduce \mathcal{T} to a 1-vertex triangulation or ideal triangulation using the algorithm in Proposition 10.3 in [10] for example. We sketch the procedure now.

First suppose that S is closed. We will show how to reduce the number of vertices of \mathcal{T} to 1. If there is an edge with the same triangle on both sides, then one endpoint of this edge is a degree one vertex v . The remaining edge has distinct triangles on each side. Hence, a flip may be performed along the latter edge, which increases the degree of v to 2. Then a further flip may be performed, increasing the degree of v to 3. A 3-1 Pachner move may then be performed, which removes the three triangles around v and replaces them with a single triangle. This reduces the number of vertices. Hence, we may assume that each edge has distinct triangles on each side. Now in any triangulation of the closed surface S , the average vertex degree is at most $\max\{6, 6 - 6\chi(S)\}$, and so there is always a vertex v with degree at most this quantity. If v is not the unique vertex of the triangulation, there is an edge emanating from v that ends on a distinct vertex. If every edge emanating from v has this form, then we can perform a flip along any of these edges to reduce the degree of v by 1. On the other hand, if there is some edge emanating from v that ends on v , then we can find a triangle with two vertices equal to v and one vertex not equal to v . If we perform a flip along the v - v edge, then we reduce the degree of v by at least 1. So we can reduce the degree of v to at most 3, and it can then be removed. Thus, we eventually obtain a 1-vertex triangulation. The number of moves we used was at most a polynomial function of $\chi(S)$ and $|\mathcal{T}|$.

A similar procedure can be used when S has non-empty boundary. We can attach triangulated discs to ∂S to form a triangulation of a closed surface \bar{S} , thereby increasing the number of triangles by at most a factor of 2. Pick a vertex in each of these discs. We can then remove all remaining vertices using flips and 3-1 Pachner moves, just as in the above procedure. The result is a triangulation of \bar{S} with one vertex in each of the new discs and no others. Now remove these vertices to form

an ideal triangulation of S . Again, the number of moves we used was at most a polynomial function of $\chi(S)$ and $|\mathcal{T}|$.

During this process, we can keep track of the 1-manifolds C_1 and C_2 . In the case of flips, this is straightforward. But in the case of a 3-1 Pachner move, non-normal arcs of C_1 or C_2 may be created in the new triangle. However, these can easily be normalised using `NORMALISE STANDARD 1-MANIFOLD IN A TRIANGULATION`. Each flip increases $w(C_1)$ and $w(C_2)$ by at most a factor of 2. Hence, $\log(w(C_1))$ and $\log(w(C_2))$ increase by at most a polynomial function of $\chi(S)$ and $|\mathcal{T}|$.

Denote this new 1-vertex triangulation or ideal triangulation by \mathcal{T}' .

We now apply `SIMPLIFY NORMAL 1-MANIFOLD` from Theorem 7.3 to \mathcal{T}' and the 1-manifold C_2 . This converts \mathcal{T}' to a new 1-vertex or ideal triangulation \mathcal{T}'' , and it converts C_2 to a properly embedded normal 1-manifold that intersects each edge of \mathcal{T}'' at most twice. (Here, we have used that C_2 is connected.) This triangulation \mathcal{T}'' is obtained from \mathcal{T}' by a sequence of flips and twists. At each step, we can keep track of the 1-manifold C_1 , until it is a normal 1-manifold in \mathcal{T}'' . We can subdivide \mathcal{T}'' to form a triangulation \mathcal{T}''' of S that contains C_2 as a subcomplex, and where $|\mathcal{T}'''| \leq 9|\mathcal{T}''|$. We can also compute C_1 as a normal curve in this refinement. We apply `MINIMAL POSITION FOR 1-MANIFOLDS` from Theorem 6.3, giving a 1-manifold C'_1 that is isotopic to C_1 and in minimal position with respect to C_2 . Their number of intersection points is the required output from our algorithm. \square

Proof of Theorem 1.2. Again we are given a triangulation \mathcal{T} for a compact orientable surface S , and normal co-ordinates for two properly embedded 1-manifolds C_1 and C_2 . We wish to compute whether they are isotopic.

Again we may assume that C_1 and C_2 are connected, for the following reason. Using `COMPONENTS OF STANDARD 1-MANIFOLD` (Theorem 3.4), we may compute a list of components of C_1 and C_2 . These fall into at most $O(|\chi(S)|)$ topological parallelism classes. We can determine these parallelism classes, since two components C'_1 and C''_1 of C_1 , say, are topologically parallel if and only if they cobound an annulus or square. We can determine this by using `CUT ALONG STANDARD 1-MANIFOLD` (Theorem 4.6) to compute the reduced handle structure on $S \setminus (C'_1 \cup C''_1)$ and then determining the topological type of each component of this surface, as well as the information about which parts of the boundary came from C'_1 and C''_1 . So we now take one representative C'_1 from a topological parallelism class of the components of C_1 . We also run through each topological parallelism class of components C'_2 of C_2 in turn. Suppose that we can determine whether C'_1 and C'_2 are isotopic. If C'_1 is not isotopic to any component of C_2 , then the algorithm terminates with a negative answer. On the other hand, if C'_1 is isotopic to some component C'_2 of C_2 , then we determine the number of parallel copies of these components in C_1 and C_2 . If these numbers are different, we again terminate with a negative answer. On the other hand, if there are the same number of components, then we remove these components from C_1 and C_2 and repeat.

So assume that C_1 and C_2 are connected. As above, we modify \mathcal{T} to a 1-vertex triangulation or ideal triangulation \mathcal{T}' . We then apply `SIMPLIFY NORMAL 1-MANIFOLD` from Theorem 7.3 to \mathcal{T}' and the 1-manifold C_2 , and then we subdivide to give a triangulation \mathcal{T}''' in which C_2 is simplicial. We also keep track of an isotopic copy of C_1 in \mathcal{T}''' . We apply `MINIMAL POSITION FOR 1-MANIFOLDS` from Theorem 6.3, giving a 1-manifold (also called C_1) that is isotopic to C_1 and in minimal position with respect to C_2 . If C_1 and C_2 intersect in minimal position,

they are not isotopic. On the other hand, if they do not intersect, then we can use CUT ALONG STANDARD 1-MANIFOLD to compute the reduced handle structure on $S \setminus (C_1 \cup C_2)$ and then determine whether any component is an annulus or square interpolating between C_1 and C_2 . There is such a component if and only if C_1 and C_2 are isotopic. \square

8. SURFACES CONTAINING A PATTERN

In this section, we will consider a closed surface S containing a pattern P , which is defined as follows.

Definition 8.1. A *pattern* P in a closed surface S is a disjoint union of simple closed curves and embedded trivalent graphs.

The reason for considering such a situation is that the 3-manifolds appearing in a hierarchy naturally inherit a pattern in their boundary [9, 12]. The analysis that we present here will be useful in future work of the author on hierarchies.

We now generalise many of the definitions from earlier in the paper to the setting of patterns.

We will consider 1-manifolds C properly embedded in a closed surface S with a pattern P . After a small isotopy of C , we may assume that C is disjoint from the vertices of P and intersects its edges and simple closed curves transversely. Henceforth, we will always assume this to be the case.

Definition 8.2. Let C be a properly embedded 1-manifold in the closed surface S , and let P be a pattern in S . Then a *bigon disc* is a disc D embedded in S such that ∂D is the concatenation of arcs α and β where $\alpha = D \cap C$ and $\beta = D \cap P$ is disjoint from the vertices of P . A *trigon disc* is a disc D embedded in S such that ∂D is the concatenation of arcs α and β , where $\alpha = D \cap C$ and $\beta = D \cap P$ intersects a single vertex of P in its interior. (See Figure 17.)

Definition 8.3. We say that C and P are in *locally minimal position* if they have no bigon or trigon disc.

By analogy with the simplification number and bigon number, we present the following definition.

Definition 8.4. The *bi-trigon number* $\text{bi-trigon}(C, P)$ of C and P is the number of trigon discs and plus twice the number of bigon discs.

We now consider what the right analogue of an isotopy region should be.

Definition 8.5. Let D be a bigon or trigon for C and P . Then the *associated isotopy region* is a disc R embedded in S that is the union of D and a subset of the form $I \times [-1, 1]$ for a closed bounded interval I , with the following properties:

- (1) the intersection between D and $I \times [-1, 1]$ is the arc $\partial D \setminus C$; it is also a component of $\partial I \times [-1, 1]$;
- (2) the intersection between C and $I \times [-1, 1]$ is $I \times \{-1, 1\}$;
- (3) the intersection between the vertices of P and $I \times [-1, 1]$ lies on $I \times \{0\}$;
- (4) each component of intersection between $I \times [-1, 1]$ and an edge of P is of the form $I' \times \{0\}$ for a closed bounded interval $I' \subseteq I$, or $\{*\} \times [-1, 0]$, or $\{*\} \times [0, 1]$ or $\{*\} \times [-1, 1]$, for a point $* \in I$, or possibly a union of such arcs.

Furthermore, we require that $I \times [-1, 1]$ is maximal with these properties, up to isotopy preserving P and C . (See Figure 17.) Let C' be the 1-manifold obtained from C by isotoping C across R . We say that D and R are *self-returning* if the component of $C' \setminus P$ containing $\partial R \setminus C$ is topologically parallel to $D \cap C$ and all the arcs of $C \setminus P$ between them are also parallel to $D \cap C$.

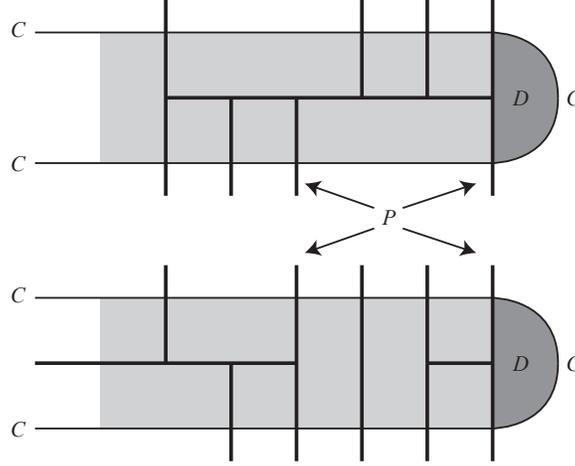


FIGURE 17. Two possibilities for the isotopy region associated with the trigon D .

The main property of this definition is that isotopy regions are extendable, in the following sense. Suppose that R is a region satisfying the above properties, but without the condition of maximality. Suppose that we were to isotope C across R creating a 1-manifold C' . Suppose that C' has a bigon or trigon disc D' that intersects $C' \setminus C$ and with interior disjoint from R and C . Then a small regular neighbourhood of $R \cup D'$ in $S \setminus C$ would also satisfy the definition above (possibly without maximality). (See Figure 18.) Hence, we deduce that when R is indeed maximal, then there is no such bigon or trigon disc D' . This is a key step in the proof of the following result, which is an analogue of Lemmas 5.4 and 6.11.

Lemma 8.6. *Let C be a 1-manifold properly embedded in a closed surface S containing a pattern P . Let D be a bigon or trigon for C and P . Let C' be the 1-manifold obtained by isotoping C across the isotopy region associated with D . Then $\text{bi-trigon}(C', P) \leq \text{bi-trigon}(C, P)$ and $|C' \cap P| < |C \cap P|$. Furthermore, if $\text{bi-trigon}(C', P) = \text{bi-trigon}(C, P)$, then there is an arc of $C' \setminus P$ that is topologically parallel to $D \cap C$, and all the components of $C \setminus P$ between these two arcs are topologically parallel to $D \cap C$.*

Proof. Just as in the proof of Lemma 5.4, we define a function

$$f: \{\text{bigons and trigons for } C'\} \rightarrow \mathbb{P}\{\text{bigons and trigons for } C\},$$

by declaring that for each bigon or trigon D' for C' , $f(D')$ is the set of bigon and trigon discs for C that lie in D' . There are obvious analogues of Properties (1) - (4) in the proof of Lemma 5.4 which also hold here:

- (1) $f(D') \neq \emptyset$ for every bigon or trigon disc D' for C' ;

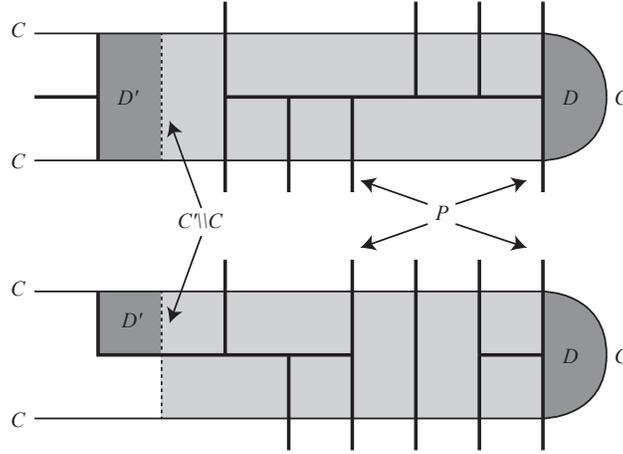


FIGURE 18. Shown are two regions R that satisfy the conditions of an isotopy region, but possibly without maximality. Shown also is $C' \setminus C$, where C' is obtained by isotoping C across R . A bigon or trigon D' for C' is also shown. In each case, a small regular neighbourhood of $R \cup D'$ in $S \setminus C$ also satisfies the definition of an isotopy region, possibly without maximality.

- (2) $f(D'_1) \cap f(D'_2) = \emptyset$ for distinct bigon or trigon discs D'_1 and D'_2 for C' ;
- (3) if D' is a bigon disc for C' , then $f(D')$ consists of bigon discs;
- (4) if $f(D') = \{D_1\}$ for a single disc D_1 that is a bigon exactly when D' is a bigon, then $C \cap \text{int}(D')$ consists of a (possibly empty) collection of parallel arcs, all parallel to $D' \cap C'$.

They quickly imply the lemma.

The only property that is not immediate is (1): $f(D') \neq \emptyset$ for every bigon or trigon disc D' for C' . We now sketch how this is proved. For a bigon or trigon disc D' for C' that is also a bigon or trigon disc for C , this is immediate. So suppose that D' is a bigon or trigon for C' that is not a bigon or trigon for C . Suppose also, for contradiction, that D' does not contain a bigon or trigon of C . Let α' be the arc $\partial D' \cap C'$. If α' is a subset of C , then it must have been the case that the interior of D' contained an arc of intersection with C , and an outermost such arc gives a bigon or trigon for C , which we are assuming is not the case. So suppose that α' is not a subset of C . It must then be incident to the arc $\partial R \setminus C$. If the interior of R is disjoint from the interior of D' , then, as observed above, a small regular neighbourhood of $R \cup D'$ in $S \setminus C$ would also satisfy the definition of an isotopy region for C (possibly apart from maximality), contradicting the fact that R is maximal. Hence, R must enter the interior of D' at the arc $\partial R \setminus C$. But this then implies that D' contains a bigon or trigon of C . \square

Definition 8.7. Let D be a bigon or trigon disc in the closed surface S , and let R be the associated isotopy region. Then an *enlargement* of D is a disc D_+ in S such that

- (1) D_+ contains R ;
- (2) $D_+ \cap C$ consists of arcs topologically parallel to $R \cap C$;

- (3) one of these arcs is $\partial D_+ \cap C$;
- (4) the remainder of ∂D_+ contains $\partial R \setminus C$ and has no points of intersection with P other than $(\partial R \setminus C) \cap P$.

The *maximal enlargement* of D is an enlargement that is not contained within a bigger enlargement, up to ambient isotopy preserving P and C . The *associated isotopy* of C removes the arcs $D_+ \cap C$ and replaces them with parallel copies of $\partial R \setminus C$.

The following is an analogue of Lemma 5.7.

Lemma 8.8. *Let D be a bigon or trigon for C and P . Let C' be the 1-manifold that is obtained from C by an isotopy across a maximal enlargement of the isotopy region associated with D . Then $\text{bi-trigon}(C', P) \leq \text{bi-trigon}(C, P)$.*

The following is analogue of Lemma 5.8, with essentially the same proof.

Lemma 8.9. *Let D be a bigon or trigon for C and P that is not self-returning. Let $e(P)$ be the number of edges of P . Then after performing $O(|\chi(S)| + e(P) + \text{bi-trigon}(C, P))$ isotopies across maximal enlargements of isotopy regions containing D , the result is a 1-manifold C' with $\text{bi-trigon}(C', P) < \text{bi-trigon}(C, P)$.*

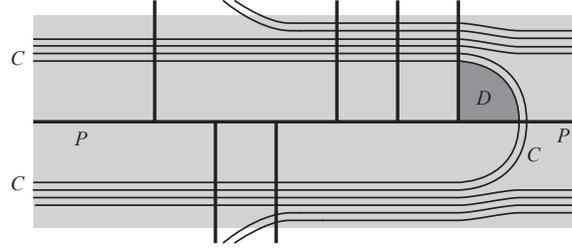
There are also analogues of Lemmas 5.10 and 5.11. In the case where a trigon or bigon disc D is self-returning and where its isotopy region R has no arc of $(R \cap C) \setminus P$ parallel to $D \cap C$, we have the following analogue of a Reeb region.

Definition 8.10. Let S be a closed surface containing a pattern P . Let C be a 1-manifold properly embedded in S . A *generalised Reeb region* is a subsurface A of S homeomorphic to an annulus or Möbius band $S^1 \widetilde{\times} [-1, 1]$ and with the following properties:

- (1) the intersection between the vertices of P and $S^1 \widetilde{\times} [-1, 1]$ lies on $S^1 \widetilde{\times} \{0\}$;
- (2) each component of intersection between $S^1 \widetilde{\times} [-1, 1]$ and an edge of P is of the form $I' \widetilde{\times} \{0\}$ for a closed bounded interval $I' \subseteq S^1$, or $\{*\} \widetilde{\times} [-1, 0]$, or $\{*\} \widetilde{\times} [0, 1]$ or $\{*\} \widetilde{\times} [-1, 1]$, for a point $*$ $\in S^1$, or possibly a union of such arcs.
- (3) A contains a bigon or trigon D in its interior;
- (4) the intersection between A and C consists of a union of essential arcs in A , all of which are topologically parallel;
- (5) each such arcs is a concatenation of three arcs $\alpha_-, \alpha_0, \alpha_+$;
- (6) α_- starts at a point of ∂A and then runs parallel to ∂A in some direction, possibly winding several times around that component but winding at least once around that component;
- (7) α_0 is topologically parallel to $D \cap C$;
- (8) α_+ starts at a point of ∂A opposite $\partial A \cap \alpha_-$, and winds around ∂A the same number of times as α_- and in the same direction.

A *generalised Reeb isotopy* is supported in a generalised Reeb region as above. It replaces $A \cap C$ with essential arcs in A of the form $\{*\} \widetilde{\times} [-1, 1]$ for a point $*$ $\in S^1$. A Reeb region A is *maximal* if any other Reeb region containing it is isotopic to A via an isotopy preserving C and P throughout. (See Figure 19.)

Theorem 8.11 (LOCALLY MINIMAL POSITION FOR 1-MANIFOLD AND PATTERN). *There is an algorithm that takes, as its input, the following:*

FIGURE 19. A generalised Reeb region associated with a trigon D .

- (1) a triangulation \mathcal{T} or handle structure \mathcal{H} of a closed surface S ;
- (2) a pattern P properly embedded in S that either is a subset of the 1-skeleton of \mathcal{T} or that intersects \mathcal{H} in the following way: it is disjoint from 0-handles; it intersects each 1-handle in the empty set or in the co-core of the handle; it intersects each 2-handle in the empty set, or in a properly embedded arc, or in a graph that consists of a single vertex of P with three arcs running to the boundary of the 2-handle;
- (3) a standard 1-manifold C ;

and provides as its output a normal 1-manifold \bar{C} that is obtained from C by an isotopy, plus possibly removing some components that bound discs. This 1-manifold \bar{C} is in locally minimal position with respect to P and has weight that is bounded above by $w(C)(1 + 2|\mathcal{T}|)$ or $w(C)(1 + 2\|\mathcal{H}\|)$. The running time is bounded above by a polynomial function of $|\mathcal{T}|$ or $\|\mathcal{H}\|$, $\text{simp}(C)$ and $\log w(C)$.

Proof. This follows the same format as the proof of Theorems 2.15 and 6.3.

First we search for a generalised Reeb region, arising from a bigon or trigon, and we perform the associated isotopy, creating a 1-manifold C' . If C' has a bigon or trigon in the generalised Reeb region, denote it by D' . If there is no such disc, then the bi-trigon number has gone down, and we return to the start. If there was no generalised Reeb region, then we pick any bigon or trigon D' and set $C' = C$. We then perform the isotopy of C' across the maximal enlargement of the isotopy region associated with D' . If this reduces the bi-trigon number, then we return to the start of the loop. On the other hand, if the bi-trigon number is unchanged, then the new curve C'' has an arc of $C'' \setminus P$ that runs parallel to $D' \cap C'$. An outermost such arc bounds a bigon or trigon. We repeat the above procedure with this disc. We continue until the bi-trigon number decreases, at which point we return to the start of the loop. This continues until there are no more bigons or trigons, at which point the resulting 1-manifold C'' is in locally minimal position with respect to P .

We claim that C'' has weight at most $w(C)(1 + 2|\mathcal{T}|)$ or $w(C)(1 + 2\|\mathcal{H}\|)$. This is proved in the same way as in the argument for Theorem 6.3. Each isotopy performed above replaces sub-arcs of C with new arcs. Each of the new arcs has weight at most $2|\mathcal{T}|$ or $2\|\mathcal{H}\|$ and it intersects P strictly fewer times than the original arc. The inequalities in the proof of Theorem 6.3 then give that the resulting 1-manifold C'' has weight at most $w(C) + 2|C \cap P||\mathcal{T}|$ or $w(C) + 2|C \cap P|\|\mathcal{H}\|$, and this is at most the required bound.

Finally, we normalise C'' using NORMALISE A STANDARD 1-MANIFOLD IN A HANDLE STRUCTURE (Theorem 2.15) or NORMALISE STANDARD 1-MANIFOLD IN A TRIANGULATION (Theorem 1.3). This takes polynomial time, and the resulting

normal 1-manifold \overline{C} has weight at most that of C'' and is still in locally minimal position with respect to P . \square

REFERENCES

- [1] Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Trans. Amer. Math. Soc.*, 358(9):3821–3850, 2006. [6, 12]
- [2] Mark C. Bell. Simplifying triangulations. arXiv:1604.04314, 2016. [4, 39]
- [3] Mark C. Bell and Richard C. H. Webb. Applications of fast triangulation simplification. arXiv:1605.03514, 2016. [2, 4, 6, 39]
- [4] Hsien-Chih Chang and Arnaud de Mesmay. Tightening curves on surfaces monotonically with applications. *ACM Trans. Algorithms*, 18(4):Art. 36, 32, 2022. [2, 5, 6]
- [5] Vincent Despré and Francis Lazarus. Computing the geometric intersection number of curves. In *33rd International Symposium on Computational Geometry*, volume 77 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 35, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2017. [2, 5, 6]
- [6] Loïc Dubois. Making Multicurves Cross Minimally on Surfaces. In *32nd Annual European Symposium on Algorithms (ESA 2024)*, volume 308 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 50:1–50:15, 2024. [2, 5, 6]
- [7] Ivan Dynnikov. Counting intersections of normal curves. *J. Algebra*, 607:181–231, 2022. [2, 5, 6]
- [8] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012. [31]
- [9] W. Haken. Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I. *Math. Z.*, 80:89–120, 1962. [6, 42]
- [10] Marc Lackenby. The efficient certification of knottedness and Thurston norm. *Adv. Math.*, 387:Paper No. 107796, 142, 2021. [12, 15, 40]
- [11] Marc Lackenby and Mehdi Yazdi. Bounds for the number of moves between pants decompositions, and between triangulations. arXiv:2401.14233, 2024. [6, 39, 40]
- [12] Sergei Matveev. *Algorithmic topology and classification of 3-manifolds*, volume 9 of *Algorithms and Computation in Mathematics*. Springer, Berlin, second edition, 2007. [6, 42]
- [13] Marcus Schaefer, Eric Sedgwick, and Daniel Štefankovič. Algorithms for normal curves and surfaces. In *Computing and combinatorics*, volume 2387 of *Lecture Notes in Comput. Sci.*, pages 370–380. Springer, Berlin, 2002. [2, 4]
- [14] Marcus Schaefer, Eric Sedgwick, and Daniel Štefankovič. Computing Dehn twists and geometric intersection numbers in polynomial time. In *Proceedings of the 20th Annual Canadian Conference on Computational Geometry, CCCG 2008*, 01 2008. [2, 4, 5, 6]
- [15] Saul Schleimer. Polynomial-time word problems. *Comment. Math. Helv.*, 83(4):741–765, 2008. [2]

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD,
WOODSTOCK ROAD, OXFORD OX2 6GG, UNITED KINGDOM